Uncertainty & Learning in Global Climate Analysis

Alexander Lorenz

vorgelegt an der Fakultät VI – Planen Bauen Umwelt
der Technischen Universität Berlin
zur Erlangung des akademischen Grades
Doktor der Wirtschaftswissenschaften (Dr. rer. oec.)

Berlin 2011

Contents

Abstract

Climate change, the *21st* centuries challenge for cooperative human decision making, is surrounded by large uncertainties concerning the scientific understanding of the climate system, of climate change induced changes of natural and social systems and of the impacts of those changes on human economic activities and human welfare in general. Parts of these uncertainties will be resolved as science advances and new observations are made. This learning will allow to refine the decisions undertaken to cope with the climate problem.
This thesis is dedicated to examine the role of uncertainty and future learning in the formal assessment of optimal global mitigation strategies for global warming. The central contributions of this study are contained within three research articles.
The first article investigates the validity of the cost-effectiveness framework when applied to the case of climate targets under uncertainty and future learning. The study highlightens two major conceptual problems of this formalism, namely the possibility of negative value of information and infeasibility of the whole decision criterion. As a consequence an alternative decision framework is proposed, the so called cost-risk analysis, that avoids those conceptual problems but still remains based on climate targets.
The second article is motivated by the clash between the general scientific intuition that epistemic uncertainties about the climate system and climate damages should play a major role in determining optimal mitigation policies (and the resulting welfare gain compared to doing nothing) and the results from the integrated assessment models that show only insignificant influence of those uncertainties. We introduce a method of assessing the importance of uncertainty both in its impact on optimal policy and in its impact on the welfare gain from acting upon climate change. We then use a representation of the integrated assessment model MIND that allows to link the decomposed value of climate policy to the structural form of the functions representing the climate cause-effect chain, thereby understanding the negligible effect of uncertainty from the model structure. Finally we propose some changes to the model structure that result in large impacts from including uncertainty.
The third article investigates the circumstances under which the anticipation of future learning about tipping-point-like threshold climate damages would be important for the determination of near term mitigation decisions. We show that this is only the case if the learning occurs within a narrow *anticipation window*. In this case far stronger near term mitigation is optimal to keep the option open to avoid the threshold in case it turns out to lead to severe damages. The location and width of this window is found to be sensitive to the DM's flexibility to reduce emissions. If reducing this flexibilitiy in the MIND model, may this represent political or social barriers, the anticipation window moves towards the present and broadens considerably, thereby increasing the importance of including future learning into the analysis of climate change.
The articles are put into perspective by an introduction into the field that lays out the general linking research questions and general conclusions.

Zusammenfassung

Der Klimawandel, als zentrale Herausforderung des 21. Jahrhunderts für globale Kooperation, ist gekennzeichnet durch enorme Unsicherheiten im wissenschaftlichen Verständnis des Klimasystems, klimainduzierter Veränderungen natürlicher und sozialer Systeme sowie der Folgen dieser Veränderungen für menschliches Wirtschaften und die allgemeine Wohlfahrt. Teilweise werden diese Unsicherheiten durch Fortschritte der Wissenschaft und neue Beobachtungen aufgelöst werden können. Dieses zukünftige Lernen wird es ermöglichen, getroffene Entscheidungen zum Klimaschutz zu revidieren und an neue Situationen anzupassen.

Diese Dissertation widmet sich der Untersuchung der genauen Rolle dieser Unsicherheiten und der Möglichkeit zukünftigen Lernens für die formale Analyse optimaler Vermeidungsstrategien des Klimawandels. Die zentralen Beiträge dieser Arbeit sind in drei wissenschaftlichen Artikeln dargelegt.

Der erste Artikel untersucht axiomatische Zielkonflikte bei der Anwendung der so genannten "Kosten-Effektivitäts" Analyse auf Klimaziele unter Unsicherheit. Die Studie stellt zwei zentrale konzeptionelle Probleme dieses Formalismus fest, wenn man zusätzlich die Möglichkeit zukünftigen Lernens einbezieht: die Möglichkeit, dass zusätzliche Information negativen Wert zugeschrieben bekommt und die Möglichkeit der Unlösbarkeit des ganzen Entscheidungskriteriums. Als Konsequenz wird ein alternatives Entscheidungskriterium vorgestellt, die sogenannte "Kosten-Risiken" Analyse. Diese basiert immer noch auf der Angabe einzuhaltender Klimaschranken, vermeidet jedoch die benannten Probleme.

Die Motivation für den zweiten Artikel liefert der Widerspruch zwischen einer wissenschaftlichen Intuition und den aktuellen Modellergebnissen. Die Intuition sieht einen starken Einfluss epistemischer Unsicherheiten auf die Bestimmung optimaler Vermeidungsstrategien und deren Einfluss auf die Wohlfahrt (im Vergleich zum "Nichtstun"). Die Modelle hingegen zeigen nur einen marginalen Einfluss dieser Unsicherheiten. Diese Studie entwickelt eine Methode, mit der man die Wichtigkeit von Unsicherheit bestimmen kann, sowohl im Sinne des Einflusses auf die optimale Politik als auch im Sinne der Wohlfahrtsveränderung, die diese Politik hervorruft. Weitherin wird eine Darstellung des Modells MIND verwendet, die es ermöglicht, die vorgestellte Metrik zur Messung der Wichtigkeit mit der Struktur der Funktionen zu verbinden, die die Kausalkette des Klimawandels im Modell abbilden. Damit kann man die insignifikante Rolle von Unsicherheit direkt aus der Modellstruktur ableiten. Davon ausgehend testen wir einige, in der Literatur diskutierte, Änderungen der Modellstruktur bezüglich ihres Einflusses auf die Wichtigkeit von Unsicherheit.

Der dritte Artikel untersucht die Umstände unter denen die Antizipation zukünftigen Lernens über Klimaschäden aus dem Überqueren eines sogenannten Tipping Points einen signifikanten Einfluss auf die kurzfristige Vemeidungsstrategie ausübt. Wir zeigen, dass dies nur der Fall ist, wenn das Lernen in einem engen "Antizipations-Zeitfenster" stattfindet. In diesem Fall ist eine striktere kurzfristige Vermeidungsstrategie optimal, die die Option erhält, im Falle schlechter Nachrichten

über die Schäden des Tipping Points selbigen nicht zu überqueren. Die Lage und Breite dieses "Antizipations-Zeitfensters" ist stark abhängig von der Flexibilität, die Treibhausgasemissionen zu reduzieren. Wenn man diese herabsetzt, beispielsweise um politische oder soziale Barrieren zu representieren, so bewegt sich das "Zeitfenster" näher zur Gegenwart und verbreitert sich deutlich. Damit wird die Wichtigkeit von Antizipation für kurzfristige Entscheidungen erhöht.

