

Operationalizing climate targets under learning: An application of cost-risk analysis

Delf Neubersch · Hermann Held · Alexander Otto

Received: 29 June 2013 / Accepted: 4 August 2014

© The Author(s) 2014. This article is published with open access at Springerlink.com

Abstract Cost-Effectiveness Analysis (CEA) determines climate policies that reach a given climate target at minimum welfare losses. However, when applied to temperature targets under climate sensitivity uncertainty, decision-makers might be confronted with normatively unappealing negative expected values of future climate information or even infeasible solutions. To tackle these issues, Cost-Risk Analysis (CRA), that trades-off the costs for mitigating climate change against the risk of exceeding climate targets, has been proposed as an extension of CEA under uncertainty. Here we build on this proposition and develop an axiomatically sound CRA for the context of uncertainty and future learning. The main contributions of this paper are: (i) we show, that a risk-penalty function has to be non-concave to avoid counter-intuitive preferences, (ii) we introduce a universally applicable calibration of the cost-risk trade-off, and (iii) we implement the first application of CRA to a numerical integrated assessment model. We find that for a 2°-target in combination with a 66 % compliance level, the expected value of information in 2015 vs. 2075 is between 0.15 % and 0.66 % of consumption every year, and can reduce expected mitigation costs by about one third. (iv) Finally, we find that the relative importance of the economic over the risk-related contribution increases with the target probability of compliance.

Electronic supplementary material The online version of this article (doi:10.1007/s10584-014-1223-z) contains supplementary material, which is available to authorized users.

D. Neubersch · H. Held (✉)

Research Unit Sustainability and Global Change, Centre for Earth System Research and Sustainability, University of Hamburg – KlimaCampus, Grindelberg 5, 20144 Hamburg, Germany
e-mail: hermann.held@zmaw.de

D. Neubersch

International Max Planck Research School on Earth System Modelling, Grindelberg 53, 20146 Hamburg, Germany

A. Otto

Environmental Change Institute, University of Oxford, South Parks Road, Oxford, OX1 3QY, UK

Keywords Cost-risk analysis · Value of information · Climate target · Climate economics · Optimal climate policy · Integrated assessment · Uncertainty

1 Introduction

Integrated Assessment Models (IAMs) are formal representations of the interconnected socio-economic and climate systems. IAMs are tools for investigating the interactions between human activities and the environment and implications of anthropogenic climate change. A subset of IAMs (e.g. DICE (Nordhaus 2008), WITCH (Bosetti et al. 2006), REMIND (Luderer et al. 2013)) apply normative decision theory to determine (welfare) optimal decisions to avoid dangerous anthropogenic interference with the climate system, as called for by the United Nations Framework Convention on Climate Change (UNFCCC).

A straight forward application of overall welfare optimization is a Cost-Benefit Analysis (CBA) (e.g. Nordhaus 2008) that weighs the welfare losses incurred by mitigation action against the welfare benefits from avoiding climate change induced damage. However, a series of challenges complicate the application of CBA to the climate problem including a difficult comparison of types of damage as well as the existence of large uncertainties (IPCC AR5 WG-III Chapter 3: Kolstad et al. (2014); in detail see e.g. Stern (2013) and Nelson (2013)). Hence, according to part of the scientific community, more research into impact models is required to address the issue of finding an approximate aggregate damage function (even on a sectoral level for a single region).

The time that is at disposal for reaching a global agreement on binding emission reduction targets (or cumulative emissions) might be shorter than needed to develop adequate CBA-based IAMs, if the possibility of limiting the rise of global mean temperature to values around 2 °C compared to the pre-industrial value shall be preserved (UNFCCC 2011). In that sense decision-aiding cannot wait for the advent of better impact models.

One way of dealing with the shortcomings described above is to switch focus to scenario frameworks and notions of robust decision making (Lempert et al. 2006), thereby abandoning the appealing notion of optimality. Part of the community sees climate targets as a viable alternative (IPCC AR5 WG-III Chapter 6: Clarke et al. (2014)) in the light of deep (“Knightian”) uncertainty and preferences in-line with the precautionary principle (Iverson and Perrings 2012). Obviously, the choice of a specific climate target is influenced in part by intuitive expectations of potential climate-induced damage. Yet the construction of a climate target avoids having to formalize the totality of all global warming impacts and their respective uncertainties – a task that would pose so far unresolved conceptual and formal problems when having to be expressed as Knightian uncertainty. Instead, this so far unresolved formal task is often replaced by a climate target (or “guard rail”) that we interpret as an informal convolution of sparse impact information and decision preferences under Knightian uncertainty, as long as no more robust method is available. Thus, Cost-Effectiveness Analysis (CEA) of (exogenous) climate targets emerged as a method that avoids the difficulties of defining a climate change damage function and tries to separate the evaluation problem from the policy analysis (Patt 1999; Luderer et al. 2012). If policy makers agree on a climate guardrail, the most cost-effective policy can be determined. If the costs (monetary or otherwise) of a policy intervention are found to be “sufficiently small”, this could open the venue for societal action (Patt 1999).

However, the situation is further complicated by considering (and anticipating) future changes in our knowledge about the Earth system response to emissions by observations

(Kelly and Kolstad 1999), the assimilation of paleo information (Schneider von Deimling et al. 2006; Lorenz et al. 2009) and improvements in theory and modeling of sub-scale processes.

The possibility of receiving new information has been shown to change the resulting optimal policy substantially (e.g. O'Neill and Melnikov 2008; Webster et al. 2008). Allen and Frame (2007) even argue that a substantial fraction of climate response uncertainty could be compensated by adopting an adaptive climate policy that reacts to learning after 2050 in order to achieve ambitious long-term climate targets. Including the possibility of future learning into an IAM enables calculating the economic value of new information about the climate response (e.g. Nordhaus and Popp 1997; Lorenz et al. 2012b). This is also critical in assessing investment opportunities into technologies and projects aimed at providing better climate information (e.g. projects such as GEO-BENE (EU-FP6) (Fritz et al. 2008) and EuroGEOSS (EU-FP7) (Pearlman et al. 2011) try to assess the return on investment for projects that mainly gather and distribute information).

No temperature target can be achieved with absolute certainty given the current specifics of an uncertain temperature response (IPCC 2013). To keep climate targets operational, an associated probability also needs to be specified. Once such a target is formalized, CEA can be applied to temperature targets by using Chance Constrained Programming (CCP) (Held et al. 2009). However, CCP is faced with two problems when considering the advent of new information, both of which are discussed in Schmidt et al. (2011) and references therein. We give a short overview of the problems to clarify the importance of our contribution.

Firstly, the prescription of CCP might lead to infeasible solutions after learning. After having resolved uncertainty, for some of the possible states of the world the decision maker will not be able to comply with the guard rail.¹ Secondly, as first indicated by Blau (1974), within a CCP decision framework anticipated learning could be allocated a negative expected value and would be rejected accordingly if optional. The historical discussion is thoroughly documented in Schmidt et al. (2011) and the interested reader can also refer to LaValle (1986) and Wakker (1988) for a theoretical discussion of CCP and negative value of information. Such a negative value is normatively unappealing and does not enable the above mentioned necessity of gauging the value of further climate observations.

In resolving both conceptual problems of CEA/CCP at once, Schmidt et al. (2011) proposed "Cost-Risk Analysis" (CRA) in which a trade-off is made between the risk of overshooting climate target and the economic cost of switching to renewable energy. Thereby, they provide a normative decision criterion that does not explicitly require a climate damage function, but builds on a consensus-based climate target. Thus, CRA preserves the structural equivalence to the expected utility framework (see also Eq. 3), which in turn ensures non-negative expected values of information (Gollier 2004).

CRA can be interpreted as a consequence of Morgan et al. (1992), who suggest to model a complex system the more parsimonious, the sparser the level of the modeler's understanding of the system. Accordingly, CRA models the preference order of the community supporting a temperature target in view of poorly accepted damage functions (Schmidt et al. 2011). Its main service lies in harmonizing CEA and anticipated future learning. It necessarily requires a climate target as an input and hence, by definition, does not represent a means to put a climate target on a more objective basis (which only ongoing impact-oriented research can do). Some further remarks about the limitations of CRA might be in order. Firstly, the nature

¹Part of the literature "bypasses" this conceptual difficulty imposed on CEA by artificially truncating the upper tail of the distribution of climate sensitivity.

of cost-risk trade-off is subject to ambiguity. Apart from an argument of parsimony, we stress that the model that we use is the most conservative one (in using the least convex risk measure). Secondly, CRA silently assumes that any global warming impact information is encapsulated in the temperature target. It remains silent about how decisions might change once impact research achieved major breakthroughs.

This study addresses the fraction within the climate economics community that in fact currently focuses on CEA as opposed to CBA. It completes the axiomatic framework of CRA in highlighting and eliminating gaps in the definition and calibration of the risk function, and presents the first IAM-based operationalization of CRA to the mitigation of climate change.

Our study makes four major contributions: (i) it shows that to avoid the sacrifice of mitigation efforts in certain learning cases one needs to introduce a risk metric that is non-concave in temperature. The boundary case, i. e. a linear function, is equivalent to the concept of discounted expected degree years. Degree years (integrating the exceedance of the temperature target over time) have been introduced by Schneider and Mastrandrea (2005) as a risk metric. (ii) This study presents a universally applicable calibration procedure of the preference for the trade-off between economic utility and climate risk. As one possible pragmatic interpretation of this procedure we refer to the preference implicit in the 17th Conference of the Parties (COP) to the UNFCCC (UNFCCC 2011) of staying below 2 °C with a (likely) 66 % probability. (iii) We apply the framework to an IAM for the first time and conduct a sensitivity study on how normative parameters, timing of learning and calibration affect the expected value of information. While, intuitively, stricter targets and earlier learning increase the value of information in most cases, we also showcase circumstances under which stricter targets can actually decrease the value of information. (iv) Finally, we show that the value of information consists of economic and risk-related value and that the composition depends on the climate target.

For demonstration and further development of CRA we employ the IAM MIND-L (Model of Investment and Technological Development including Learning) (Edenhofer et al. 2005; Held et al. 2009; Lorenz et al. 2012b). The remainder of the paper is structured as follows: the theoretical framework of CRA is detailed in Section 2 and the application is treated in Section 3. We conclude and discuss our results in Section 4. Online Resource 1 includes background calculations and further substantiating figures.

2 The cost-risk framework

At the heart of CRA lies a trade-off between the cost of mitigation and the excess risk introduced by temperatures above a guard rail. This section introduces the formalism required for the discussion, building upon Schmidt et al. (2011) and Schneider and Mastrandrea (2005) to make three points: (i) the standard formulation of the CRA shown in Schmidt et al. (2011) leads to the possibility of counter-intuitive sacrificing of the climate after a learning event, (ii) using degree years as a risk measure prevents the sacrifice and furthermore preserves all envisioned benefits of a CRA framework, and (iii) the trade-off parameter can be calibrated using targets discussed in politics.