Eingefasst werden die Artikel von einer Einleitung in das generelle Forschungsumfeld, die auch die zentralen, verbindenden, Forschungsfragen einführt, und einer Zusammenfassung, die Schlussfolgerungen und weitere Forschungsschritte vorstellt.

Acknowledgements

I want to thank many colleagues from the Potsdam-Institute for Climate Impact Research from whom I learned much and who contribute to an outstandingly open, stimulating and constructive working environment from which I profited in many aspects.

Especially, I would like to thank Hermann Held, whose enthousiastic and impressive presentation of the “inverse problems in climate science” inspired me to dive into the fascinating domain of climate science. He offered me the opportunity to work with him and learn from him, first on a masters project, and in continuation on this thesis. Without his scientific, intellectual, moral, and financial support this thesis would simply not have been possible.

In equal intensity I want to thank Elmar Kriegler, for his collaboration, help, and profound advise as well as for his continuous and increasing engagement in this PhD project, that helped me to follow some of the many ideas to fuition. Without him, this thesis might have been possible but would possibly not have been finished in its current form.

I would also like to thank Ottmar Edenhofer for supervising and examining this thesis as well as for providing the whole research domain with the leadership that enables the constructive, positive environment I could enjoy over the last four years.

With particular gratitude I want to mention my collaboration with Matthias Schmidt. With many discussions, a continuous intellectual sparring, and an enviable motivation, he helped me to live through the unavoidable downs of a PhD. Thanks for your collaboration.

Special thanks are due to Kai, Martin, Matthias, and Franziska for sharing an office with me, as well as to Lavinia, Kristin, Odo and Islay for many relaxing counterpoints to the office.

I gratefully acknowledge funding by a scholarship of the German National Science Foundation, from December 2009 until the finalization of this thesis.

And last, but surely not least, I want to thank Friederike, for your support, for your love, for everything.

Chapter 1

Introduction

Anthropogenic climate change is one of the most prominent and most pressing global problems of the unfolding century. While the basic principles of the cause-effect chain from human economic activity to changing global climate conditions are well understood, large uncertainties remain in the quantification of potential future evolution pathways of the climate system and the resulting consequences for human civilization. However, during the last decades an enormous amount of insights and evidence have been gathered and our understanding of the complex socio-economic-climate system advances. Hence uncertainties and the expectation of future advances in system knowledge are prominent features of the global climate problem.

This thesis aims at contributing to the assessment of the importance of these features of uncertainty and learning for the integrated assessment of climate change mitigation policies. This will help to clarify the justification of their prominence in the ongoing debate. This chapter introduces the broader context of this thesis and clarifies its objectives. The basic physics of climate change is reviewed in Section 1.1. Sections 1.2 and 1.3 introduce the economic approach to the analysis of climate change. The issues rising from the recognition of uncertainties and learning about key components of the combined socio-economic-climate system and the ways these problems are approached in the literature are reviewed in Section 1.4. Within in this context, Section 1.5 introduces the objectives and structure of the thesis.

1.1 The Physics of Climate Change

The basic processes of anthropogenic global warming are well understood. A number of human activities like pasturing land, cutting down forests, producing livestock, and especially the burning of fossil fuels, leads to an increased flow of greenhouse gases (e.g. carbon dioxide, methane, etc) into the atmosphere. This increases the atmospheric concentration of those gases above the pre-industrial equilibrium level. As those greenhouse gases are transparent for incoming short wave radiation from the sun but nearly opaque for the long wave heat radiation from the surface of the earth a higher concentration of greenhouse gases perturbs the energy balance between the incoming solar shortwave radiation and the outgoing longwave radiation which determines the earth's surface temperature. The resulting radiative forcing at the top of the atmosphere leads to an increase in GMT (until new equilibrium of the radiation balance). This basic concept of the greenhouse effect was already described by Arrhenius (1896).

From the industrial revolution (1850) the atmospheric concentration of greenhouse gases (measured in CO_2 equivalents) has increased by about 50%, from 287 ppm (parts per million) to 433 ppm in 2004 (Solomon et al., 2007). Global emissions of carbon dioxide equivalent greenhouse gases as of 2004 amounted to 49 Gt (Trenberth et al., 2007) and have been increasing by about 4% per year ever since. In 2009 the CO_2-only emissions summed up to 29.5 Gt (compared to 27.7 Gt in 2005) (IEA, 2011). This led to an addition of 2 ppm of CO_2 in the atmosphere per year on average between 2004 and 2009 (Stern, 2007), with a current atmospheric CO_2 concentration of 389 ppm (as measured at Mauna Loa in September 2011, Conway & Tans, 2011). In the same period (1850-2005), global mean temperature increased by 0.76 ± 0.19 °C, at an accelerating rate. The warming over the last 50 years is almost double that of the previous 100 years (Trenberth et al., 2007). Given our sound understanding of the thermodynamics of the climate system it is very unlikely that those changes in temperature can be explained without external forcing (like changes in solar irradiation, volcanic forcing, and anthropogenic activities). As temperature has been strongly increasing when non-anthropogenic factors would be likely to have produced a cooling effect, attribution studies find it very likely that it has been anthropogenic greenhouse gas emissions that have caused most of the observed warming since the mid-20th century (p60, Solomon et al., 2007).

During the last century, and especially within the last decades, a large number of changes in both the global and regional parts of the climate systems have been detected: melting of glaciers all over the world, changes in regional temperature and precipitation patterns, a rise in the global sea level of about 25 cm, and an increase of 0.7 °C in the global mean temperature within the last century (IPCC, 2007). Due to the large inertia of the climate system, stemming from the enormous capacity of the oceans to store heat and the already large, accumulated athmospheric carbon pool, this warming would continue for centuries even if greenhouse gas emissions were stopped immediately. Following Ramanathan & Feng (2008), when keeping athmospheric CO_2 concentration constant at the 2005 level (but reducing the emission of aerosols over the 21st century), we would be committed by the current cumulative emissions to an expected global warming of $2.4(1.4-4.3)$ °C above pre-industrial temperatures.

Although the thermodynamics of global warming are well understood, prognoses of future changes in the different climate subsystems are inherently difficult, due to some less understood processes that define the dynamics of the climate system. The sign and amplitude of different feedback mechanisms like the water vapor feedback and the radiative forcing from changing cloud cover are still highly uncertain. Additionally a large number of uncertain parameters in our representation of the climate system influence the cause-effect chain of global warming. Prominent examples of uncertainties are the so-called climate sensitivity (Roe & Baker, 2007), the equilibrium change in mean global surface temperature due to a doubling of atmospheric CO_2 concentration from the pre-industrial level, and the transient climate response (Stott et al., 2006), that describes the immediate temperature response to a 1% yearly increase in CO_2 concentration at the point where the doubled concentration level is reached. Projections of global mean temperature within the 21st century are also highly uncertain. In the absence of any climate policy, the IPCC predicts a global warming of $4.0(2.4-6.4)$ °C above pre-industrial mean temperatures to be reached by the year 2100 (Meehl, 2007).

The role of these uncertainties and potential future improvements in our knowledge is the central matter of this thesis and will be introduced in more detail in Section 1.4.

1.2 The Impacts of Climate Change

A large number of changes in the global climate system are expected to accompany the rise in global mean temperature or follow from it. Global sea levels are expected to rise by $0.18-0.59$m by the end of this century (IPCC, 2007). Newer estimates see the potential for even greater increases (Vermeer & Rahmstorf, 2009). Due to varying regional compositions of the drivers of sea level rise (SLR) its regional extent is also expected to differ significantly (Yin, 2009). Precipitation will change both in amplitude and in the regional distribution (IPCC, 2007). Widespread mass losses from glaciers and reductions in snow cover over recent decades are projected to accelerate throughout the 21st century, reducing water availability, hydropower potential, and changing seasonality of flows in regions supplied by meltwater from major mountain ranges (IPCC, 2007).

One of the most important aspects of the impact of global warming is that for a given change in global temperature, the regional changes can differ enormously. Thus a “moderate” global warming of about 2 °C above pre-industrial temperatures can lead to tremendous changes in different climate subsystems and sensitive regions. The most severe changes in regional climates have been observed in the Arctic, where the temperature change is two times as high as the global mean (Trenberth et al., 2007). The minimum extent of the Arctic sea ice cover during summer has decreased by about 30% over the last decades (Stroeve, 2011). If this disaggregation is going to continue at the same speed, the Arctic Ocean could be ice-free in summer by 2050. Additionally, once the sea ice cover falls below a certain level of thickness, it becomes more vulnerable to changes in regional weather conditions. Hence the volatility of sea ice cover increases and thus the predictability of future changes decreases significantly. There are indications for the latitudinal position of the inner tropical convergence zone (ITCZ) changing in response to temperature gradients between the northern and southern hemisphere (Broccoli et al., 2006). As temperature increase in response to a rise in athmospheric greenhouse gas levels is expected to vary regionally, this could mean that the ITCZ, as a band of highest precipitation, could move, thus regions with once high levels of precipitation could dry up while others will encounter higher seasonal precipitation events. This change will also lead to a corresponsing movement of the arid belt that follows outside of the ITCZ.

A second important feature of global warming concerns the (partial) irreversibility of climate change . Although the projections of global mean temperature are smooth over time, i.e. a linear response to the forcing, this does not have to hold for regional changes. A number of climate subsystems have been identified that can react in a strongly non-linear fashion to changes in the overall forcing. Those so-called tipping elements (Lenton et al., 2008) were introduced to describe the components of the Earth system that can be switched – under particular conditions – into a qualitatively different state by small perturbations. Thereby the term ‘tipping point’ is used to refer to the critical point (in forcing and a feature of the system) at which such a transition is triggered. Furthermore, most of these subsystems show hysteresis behavior, which means that once such a tipping element reaches another stable state, a simple reversion of the forcing will in general not suffice to switch the system back into its original state. Examples of these systems are the thermohaline circulation in the north Atlantic, the Amazon Rainforest, and the South Asian monsoon circulation. Another example is the Greenland ice sheet.

Some of these systems are characterized by the existence of a positive feedback mechanism that causes the transition in the system, once the forcing passes the ‘tipping point’. In the case of the Greenland ice sheet, increased melting due to higher temperatures leads to a lowering of the thickness of the ice sheet. But by lowering the altitude of the maximum thickness, the surrounding temperature increases even more, which further exacerbates the melting process. Once the thickness has fallen below a certain threshold, the total disaggregation of the ice sheet cannot be stopped by simply bringing the temperatures back to normal.

Besides the regional positive feedbacks which can lead to a tipping behavior in climate subsystems, there also exist a number of positive feedbacks on a global scale. The greenhouse effect of water vapor, as part of the overall greenhouse effect, is one of the potential, global positive feedbacks. As the global temperature is increased by the greenhouse effect from airborne water vapor, even more water enters the atmosphere and thereby further strengthens the greenhouse effect. However, the strength of this feedback is not known precisely as several complications occur (Forster et al., 2007). The radiative effect from increasing water vapor content is decreasing by a spectral saturation effect. A second example of a global feedback is the dependence of the ability of the oceans to store CO_2 from the atmosphere on the global temperature itself. This ability is influenced by temperature increases through a multitude of effects (see Fung et al., 2005): Warming reduces solubility of carbon, increases ocean stratification, reduces vertical mixing, and slows the thermohaline circulation. This also impacts on the oceanic biological productivity. Partially loosing the capacity of storing CO_2 in the oceans would further increase the greenhouse effect. Finally, there are some other sources of potentially drastic global positive feedbacks. Large amounts of CO_2 and methane are stored in the permafrost soil of the northern hemisphere and even larger amounts of methane are stored at the bottom of the marine continental shelf regions in the form of methane hydrates. An increase in regional temperatures and oceanic temperature could lead to a large scale release of those greenhouse gases into the atmosphere and lead to a drastic, so-called ‘run-away greenhouse effect’ of uncertain but tremendous amplitude (e.g. Keller et al., 2007; Lenton, 2011). The bedeviling feature of these non-linear mechanisms is the inherent uncertainty about the location of the tipping points and about the consequences of a switching for other parts of the climate system. However there are some instances of abrupt changes in the past of the Earth system found in paleo-climatic data, that might represent switches of tipping elements, and there are also some methods for early warning about approaching a tipping point from observations (Dakos et al., 2008; Held & Kleinen, 2004; Lenton, 2011).