2.1 Sacrificing the climate

We start with the formulation of the minimization problem for a CRA with mitigation costs $C(E)$ and a risk function $R(E)$ both depending on the emission strategy E that is chosen

by the decision maker. In order to motivate why a further development of CRA compared to the standard scheme of Schmidt et al. (2011) is necessary and for the sake of conceptual clarity in this sub-section we restrict the discussion to a static picture in terms of cumulative emissions E as crucial control variable (Lorenz et al. 2012). Using β as the trade-off parameter the minimization problem reads:

$$\min_E \{C(E) + \beta R(E)\}. \tag{1}$$

For the function $R(E)$ Schmidt et al. (2011) proposed the probability of violation as the risk measure:

$$R(E) = \int f(\theta) \Theta(T_{\max}(E, \theta) - T_g) d\theta. \tag{2}$$

Here Θ is the Heaviside function, θ is the uncertain parameter with the probability distribution $f(\theta)$.² The guard rail is formulated in terms of temperature T_g . The maximum temperature given a certain strategy and a realization of the uncertain parameter is T_{\max} . Note that the risk function preserves a crucial property of CEA that would not demand any reduction or restriction of welfare due to temperature values below T_g .

A learning event could place the decision maker in a position where he has to weigh the cost of reducing the temperature below the guard rail and the benefit from doing so. It is easy to see that if the cost is higher than the risk described by $\beta R(E)$, the optimum is represented by a business-as-usual path (maximal emissions). The ever increasing temperatures do not create an additional penalty and the costs $C(E)$ can be minimized unchecked. An example plot of the sacrifice for the model MIND-L is shown in Online Resource 1.

2.2 A linear risk function

We propose an axiom to exclude the behavior in the case described above, hence adjusting CRA to the preference order of the proponents of climate targets.

Axiom of Sacrifice Inhibition Imagine society, in a static setting, regards a budget of emissions E^* for the world as welfare optimal. If, for some reason, new information arrives that this budget cannot be met because an amount of dE emissions has or will be emitted too much, then the optimal budget goal should be $E^{**} = E^* + dE$, not any larger value.

In a non-convex decision problem, a little change in emissions (or cost of mitigation) can cause the solution to jump to a different regime facilitating a violation of the above axiom. In a convex problem, local optima are avoided, and it is guaranteed that the solution will not jump and therefore the above axiom will not be violated. To this end, the sum of cost and risk function has to be convex in the decision variables to have a single minimum. We assume that the cost function is convex as it is in all prominent models of the climate problem known to us. This, in turn, implies that it is a sufficient condition for the function $R(E)$ to be non-concave (above T_g) for the total sum to be convex in E . Furthermore, if we want to ensure that this would even hold for *any* convex cost function, $R(E)$ necessarily needs to be convex. A linear $R(E)$ represents the boundary case between convex and concave risk functions.

²For this paper we only consider uncertainty in the climate response to an increase in greenhouse gas concentration.

2.3 Discounted expected utility maximization

To be able to conduct a dynamic analysis, the above concept can be interpreted as a linear risk function on temperature at each point in time and each state of the world.³ Integrated over time this is essentially equivalent to Degree Years (Schneider and Mastrandrea 2005). We deviate slightly to ensure that we remain in an expected utility framework by aggregating the risk in the same way as utility is aggregated, i.e. by exponential discounting. We can redefine our risk measure as measuring “expected discounted degree years”. Formulated as a maximization problem we arrive at the discounted utility welfare equation for CRA as follows:

$$\max W = \sum_{t=0}^{t_{\text{end}}} \sum_{s=1}^S p_s \left\{ \underbrace{U(t, s)}_{\text{economic}} - \underbrace{\beta R(T(t, s))}_{\text{risk-related}} \right\} e^{-\delta t}. \quad (3)$$

Here δ is the pure rate of time preference and $U(t, s)$ is the utility from consumption for time t and state of the world s . Each state of the world is associated with a probability p_s . W describes the overall welfare and the risk function $R(T)$ is calculated in each time step and state of the world as:

$$R(T) = \Theta(T - T_g) \cdot (T - T_g) \quad (4)$$

The above form of CRA now only lacks a specification for β .

2.4 Calibration

The parameter β reflects how much the decision maker cares about climate change induced risk. This is a normative decision related to the decision about safety, and could be elicited in principle. However here we highlight a procedure that is universally applicable for any temperature target. As outlined in the introduction, the conceptual clash between CEA and anticipated learning originates from the fact that presently no temperature target can be complied to with certainty, hence any temperature target needs to be specified probabilistically.⁴ If society decides on a temperature target it should also specify a corresponding probability of compliance in order to make the target operational. We propose to choose β such that the solution of CRA is efficiently in compliance with that joint temperature-probability target. Hence the change from CEA with CCP to CRA implies replacing an expected welfare optimization under a probabilistic temperature constraint by an unconstrained expected welfare optimization that is however calibrated such that the probabilistic target is still met. In each case one single output quantity is prescribed (i.e. maximum temperature) for the welfare optimization.

While CRA could be calibrated against any such target, here we would like to demonstrate CRA for a specific normative setting of particular relevance in the context of international climate negotiations. The target of constraining global mean temperature increase to 2 °C above pre-industrial conditions has been discussed prominently in the climate policy debate (Schellnhuber 2010). However, a compliance level has been specified less clearly. The 17th COP to the UNFCCC refer in their decisions to “aggregate emissions

³To go from a linear function $R(E)$ to a linear function in temperature we make use of one of the most prominent results of climate science: cumulative emissions are approximately proportional to peak temperature increase (Allen et al. 2009).