Uncertainty and potentials for future learning prevail in each of the mechanisms for climate change impacts, especially in the non-linear subsystems. Those uncertainties add another dimension to the overall uncertainty about climate change and will also be tackled in this thesis.

1.3 The Economics of Climate Change Mitigation

The impacts of climate change on the earth system are not of concern per se. During its long history, the earth has experienced many drastic or abrupt changes within the climate system, like the so-called snowball earth events (Hoffman et al., 1998) or major meteoric impact events (Alvarez et al., 1980). Thus, from a very aggregated perspective one could say that abrupt and extreme changes within the climate system are “quite normal”. Such events caused massive changes in atmospheric composition, the global mean temperature, and regional conditions for life, resulting in the extinction of 95% of all species. Anthropogenic climate change differs from those events: The observed and expected changes in natural systems take place with unrecedented speed (100’s rather than 100000’s of years). Human society has adapted to the relatively stable climate of the past holocene and whenever the climate deviated slightly from this holocenic "stable state" societies suffered and whole civilisations broke down. Therefore the changes in the natural systems due to anthropogenic climate change will have large impacts on human activities, health and conditions for human life in general. Furthermore, as human activities are the cause of the current changes in the climate system, and additional climate change could be avoided by stringent mitigation policies, the problem of anthropogenic climate change belongs to the realm of decision problems. Economic theory provides tools for evaluation of climate induced damage as well as for decision making in the climate context.

One of the major threats from climate change is not the change in the mean state of regional climate variables but the rising probability of extreme events like droughts, floods, or storms. Both kinds of impacts, changes in the mean state and a higher risk of extremes, will impact humans differently across regions, income groups, age groups and other attributes. Sea level rise, for example, will mainly impact coastal regions by increasing the risk for extreme flooding events that threaten human life, health and infrastructure. However, other regions would be affected indirectly. As a sea level rise of $0.5-1$m would increase the number of people under risk of coastal flooding by up to 200 million (NRC, 2011), large migration movements are expected to follow from climate change impacts. A moderate increase of temperature would also lead to different impacts across world regions, as the northern hemisphere (e.g. Russia, Scandinavia, or northern Germany) might even benefit from increased crop yields while arid and warm regions like southern Europe or sub-Saharan Africa will encounter drastic reductions in crop yields from reduced water availability and soil erosion (Edenhofer et al., 2010). Welfare economics provides the means to evaluate the impacts of climate change and aggregate them to a measure of losses to human welfare. However such an evaluation poses several severe challenges. Firstly, the inherent uncertainties in predicting all necessary aspects of climate change relevant to the damage assessment are immense. Secondly, the incommensurability of different impacts poses an enormous problem for their aggregation, e.g. when comparing losses in biodiversity to an x% decrease in agricultural productivity. And finally, some very intricate normative issues arise, when aggregating climate damages over time, regions and over uncertain possible future states of the world (Stern, 2007; Nordhaus, 2008a; Dasgupta, 2008): how to value future damages against current welfare, how to value damages inflicted on rich people relative to damages inflicted on poor people, and how to weigh the risk of catastrophic damages?

Economic Theory provides two important insights for managing climate change. First, the climate change phenomenon is described as a failure of markets to anticipate the social damages inflicted by economic activities. This provides the economic rational for a regulating authority to intervene and internalise the negative effects from climate change. The non existence of an effective global market regulation body leaves this intervention to be negotiated between many sovereign countries that are in a competitive situation on the global markets. This delivers an explanation as to why global action to mitigate dangerous anthropogenic interference with the climate system, as stipulated by UNFCCC, Art. 2, is so difficult to achieve although the potentially drastic consequences of unmitigated climate change are provided by science. Second, it provides some important concepts to analyze the trade-offs between current and future generations, between rich and poor regions and between the costs of mitigation and the risk of unmitigated climate damage.

Climate change related damage caused by anthropogenic greenhouse gas emissions comprise a market failure as each emitter of greenhouse gases only experiences a tiny part of the overall damage from climate change. Thus the emitter does not anticipate the causal link between her actions and the resulting negative implications. A large number of self interested actors that compete on the global market would thus not agree on any non-trivial mitigation action. This situation is equally referred to as a market failure, an externality, or a common good problem (Hardin, 1968). The solution to a common good problem is that the governmental authority regulating the market (and representing social interest) puts a price on the activity that creates the externality. Thereby the external damages get internalized and the market once again efficiently implements the social interest.

In the context of climate change this solution corresponds to the installation of a globally binding regime that puts a price on greenhouse gas emissions, either explicitly, by a carbon tax, or implicitly, by a binding limit on absolute emissions combined with a trading scheme of carbon emission permits (Weitzman, 1974). In the absence of further externalities, like technological spillovers (Leimbach & Baumstark, 2010), and without uncertainty about mitigation costs and damages, both instruments efficiently internalize the climate externality. The handling of multiple externalities and asymmetric information would require additional policy instruments and represent a field of research on their own. The practical implementation of a globally binding climate regime has not been successful up to now (COP 16, 2010). The lack of an authority regulating the global market, the diversity of interests of the different major economic actors, the inherent difficulty of including the interests of future generations, and the difference between those actors who cause(d) climate change and those who will experience the damage are major obstacles for such and agreement. However, at least the UN have collectively recognized the issue of climate change mitigation by (loosely) committing to a target of constraining global warming to a 2°C increase from the pre-industrial global mean temperature (COP 16, 2010). Furthermore, unilateral commitments for mitigation of greenhouse gas emissions have been piloted by several countries (www.climateactiontracker.org). However, one can entertain some doubt as to whether those actions will suffice to counteract climate change.

Focusing on the trade-offs inherent to the question of the appropriate level of mitigation action, economic theory provides two different approaches. The first one includes a full monetization and aggregation of climate damages as well as mitigation costs and aims at a formal cost-benefit analysis of mitigation actions (as done e.g. by Nordhaus, 1994 , 2008a). The second approach refrains from monetizing damages and provides an assessment of climate related risks instead. By fixing the risk exposure to a certain level, cost efficient mitigation policies can be derived (Schneider & Mastrandrea, 2005; Held et al., 2009). While the choice of an optimal policy is included in the first approach and has to be done ex post within the second approach, economists use to emphasize that the main benefit of both methods lies in the exploration of the consequences of alternative policy scenarios rather then delivering the one optimal solution (Weitzman, 2011). When considering the problems in the assessment of damage and the dependence on normative settings discussed above, a healthy warning about the use of formal cost-benefit analyses seems more than appropriate.

The formal implementation of both approaches is done in so-called integrated assessment models (IAMs), which include a representation of the socio-economic and the climate system. This way both the impact from economic activities on the climate system and the impact from climate change on economic activity are captured within one modeling framework. The pioneering work in this field is the DICE model by Nordhaus (1994). Integrated Assessment models vary in the detail and methodology of the description of the single subsystems as well as in the choice of normative settings. Large parts of controversies about differing results from IAMs can actually be related to the choice of normative parameters (Stern, 2007; Nordhaus, 2006). What they have in common though, is that they mostly only consider mitigation in terms of reductions in GHG emissions (exceptions are e.g. Ingham et al., 2007, Bosello et al., 2010, who also consider adaptation). Taking into account the whole chain of causes and effects of climate change, along the so-called Kaya identity (see Waggoner & Ausubel, 2002, a comprehensive set of policy levers is thinkable to reduce climate change impacts: Limiting population growth, per capita economic output, energy intensity of economic activity, carbon intensity of energy production (and other activities like agriculture and transport), or capturing greenhouse gases at the production site, combined with storing it (Carbon Capture and Storage), would decreases greenhouse gas emissions. Methods like the enhancement of oceanic carbon uptake through iron fertilization (Boyd et al., 2007) and the reduction of atmospheric carbon content through “air capture” (Lackner et al., 2008) are aiming at managing the carbon cycle directly to avoid an increase in athmospheric concentration of greenhouse gases. The atmospheric radiation balance could be influenced by emitting aerosols or enhancing low level cloud cover by emission of sea salt particles, both of which have a cooling effect. Another possibility would be the installation of orbital “sun blocking” facilities, like mirrors. While the carbon management is aimed at undoing the harm of anthropogenic carbon emissions, the radiation management only aims at combating global warming and leaves out other effects, like the acidification of the oceans (Keith, 2001). Formerly both groups of methods, carbon cycle and radiation management have been summarized under the term “geo-engineering”. As these options come with potential side effects of highly uncertain types and magnitudes, from a risk perspective they are to be seen as a measure of last resort rather than a primary policy option (Victor, 2009). Finally there will be the necessity to cope with the impacts of climate change by adapting the the changing natural systems and by building up resilience towards extreme events.

1.4 Uncertainty and Learning in Global Climate Analysis

As highlighted by the different examples above, uncertainty is a pervasive feature of global climate policy analysis, because it is inherent in all parts of the cause effect chain of the integrated assessment of climate change: from future population change, over economic development, energy intensity of economic activity, intensity in greenhouse gas emissions from the generation process of final energy carriers, the carbon cycle, the temperature response to increases greenhouse gas concentrations, and the impacts of a changing carbon content and temperature of the environment on biological, social and economic systems. Our understanding of every single step of this process chain is hampered by uncertainty.

Within the integrated assessment of climate change, these uncertainties influence both, the stringency of optimal mitigation policy and the welfare implications from different possible mitigation strategies. The amplitude and direction of these impacts determine the "importance" of accounting for uncertainty and learning within the integrated assessment of climate change.

Uncertainty influences optimal climate policy in different ways. Decision makers (DMs) are generally modeled to be risk averse. This means that they dislike uncertainty, i. e. they are worse off when given an expectation of different possible, but uncertain, outcomes instead of the outcome for the expected value of the underlying parameter determining the outcomes. This attitude is represented by a concavity in the DM's utility function, which describes how much welfare, or utility, the DM derives from a given level of consumption of goods. If uncertainty is directly represented on the level of consumption, the DM's risk aversion is the only reason for disliking uncertainty. However, if the DM is uncertain about a parameter that itself influences the consumption level, the concavity of utility with respect to this uncertain parameter can also arise from a concavity in the dependency of the consumption level on the uncertain parameter. For the optimal stringency of climate policy the differential effect of uncertainty between different decisions is important: is the uncertainty the DM dislikes greater or smaller when committing to ambitious mitigation instead of following a business as usual scenario? The more ambitious the mitigation scenario, the less the deviation of the climate system from its current state and the less the uncertainties about climate change induced damages. On the other hand a strict mitigation scenario induces uncertainties in the social costs of mitigation. It remains an empirical question which of the mentioned effects dominates the analysis.

The uncertainties about the socio-economic system and the climate system might change over time due to observations or improvements in scientific understanding of the underlying system. Future learning, in general, involves two effects on optimal climate policy. After receiving new information the DM might want to adjust her policy to the new state of knowledge. In the absence

of a multi agent framework, when only one DM plays a "game against nature" she will always be better off when she has the possibility to learn simply because there might be some cases where adjusting the policy after learning leads to welfare gains against keeping to the default "no learning policy", but at least she can always simply chose the default policy so she cannot be worse off (Gollier,2004).
In addition to the post-learning adjustments in the policy, the DM might also want to change his near-term policy anticipating future learning. This might be due to irreversibility following the decisions. For example the DM might want to strengthen her effort in near-term mitigation to keep the option open to stay below a climate tipping point if she will learn about the location or severity of such a threshold in the future. On the other hand she might want to decrease her effort in the near term to avoid sunk costs of early mitigation in case she learns in the future that climate change damage is not as severe as expected. This has been used as an argument to postpone investments into mitigation of GHG emissions until more is known about the impacts of climate change (for a theoretic presentation of the argument see (Baker, 2006).
Which side plays out to be dominant depends on the representation of flexibility and irreversibility in decisions and on the representation of uncertainties and future learning possibilities, hence it is an empirical question.

It is not only an empirical question whether uncertainty and learning lead to more or less stringent short-term emission reductions but also how significant those adjustments of optimal climate policy are in terms of additional welfare gains, compared to the overall net social benefit of acting upon climate change. As the comprehensive handling of uncertainties and future learning possibilities highly complicates the integrated assessment of optimal mitigation policy, it has to be asked how much we would lose in terms of welfare when applying a suboptimal policy, derived from an analysis that neglects uncertainty or learning.