⁴Caused by a presently infinitely-tailed probability density function of climate sensitivity.

pathways consistent with having a likely chance of holding the increase in global average temperature below 2 °C or 1.5 °C above pre-industrial levels” (UNFCCC 2011). We pragmatically use the interpretation of “likely” as implying a probability of observing the target of at least 66 % as recommended in Mastrandrea et al. (2010). We call this probability “safety” throughout the paper. Building on this assumption and the statement from COP17, we formulate the climate target in our analysis as follows:

Calibration target: a chance of at least 66 % to restrict temperature anomalies to a maximum of 2 °C without considering future reduction of uncertainty.

Hereby, we calibrate β in-line with one potential interpretation of the COP’s agreement. For a discussion of why we prefer to conduct the calibration in a scenario without learning see Online Resource 1. The next chapter deals with the application of this framework to MIND-L to calculate the value of information and analyze changes in policy.

3 Application

3.1 Setting

For the numerical analysis, we use the climate-energy-economy model MIND-L in the form presented by Lorenz et al. (2012b). For further information about this Ramsey-type model please consult Online Resource 1. The uncertainty in the model stems from climate sensitivity that is assumed to be distributed in a log-normal form: $\text{pdf}(\text{CS}) = \mathcal{LN}(0.973, 0.4748)$ (Wigley and Raper 2001). A perfect correlation to climate response time scale is assumed, to calculate transient climate response from the climate sensitivity as described in Lorenz et al. (2012b).

We consider perfect learning in the analysis as a good first approximation of the more general case. Please refer to Online Resource 1 for a description of our implementation. Furthermore, we will concentrate on the value of learning and only briefly touch upon the effect of anticipation, which has been studied at length elsewhere (Ulph and Ulph 1997; Lorenz et al. 2012b).

We define three special cases of the welfare optimization problem (similar to Webster et al. 2008; Lorenz et al. 2012b):

- BAU (business-as-usual) is a scenario in which climate change is completely ignored and there are no penalties related to it. This counter-factual reference case is associated with the highest welfare, as the existence of climate change represents a net loss in all other scenarios.
- NOLEARN is a scenario that assumes that no better information about the climate system is found in time and our decisions on how strongly to invest in mitigation have to be made with the information we have today. There is a climate problem and the decision maker includes the cost-risk trade-off but one decision has to be made for all states of the world. The trade-off parameter β is calibrated to reach a specified target.
- LEARN(2015...2075) assumes that at a specific learning point the decision maker learns the true state of the world and can act accordingly. The trade-off parameter is kept at the value which was used for the corresponding NOLEARN case.

We use a constant relative risk aversion utility (welfare) function with a risk aversion of η . Throughout the application of CRA we change four parameters: (i) risk aversion (η), (ii) pure rate of time preference (δ), (iii) calibration target (safety) and (iv) learning point. The base model is characterized by a risk aversion of 2, a time preference of 1 % pa and

a target of 2 °C with 66 % safety. Calculated values are given as percentage changes of the Certainty and Balanced Growth Equivalent (CBGE, see Online Resource 1 and Anthoff and Tol (2009) for details). It represents the initial consumption change that is necessary to create the difference in welfare assuming constant equal growth rate of consumption.

3.2 Comparison to CEA

To put the model in perspective we compare it to classic CEA with CCP. The analysis lets us conclude that CRA and CEA are similar in behavior until 2050 from which point onwards there is an offset by which CEA suggests less mitigation. The mitigation cost (consumption loss from switching to renewable energies) for CEA is 1.3 % CBGE compared to 1.5 % for CRA.⁵ Note that this is only the economic part of the welfare loss from risk and mitigation which is around 3.3 % CBGE. We will come back to this distinction in Section 3.5. See Online Resource 1 for a series of plots comparing CEA and CRA.

We explain the discrepancy in costs as follows. The decision maker in CRA recognizes that crossing the target for longer periods of time is bad and reduces temperatures in the future. In CEA there is no “risk”, just cost effectiveness and temperatures remain at 2 °C with 66 % safety. We see CRA as a much more rational way of staying below a guardrail with a certain probability considering that the decision maker does not re-invest into fossil fuels in 2050 to prevent the temperatures from dropping.

We conduct a CEA with 5 different chance constraints 40 %, 50 %, 60 %, 66 %, 70 % on a 2 °C target. We compare this to a CRA with calibration targets equivalent to the chance constraints in the CEA. We find that over the whole range of safeties the solutions are similar up to year 2050. Hence, although CEA has conceptual problems, the scenarios that have already been generated (IPCC 2014b: WGIII, Chapter 6), provide good proxies of efficient next-decades investments. See the Online Resource 1 for a plot.

3.3 Expected value of perfect information

After analyzing the situation of a CRA without learning, we apply CRA to demonstrate the consistent handling of learning. The first application is to find the expected value of perfect information (EVPI) about the climate response. We show that the value declines with decreasing stringency of the target and later learning times. We measure the value by comparing the welfare in a case of early learning (e.g. LEARN(2015)) with the welfare of a reference case of late learning (LEARN(2050) or LEARN(2075)). The EVPI in 2015 instead of 2075 for example is referred to by EVPI(2015, 2075). The learning scenario assumes that the decision maker anticipates the learning event. Without this assumption, the EVPI is reduced by 2–6 %, depending on the learning point.

We investigate two possibilities of changing the target: the threshold of dangerous anthropogenic interference (or guard rail) and the probability with which to stay below the guard rail (safety) are altered (individually). Figure 1 displays the results and immediately confirms that learning later decreases the EVPI. Panel a of Fig. 1 demonstrates that decreasing the guard rail from 2.25 °C to 2 °C (increasing stringency) only increases the value of information if the information arrives before 2020. For scenarios without learning there is an equivalence between varying guard rails and varying safeties which is shown in panel c of

⁵ MIND-L's costs lie in the mid-range of IAMs as analyzed by Edenhofer et al. (2010).

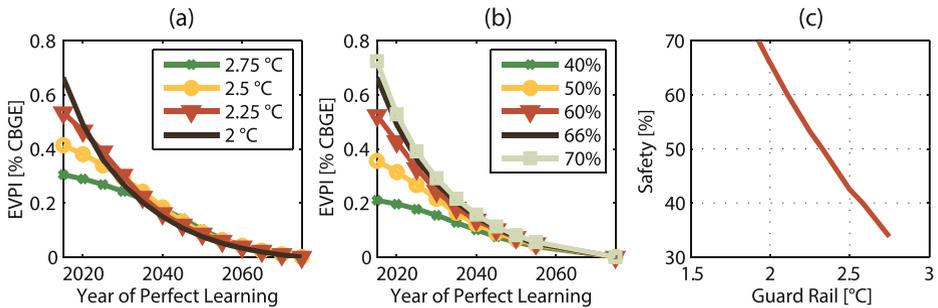


Fig. 1 EVPI depending on learning time point for (a) different target guardrails, all with 66 % safety and (b) different probabilities of abiding by a 2 °C target. (c) shows the equivalence between varying targets with a fixed safety (66 %) and varying safety with a fixed target (2 °C), i.e. keeping a 2.5 °C target with 66 % safety gives the same result as keeping a 2 °C with 42 % safety. This only applies for the NOLEARN case

Fig. 1. If learning is included this is no longer the case because in the state of perfect information temperature paths will gravitate toward the guard rail and if it is higher, this will cost less. For varying safeties, increasing stringency of the target means that the EVPI rises, as shown in panel b of Fig. 1.

For the remainder of the paper we only consider changing the safety target and not the guard rail which we keep at 2 °C, in line with DAI-EU from Schneider and Mastrandrea (2005). We also conducted sensitivity studies regarding power-law exponent of the risk function (from 1 to 4) and find changes of EVPI(2015, 2075) within 6 % only.

3.4 Effect of learning on optimal policy and risk

We are interested in the expected effect of gaining information and how it changes with the timing of learning. Figure 2a shows that the expected discounted degree years increase if information arrives later. This effect vanishes for stricter targets because learning does not influence the paths above the guard rail much, as mitigation is already close to maximum. Figure 2b shows the expected cumulative emissions and underlines that the earlier the information arrives the more can be emitted in expectation and therefore, the cost is reduced.

Interestingly, for 40 % safety, the emissions stay relatively stable although the degree years change drastically. The reason being that going from LEARN(2075) to LEARN(2015) the increase in emissions in states of the world with low temperatures, balances the extra mitigation effort in high temperature worlds. However, only the extra mitigation efforts influence and reduce the risk.

3.5 Origin of EVPI

By having a closer look at where the change in welfare originates from (either from a change in consumption or a change in temperature, see Online Resource 1 for details) we can divide the EVPI into an economic and a risk-related value. These terms refer to the parts of Eq. (3). Referring back to Fig. 2, it is evident that as the safety target decreases the change in the economic part (emissions) plays less of a role than the change in risk (degree years) if late and early learning are compared.

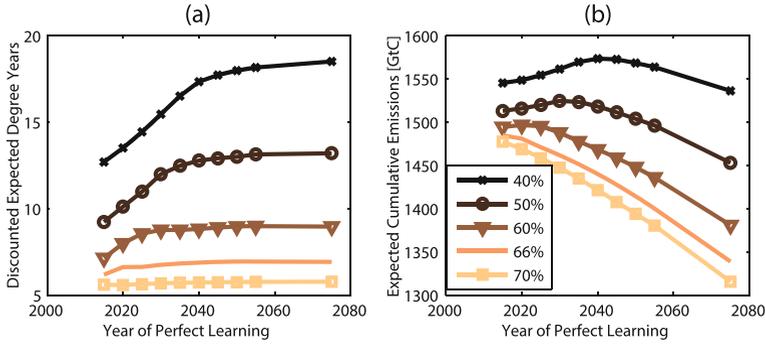


Fig. 2 Effect of learning point on policy and risk, (a) shows the change in expected discounted degree years for different learning points and safety values and (b) shows the expected cumulative emissions in 2200 for different learning points and safety values

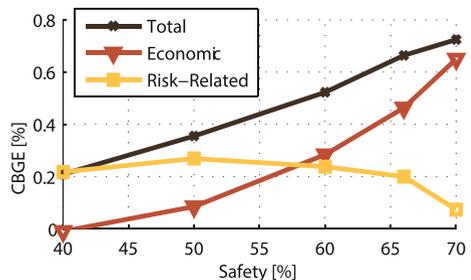
This can be verified by plotting the economic and risk-related parts of EVPI(2015, 2075) for changing safeties as shown in Fig. 3. For low safeties, the entire value comes from reducing the risk, whereas for high safeties, most of the value comes from increasing emissions (in low climate sensitivity states of the world). The reason for this behavior is that for low safeties the base level of emissions is high (see Fig. 2) and therefore, the value gained from increasing emissions is small. For high safeties an inverse argument holds. In summary, for more stringent targets, the EVPI is a real monetary value which decreases the mitigation costs, whereas for less stringent targets it is not.