The literature that is concerned with epistemic uncertainty and future learning can be divided into studies on the effect of uncertainty and on the effect of learning. Thereby the focus is mostly on the effect of uncertainty and learning on optimal (near term) decisions and less on the welfare effects from the adjustments of optimal decisions due to the accounting for uncertainty and learning.

Within integrated assessment models (IAMs) the effect of uncertainty about the climate system or about climate damage is found to only lead to small changes in the optimal stringency of mitigation policy compared to a situation where all uncertain parameters are fixed to their expected value (Peck & Teisberg, 1993; Webster et al., 2008). This result changes when uncertainty about the normative parameters of the decision framework is included (Pizer, 1997). This uncertainty leads to strong changes in optimal policy (up to 30% in cumulated emissions). Another source for high impacts of uncertainties is the inclusion of so-called "fat tailed distributions" of climate response or climate induced damage, as investigated by Weitzman (2009). He showed that in this case the existence of low probability, high impact events from climate change leads to significantly higher mitigation efforts and can even dominate the cost-benefit analysis in the sense that, under certain assumptions, society would be willing to spend almost all of its GDP to prevent a very unlikely catastrophic future. Schmidt et al. (2011a) show that even without including fat tailed distributions uncertainty can have a substantial effect on optimal policy, if the heterogeneity of the

climate damage distribution across the global population is taken into account.
The question of whether the adjustment of an optimal policy to uncertainty leads to significant changes in the welfare gain from acting upon climate change has received less attention (an exception is Pizer, 1997).

Chapter 3 builds upon this example and embarks on the issue of determining the importance of explicitly accounting for uncertainty within the integrated assessment of climate change. The benefit of accounting for uncertainty can be evaluated by comparing the outcome of a best guess policy in a deterministic setting with an expected value maximizing policy under uncertainty. Chapter 3 argues that accounting for uncertainty is important if it leads to significant changes in optimal mitigation policy which in turn leads to significant changes in the net welfare benefit from mitigation action. A framework is developed that allows to analyse this measure of the importance of accounting for uncertainty within an integrated assessment model and to relate the importance to the functional structure of the climate cause effect chain. The negligible magnitude of welfare changes due to introduction of uncertainty found in many other studies is confirmed and related to compensating factors from different steps within the cause effect chain.

Within the literature, the effect of future learning on optimal near-term mitigation decisions is termed the so-called "anticipation effect". Studies employing the expected utility maximization framework found that this anticipation effect on near-term mitigation decisions is small (Peck & Teisberg, 1993; Ulph & Ulph, 1997; Webster, 2002; Webster et al., 2008; O'Neill & Melnikov, 2008). Again, the welfare effect from this anticipative behavior has received less attention, which is understandable in the absence of an effect on decisions. Within the risk management approach mentioned above (Section 1.3), that maximises welfare under the constraint of limiting the risk of crossing a climate treshold to a certain probability, a far higher impact of the anticipation of future learning on near-term decisions and resulting welfare gains has been found (Webster et al., 2008; Bosetti et al., 2008). However, as is shown in Chapter 2 (Schmidt et al., 2011b), adopting this approach leads to axiomatic inconsistencies when learning is included. Keller et al. 2004 have found that the impact of the anticipation of future learning strongly increases if uncertainty about highly non-linear climate damage, e.g. from tipping elements in the climate system, are included. While the theoretic literature agrees that a strongly non-linear relationship between the decisions of a problem and the resulting welfare, together with irreversibilities, can lead to anticipation effects (see e.g. Baker, 2002) a systematic analysis of the significance of these anticipation effects is still missing.

Chapter 4 enters the discussion at this point and develops a framework for investigating the importance of anticipation of future learning, both in terms of changes in optimal decisions and in resulting welfare gains. Building upon the work of Keller et al. (2004) it first introduces a notation of the so-called expected value of anticipation that results from a decomposition of the overall expected value of future information into the welfare gain from pre-learning and post-learning decisions adjustments. While the overall expected value of future information is important when comparing the importance of the reduction of different possible uncertainties, the expected value of anticipation measures the importance of explicitly including future learning into the analysis of an optimal short-term policy. Chapter 4 shows that anticipation can become crucial both in terms of necessary adjustments of pre-learning emissions and resulting welfare gains if learning about an irreversible threshold is included. Conditions on the time of learning and the threshold characteristic are determined, for which this is the case. They can be summarized as a narrow "anticipation window".

1.5 Thesis Outline

Summarizing the discussion above, this thesis contributes to the question of whether uncertainty and learning play an important role in the integrated assessment of climate change by developing a framework for testing the importance within complex models that allows to relate the magnitude of the welfare effect from uncertainty and learning to the functional model structure. The framework is applied to investigate the circumstances under which anticipation of future learning about tipping-point-like threshold damages leads to significant changes in the optimal near-term mitigation policy and corresponding welfare improvements. Additionally the viability of different decision frameworks for the investigation of optimal mitigation under uncertainty and learning is investigated. This analysis is conducted within the three core chapters of the thesis (Chapters 2-4). The Chapters will be outlined in the following and the author's contributions to each of the single articles will be mentioned.

Chapter 2: This chapter is to be seen as a methodological prerequisite for the central analysis of this thesis as it investigates the acceptance of a growingly popular decision criterion for the analysis of uncertainty and future learning.[1] Climate Targets are becoming ever more influential as witnessed by the recent adoption of the 2°C target by UNFCCC COP 15. As a consequence, many studies limit themselves to finding least-cost solutions to achieve these targets in a cost-effectiveness analysis. This article first argues that the 2°C target, for instance, is only meant to be met with a certain probability if uncertainty about global warming is taken into account. Meeting it with certainty would simply be too costly or even impossible. Cost-effectiveness analysis for the resulting probabilistic targets is then shown to imply major conceptual problems that prevent the consideration of learning about uncertainty, which constitutes an essential part of the problem. The article therefore proposes an alternative decision criterion that performs a trade-off between aggregate mitigation costs and the probability of crossing the target. This criterion avoids the conceptual problems of cost-effectiveness analysis and is still to some extent based on given climate targets. This article has been published as "Schmidt, M.G.W., A. Lorenz, H. Held, E. Kriegler 2011. Climate Targets under Uncertainty: Challenges and Remedies. Climatic Change: Letters 104 (3-4): 783-791". M.G.W. Schmidt conceived the idea for this research, performed the analysis and wrote the article. The co-authors, and A. Lorenz in particular, contributed with extensive discussions concerning all stages of the analysis as well as with several internal revisions of the manuscript.