3.6 Sensitivity of value of learning

Varying the normative parameters and combining them by the Ramsey formula for the social discount rate $d = \eta g + \delta$, allows us to plot the EVPI for different social discount rates. The growth of consumption g is endogenous in the model and lies around 2.2 % pa. Refer to Online Resource 1 for numerical values of η and ρ .

Figure 4a shows the EVPI for different learning points and social discount rates. Figure 4b shows the times at which the value has fallen to half the value of 2015. All of these simulations are compared to LEARN(2075). Changing the reference hypothesis (to, for example, 2050) hardly influences the results (compare Fig. 4c). With increasing social discount rate, the half-time of information decreases from around 15 years to just above 5 years. This is a dramatic decrease and expresses the urgency of action with respect to the 2 °C target.

Fig. 3 EVPI(2015, 2075) divided into economic and risk-related value and plotted against the safety that the model was calibrated against. The stricter the target (higher safety) the more value comes from higher emissions. If the safety is low, then learning will hardly affect the economy and only reduce the risk



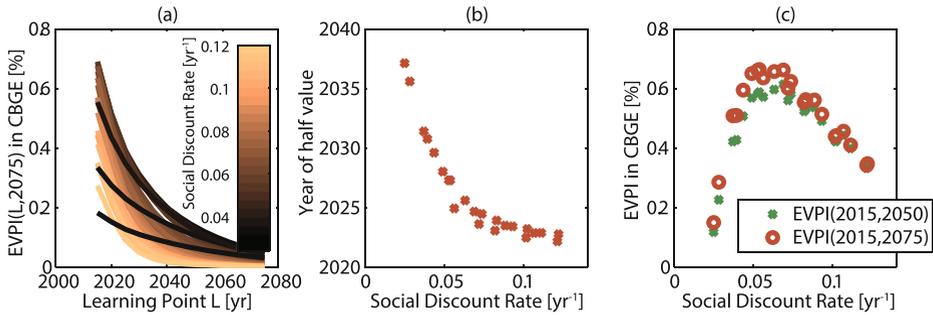


Fig. 4 Panel (a) shows a rapid decline of EVPI as the time of learning is moved into the future. Panel (b) shows the year in which the EVPI is halved in relation to EVPI(2015, 2075). The half-time decreases with increasing social discount rate. Panel (c) shows EVPI for learning in 2015 compared to learning in 2050 (green marks) and 2075 (red circles)

The EVPI for learning early in 2015 compared to 2050 and 2075 is shown in Fig. 4c. The concave behavior is due to the changing discount rate moving the focus of the aggregation along time. We are comparing learning between 2015 and 2075, implying that most of the differences between the scenarios lie in this span. Therefore, very high discount rates devalue the benefit of learning as most weight is on pre-learning welfare. Vice versa, low values of d devalue the benefit as most weight is put on times after learning occurs.

This result is astonishing, because intuitively the origin of the effect is sought for in the dynamic adjustments of decisions to a change in preference, for example, an increase in EVPI as the risk aversion is increased (Eeckhoudt and Godfroid 2000). This effect is very small and dominated by the discounting effect described above. The EVPI reaches its maximum around a d of 5–6 % pa, which is also the value in the base run. Unsurprisingly, the choice of normative parameters is of relevance for the calculation of monetary values.

4 Conclusion and discussion

We emphasize the need for a hybrid decision support instrument that combines the formal ability of Cost-Benefit Analysis (as Expected-Utility Maximization) to deal with anticipated future learning with the option to intuitively represent climate damages in a Cost-Effectiveness Analysis. Hence we address a decision-making community that sees climate targets as a viable bridging concept as long as, in their view, sophisticated impact models are lacking. Schmidt et al. (2011) have suggested such a hybrid instrument in terms of “Cost-Risk Analysis” (CRA). It represents a trade-off between expected welfare losses due to mitigation and avoided risk based on the exceedance of a climate target, e.g. the 2 °C target.

We implement the scheme of CRA as proposed in Schmidt et al. (2011) in an integrated assessment model for the first time, and discover solutions that would abruptly switch to business-as-usual emission behavior after learning. This clearly contradicts the preference order of proponents of climate targets. Hence we propose an axiomatically more satisfying version of CRA avoiding that difficulty and yet preserving the ability to deal with future learning.

Furthermore, by delivering an universally applicable procedure to calibrate the remaining trade-off parameter between expected welfare losses due to mitigation and avoided risk

we complete the set of assumptions necessary for any implementation of a CRA. We re-emphasize that no temperature target can be achieved with certainty, hence should always be formulated in conjunction with a probabilistic compliance level. We propose to calibrate above trade-off parameter such that the joint temperature-probabilistic compliance level target is met efficiently. For demonstrative purposes we calibrate the trade-off parameter at a probabilistic no-learning interpretation of the 2 °C target, deduced from long term climate policy ambition as documented at the 17th COP to the UNFCCC (UNFCCC 2011). That way we have found a dynamically consistent way to formulate the value function of the 2 °C community with respect to climate targets which they seemed to pursue at the climate conferences. Hereby, we also circumvent the problem of finding an explicit value function such as those for ecological systems or ice sheets.