Chapter 3: A common result from the cost-benefit literature is that the inclusion of uncertainty

[1]The following description of Chapter 2 is taken from Schmidt (2011).

about the climate system response to anthropogenic greenhouse gas emissions and the climate induced damages has only small effects upon the stringency of optimal climate mitigation and the resulting welfare gain from adjusting the policy to uncertainty. This negligibility of "normal" uncertainty within the integrated assessment of climate change clashes with the common intuition, at least of the authors, that uncertainty has to be of great importance. This study presents new insights into the source of the negligible uncertainty effect. Unlike previous studies, we go beyond general findings on sufficient conditions for a negligible uncertainty effect. Since such conditions are not fulfilled by integrated assessment models, the magnitude of the uncertainty adjustment effect becomes an empirical question. We present a method to analyze the importance of uncertainty by tracing it through the cause-effect chain from greenhouse gas emissions to temperature change to induced climate damage to welfare implications. This allows us to explain the negligible uncertainty effect as a result of compensating factors in the cause-effect chain. More concretely, we introduce a decomposition of the overall benefit of climate policy into single components: The benefit of the best guess policy, the re-evaluation of the best guess policy under uncertainty, and the value of adapting the optimal policy to uncertainty. We extend the decomposition to the case of perfect learning and analyze the relative importance of all components in the MIND model. Additionally we project the complex integrated assessment model to an a-temporal marginal cost-benefit picture. This allows us to connect the different components of the overall benefit of climate policy to the functional form of the marginal benefits and costs. This understanding of the missing uncertainty effect allows to identify changes in the formulation of the climate cause effect chain that would lead to significant impacts from uncertainty. Examples are more convex climate damages (e.g. exponential damages) or a less concave (e.g. linear) response of temperature to cumulated emissions. This article has been submitted to Climate Change Economics as "Lorenz, A., E. Kriegler, H. Held, M.G.W. Schmidt. How important is Uncertainty for the Integrated Assessment of Climate Change?". The research question for this article was jointly developed by all four authors. A. Lorenz developed the article design, conducted the analysis and wrote the article. The co-authors, and especially E. Kriegler, contributed with extensive discussions and with several thorough internal reviews of revised versions of the manuscript.

Chapter 4: Uncertainty, and especially learning about strongly non-linear, tipping-elements-like damages, might have a strong influence on near-term mitigation decisions. This article systematically investigates both, the changes in optimal near-term mitigation effort and the associated welfare gain relative to the overall benefit from learning due to the anticipation of future learning about the threshold damage. The analysis, conducted within the IAM MIND shows that learning about threshold damage is of significance if and only if the learning happens within a specific, narrow, "anticipation" time window. In this case, the additional early mitigation effort keeps the option open to prevent crossing the threshold if the future learning reveals severe threshold damage. Future learning has no significant effect on near-term policy otherwise. Within the "anticipation window" the welfare gain from anticipating future learning is significant and contains nearly the complete value of information. Learning is still valuable, but not its anticipation, if it happens earlier (outside the anticipation window), if the DM is flexible enough to react after the information has arrived, conditional on the message. Additionally, the article introduces some novel concepts for the analysis of the separate welfare effect from anticipation of future learning. This article has been accepted for publication in Environmental Modelling and Assessment as “Lorenz, A., M.G.W. Schmidt, E. Kriegler, H. Held. Anticipating Climate Threshold Damages.” The research question and design for this article was developed jointly by all four authors. The analysis was performed by A. Lorenz, who also wrote the main part of the article. M.G.W. Schmidt made substantial contributions to conceptualizing the results and rewriting the manuscript in several internal revisions.

Chapter 5 summarizes the results from the different articles and draws some conclusions for the importance of accounting for uncertainty and anticipating future learning in the decision process. Requirements for further analysis and future research questions are given in an Outlook.

Chapter 2

Climate Targets under Uncertainty: Challenges and Remedies*

Matthias G.W. Schmidt
Alexander Lorenz
Hermann Held
Elmar Kriegler

*This chapter has been published as Schmidt, M.G.W., A. Lorenz, H. Held, E. Kriegler 2011. Climate Targets under Uncertainty: Challenges and Remedies. *Climatic Change: Letters* 104 (3-4): 783-791

Climatic Change (2011) 104:783–791
DOI 10.1007/s10584-010-9985-4

LETTER

Climate targets under uncertainty: challenges and remedies

A letter

Matthias G. W. Schmidt · Alexander Lorenz · Hermann Held · Elmar Kriegler

Received: 16 June 2010 / Accepted: 22 October 2010 / Published online: 26 November 2010

Abstract We start from the observation that climate targets under uncertainty should be interpreted as safety constraints on the probability of crossing a certain threshold, such as 2°C global warming. We then highlight, by ways of a simple example, that cost-effectiveness analysis for such probabilistic targets leads to major conceptual problems if learning about uncertainty is taken into account and the target is fixed. Current target proposals presumably imply that targets should be revised in the light of new information. Taking this into account amounts to formalizing how targets should be chosen, a question that was avoided by cost-effectiveness analysis. One way is to perform a full-fledged cost-benefit analysis including some kind of monetary damage function. We propose multi-criteria decision analysis including a target-based risk metric as an alternative that is more explicite in its assumptions and more closely based on given targets.

1 Introduction

Climate targets have been widely discussed since the United Nations Framework Convention on Climate Change (UNFCCC 1992). More recently, the European Union (European Council 2005) and the Copenhagen Accord (UNFCCC 2009) adopted the 2°C-target, which calls for limiting the rise in global mean temperature with respect to pre-industrial levels to 2°C.

Electronic supplementary material The online version of this article (doi:10.1007/s10584-010-9985-4) contains supplementary material, which is available to authorized users.

M. G. W. Schmidt (✉) · A. Lorenz · E. Kriegler
Potsdam Institute for Climate Impact Research,
Telegraphenberg A31, 14473 Potsdam, Germany
e-mail: schmidt@pik-potsdam.de

H. Held
University of Hamburg - Klima Campus, Bundesstr. 55,
Hamburg 20146, Germany

There are large uncertainties involved in climate change. Under probabilistic uncertainty about climate sensitivity, for instance, a certain emissions policy leads to a probability distribution on temperature increases. It is in general impossible or at least very costly to keep the entire distribution below 2°C, for instance. Therefore, under uncertainty climate targets should rather be interpreted as safety constraints on the probability of crossing a certain threshold such as 2°C. Such probabilistic targets have been studied amongst others in den Elzen and Meinshausen (2005), Meinshausen et al. (2006, 2009), den Elzen and van Vuuren (2007), den Elzen et al. (2007), Keppo et al. (2007), Rive et al. (2007), Schaeffer et al. (2008).