We apply our implementation of CRA to the problem of calculating the expected value of future information about the climate response to carbon emissions. It is an indicator of importance to determine the value of additional observation campaigns, or the relative importance of different sources of uncertainty. Such calculations also help to determine the relative importance of mitigation options with differing flexibility in adjusting emissions to new information.

Overall, our results support the funding towards reducing the uncertainty on response parameters greatly in the next decades. The expected value of learning climate response perfectly in 2015 rather than in 2075 is between 0.15 % and 0.66 % of consumption every year (up to one third of mitigation cost) depending on normative parameters. It compensates 12–22 % of the overall welfare loss compared to business-as-usual, hence is a significant part of the mitigation question. Furthermore we find that already generated CEA scenarios may be excellent proxies of CRA-based scenarios, as far as next-decades efficient investment paths are concerned.

The main focus of this article is to stimulate a discussion on how climate targets should be interpreted and models calibrated accordingly, such that they could accommodate anticipated future learning in a dynamically self-consistent way. We have shown that CRA can be used productively in a learning framework and can aid policy makers in the next decades until such time that improved impact models are available.

Further analysis could look at potential effects of endogenous learning, further uncertainties or the influence of flexibility in the energy sector. We also note that CRA brings benefits for other types of analysis, not only for considering learning. For example, once CRA is calibrated against a reference case, simulations can be made even if boundary conditions force the solution to violate the target.

Acknowledgments We would like to thank Steven Dirkse for help with the GAMS code.

Open Access This article is distributed under the terms of the Creative Commons Attribution License which permits any use, distribution, and reproduction in any medium, provided the original author(s) and the source are credited.

References

- Allen MR, Frame DJ (2007) Call off the quest. *Science* 318(5850):582–3. doi:[10.1126/science.1149988](https://doi.org/10.1126/science.1149988)
- Allen MR, Frame DJ, Huntingford C, Jones CD, Ja Lowe, Meinshausen M, Meinshausen N (2009) Warming caused by cumulative carbon emissions towards the trillionth tonne. *Nature* 458(7242):1163–6. doi:[10.1038/nature08019](https://doi.org/10.1038/nature08019)

- Anthoff D, Tol RSJ (2009) The Impact of Climate Change on the Balanced Growth Equivalent: An Application of FUND. *Environ Resource Econ* 43(3):351–367. doi:[10.1007/s10640-009-9269-5](https://doi.org/10.1007/s10640-009-9269-5)
- Blau R (1974) Stochastic programming and decision analysis: an apparent dilemma. *Manag Sci* 21(3):271–276
- Bosetti V, Carraro C, Galeotti M, Massetti E, Tavoni M (2006) WITCH - A World Induced Technical Change Hybrid Model. University Ca' Foscari of Venice Economic Research Paper No 46/06. doi:[10.2139/ssrn.948382](https://doi.org/10.2139/ssrn.948382)
- Clarke L, Jiang K, Akimoto K, Babiker M, Blanford G, Fisher K, Vanden J, Hourcade C, Krey V, Kriegler E, Löschel A, McCollum D, Paltsev S, Rose S, Shukla RP, Tavoni M, vd Zwaan B, van Vuuren DP (2014) Working Group III contribution to the IPCC 5th Assessment Report "Climate Change 2014: Mitigation of Climate Change", Chapter 6: Assessing Transformation Pathways
- Edenhofer O, Bauer N, Kriegler E (2005) The impact of technological change on climate protection and welfare: Insights from the model mind. *Ecolg Econ* 54(2-3):277–292. doi:[10.1016/j.ecolecon.2004.12.030](https://doi.org/10.1016/j.ecolecon.2004.12.030)
- Edenhofer O, Knopf B, Barker T (2010) The economics of low stabilization: Model comparison of mitigation strategies and costs. *Energy J* 31:11–48
- Eeckhoudt L, Godfroid P (2000) Risk aversion and the value of information. *J Econ Educ* 31(4):382–388
- Fritz S, Scholes R, Obersteiner M, Bouma J, Reyers B (2008) A conceptual framework for assessing the benefits of a global earth observation system of systems. *Sys J IEEE* 2(3):338–348. doi:[10.1109/JSYST.2008.926688](https://doi.org/10.1109/JSYST.2008.926688)
- Gollier C (2004) The economics of risk and time. The MIT Press
- Held H, Kriegler E, Lessmann K, Edenhofer O (2009) Efficient climate policies under technology and climate uncertainty. *Energy Econ* 31:S50–S61. doi:[10.1016/j.eneco.2008.12.012](https://doi.org/10.1016/j.eneco.2008.12.012)
- IPCC (2013) Climate Change 2013: The Physical Science Basis. Contribution of Working Group I to the Fifth Assessment Report of the Intergovernmental Panel on Climate Change. Cambridge University Press
- Iverson T, Perring C (2012) Precaution and proportionality in the management of global environmental change. *Global Environ Chang* 22(1):161–177. doi:[10.1016/j.gloenvcha.2011.09.009](https://doi.org/10.1016/j.gloenvcha.2011.09.009)
- Kelly DL, Kolstad CD (1999) Bayesian learning, growth, and pollution. *J Econ Dyn Control* 23(4):491–518. doi:[10.1016/S0165-1889\(98\)00034-7](https://doi.org/10.1016/S0165-1889(98)00034-7)
- Kolstad C, Urama K, Broome J, Bruvold A, Olvera MC, Fullerton D, Gollier C, Hanemann WM, Hassan R, Jotzo F, Khan MR, Meyer L, Mundaca L (2014) Working Group III contribution to the IPCC 5th Assessment Report Climate Change 2014: Mitigation of Climate Change Chapter 3 Social. Economic and Ethical Concepts and Methods
- LaValle I (1986) On information-augmented chance-constrained programs. *Oper res lett* 4(5):225–230
- Lempert RJ, Groves DG, Popper SW, Bankes SC (2006) A general, analytic method for generating robust strategies and narrative scenarios. *Manag Sci* 52(4):514–528
- Lorenz A, Held H, Bauer E, Deimling TS (2009) Constraining ocean diffusivity from the 8.2 ka event. *Clim Dyn* 34(5):719–734. doi:[10.1007/s00382-009-0590-z](https://doi.org/10.1007/s00382-009-0590-z)
- Lorenz A, Kriegler E, Held H, Schmidt MGW (2012) How to measure the importance of climate risk for determining optimal global abatement policies. *Clim Change Econ*:3
- Lorenz A, Schmidt MGW, Kriegler E, Held H (2012b) Anticipating climate threshold damages. *Environ Model Assess* 17(1-2):163–175. doi:[10.1007/s10666-011-9282-2](https://doi.org/10.1007/s10666-011-9282-2)
- Luderer G, DeCian E, Hourcade JC, Leimbach M, Waisman H, Edenhofer O (2012) On the regional distribution of mitigation costs in a global cap-and-trade regime. *Clim Chang* 114(1):59–78. doi:[10.1007/s10584-012-0408-6](https://doi.org/10.1007/s10584-012-0408-6)
- Luderer G, Leimbach M, Bauer N, Kriegler E, Tino A, Baumstark L, Bertram C, Giannousakis A, Hilaire J, Klein D et al (2013) Description of the remind model version 1.5
- Mastrandrea M, Field C, Stocker T, Edenhofer O, Ebi K, Frame D, Held H, Kriegler E, Mach K, Matschoss P, Plattner GK, Yohe G, Zwiers F (2010) Guidance Note for Lead Authors of the IPCC Fifth Assessment Report on Consistent Treatment of Uncertainties Intergovernmental Panel on Climate Change IPCC. <http://www.ipcc.ch>
- Morgan MG, Henrion M, Small M (1992) Uncertainty. Cambridge University Press
- Nelson JA (2013) Ethics and the economist: What climate change demands of us. *Ecol Econ* 85:145–154
- Nordhaus WD (2008) A question of balance: economic modeling of global warming Yale University. New Haven, Connecticut
- Nordhaus WD, Popp D (1997) What is the value of scientific knowledge? An application to global warming using the PRICE model. *Energy J* 18(1):1–45
- O'Neill BC, Melnikov NB (2008) Learning about parameter and structural uncertainty in carbon cycle models. *Clim Chang* 89(1-2):23–44
- Patt AG (1999) Extreme outcomes: the strategic treatment of low probability events in scientific assessments. *Risk Decision and Policy* 4(1):1–15

- Pearlman J, Craglia M, Bertrand F, Nativi S, Gaigalas G, Dubois G, Niemeyer S, Fritz S (2011) Eurogeoss: an interdisciplinary approach to research and applications for forestry, biodiversity and drought. In: Proceedings of the 34th international symposium on remote sensing of environment, pp 1–4
- Schellnhuber HJ (2010) Tragic triumph. *Clim Chang* 100(1):229–238. doi:[10.1007/s10584-010-9838-1](https://doi.org/10.1007/s10584-010-9838-1)
- Schmidt MGW, Lorenz A, Held H, Kriegler E (2011) Climate targets under uncertainty: challenges and remedies. *Clim Change* 104(3-4):783–791. doi:[10.1007/s10584-010-9985-4](https://doi.org/10.1007/s10584-010-9985-4)
- Schneider SH, Mastrandrea MD (2005) Probabilistic assessment of "dangerous" climate change and emissions pathways. *Proc Natl Acad Sci USA* 102 15(44):728–35. doi:[10.1073/pnas.0506356102](https://doi.org/10.1073/pnas.0506356102)
- Schneider von Deimling T, Held H, Ganopolski A, Rahmstorf S (2006) Climate sensitivity estimated from ensemble simulations of glacial climate. *Clim Dyn* 27(2-3):149–163. doi:[10.1007/s00382-006-0126-8](https://doi.org/10.1007/s00382-006-0126-8)
- Stern N (2013) The structure of economic modeling of the potential impacts of climate change: grafting gross underestimation of risk onto already narrow science models. *J Econ Lit* 51(3):838–859
- Ulph A, Ulph D (1997) Global warming, irreversibility and learning. *Econ J* 107(442):636–650
- UNFCCC (2011) Report of the Conference of the Parties on its seventeenth session, held in Durban from 28 November to 11 Addendum. Part Two: Action taken by the Conference of the Parties at its seventeenth session. Decision 1/CP.17: Establishment of an Ad Hoc Working Group on the Durban Platform for Enhanced Action FCCC/CP/2011/9/Add 1
- Wakker P (1988) Nonexpected utility as aversion of information. *J Behav Decis Mak* 1(3):169–175. doi:[10.1002/bdm.3960010305](https://doi.org/10.1002/bdm.3960010305)
- Webster M, Jakobovits L, Norton J (2008) Learning about climate change and implications for near-term policy. *Clim Change* 89(1-2):67–85
- Wigley TM, Raper SC (2001) Interpretation of high projections for global-mean warming. *Science* 293(5529):451–4. doi:[10.1126/science.1061604](https://doi.org/10.1126/science.1061604)