The uncertainty surrounding climate change will at least partly be resolved in the future, which is called "learning". Uncertainty about climate sensitivity, for instance, will be reduced by future advances in climate science. This will change the probability of crossing a certain threshold for a given policy. But it will also allow to adjust climate policy. Since there are irreversibilities and inertia both in the climate system and the economy, it is not only important to adapt to new information but also to choose an anticipating near-term climate policy that provides flexibility to adapt to future information. There is an extensive literature on whether such a policy is more or less stringent. For an overview of the theoretical and the integrated assessment literature see Lange and Treich (2008) and Webster et al. (2008), respectively.

Cost-effectiveness analysis (CEA) determines climate policies that reach a given climate target at minimum costs. It takes targets as (politically) given and does not answer the question of what an optimal target should be in the light of the available information. In Section 2 we highlight that CEA for fixed probabilistic targets leads to major conceptual problems if learning is taken into account. Therefore, and because it is presumably part of current policy proposals anyway, we have to take into account that targets will be adjusted to new information. This demands formalizing how targets are determined based on the available information and by balancing costs and benefits in a broad sense. This is discussed in Section 3. Hence, the condensed message of this letter is that learning is an important part of the climate problem, and that if learning is taken into account, it is not a viable option to just perform CEA for a given climate target but necessary to formalize how targets should be determined.

More precisely, in Section 2 we highlight that a decision maker performing CEA for a fixed probabilistic target might be worse off with learning than without and consequently reject to learn. Furthermore, we show that she can also be unable to meet even the probabilistic interpretation of her target due to learning. We do this by using results from the literature on decision making under uncertainty and a simple example. Both problems are strong arguments for not using CEA for probabilistic targets if learning is considered.

In Section 3, we discuss ways to take the adjustment of targets to new information into account. One way is a full-fledged cost-benefit analysis (CBA) including a monetary climate damage function. CBA applied to the climate problem has numerous detractors. A main point of criticism is that CBA "conceal[s] ethical dilemmas" (Azar and Lindgren 2003) and difficult value, equity, and subjective probability judgments concerning climate impacts. Alternative approaches based on a precautionary or sustainability principle in turn do not have a clear formalization. As a middle ground, we explore multi-criteria decision analysis based on a trade-off between aggregate mitigation costs and a climate target based risk metric such as the probability of crossing the target threshold.

2 Fixed targets

Exemplarily, we consider a temperature target and uncertainty about climate sensitivity denoted by θ, but analogous results hold for any probabilistic target. The target consists of a temperature threshold Z of global warming, e.g. $Z = 2°C$, and a maximum acceptable threshold exceedance probability Q. We will also call the threshold exceedance probability the "risk" and Q the "risk tolerance". We denote the vector of greenhouse gas emissions over time by $E(t)$, the resulting temperature trajectory by $[T(E, \theta)](t)$, and aggregate mitigation costs not including any climate damages by $C(E)$. $C(E)$ can also be a utility function of costs. The risk as a functional of emissions is given by $R(E) = \int d\theta f(\theta)\Theta(T_{\max}(E, \theta) - Z)$, where $f(\theta)$ is the probability density function, $\Theta(\cdot)$ is Heaviside's step-function, and $T_{\max}(E, \theta) = \max_t[T(E, \theta)](t)$ is the maximum temperature. If yet nothing is learned about the uncertainty, CEA for the probabilistic target reads as

$$\begin{aligned} &\min_E C(E), \\ &s.t. \;\; R(E) \leq Q. \end{aligned} \tag{1}$$

Costs are minimized such that the probability of crossing the threshold, or the risk, is no larger than Q. Due to the constraint on a probability, such a problem is called a chance constrained programming (CCP) problem (Charnes and Cooper 1959). For an extensive numerical investigation of this problem see Held et al. (2009). The equivalence to Value-at-Risk constrained problems is shown in Section 1 of the Supplement.

In order to include learning, we consider a simple so called act-learn-act framework. That means the decision maker first decides on emissions before learning, denoted by $E_1(t)$, $t \leq t_l$. At time t_l with probability q_m she receives a signal or message m that is correlated with θ, and she updates her prior probability distribution $f(\theta)$ and risk metric $R(E)$ to a posterior distribution $f(\theta|m)$ and risk $R_m(E) = \int d\theta f(\theta|m)\Theta(T_{\max}(E, \theta) - Z)$ according to Bayes' rule. Subsequently she decides on emissions after learning, denoted by $E_m(t)$, $t > t_l$, which in general depend on the message that has been received. A dynamic extension of CCP then reads as

$$\begin{aligned} &\min_{E_1} \left\{ \sum_{m \in M} q_m \min_{E_m} \{C(E_1, E_m)\} \right\}, \\ &s.t. \;\; R_m(E_1, E_m) \leq Q, \; \forall m \in M, \end{aligned} \tag{2}$$

Hence, expected costs are minimized such that the posterior probability of crossing the threshold is no larger than Q for all messages m. Equation 2 is not the only way to extend CCP to an act-learn-act framework. An alternative formulation is obtained by constraining the expected value of the probability of crossing the threshold across all messages, i.e. $\sum_{m \in M} q_m R_m(E_1, E_m) \leq Q$. This alternative is also discussed below.

A similar problem to Eq. 2 was studied in O'Neill et al. (2006). For the special case of $Q = 0$, where the target has to be met with certainty, it was studied in Webster et al. (2008), Johansson et al. (2008), and Bosetti et al. (2009). $Q = 0$ is problematic because it is likely to be infeasible if the upper tail of the probability distribution of climate sensitivity is taken into account. Schaeffer et al. (2008), for instance, report a non-zero probability of crossing 2°C even if greenhouse-gas concentrations were stabilized at current levels. And even if $Q = 0$ were feasible, it would lead to very high mitigation costs and arguably does not correspond to current target proposals. Webster et al. (2008), for instance, report a cost-effective carbon tax of more than \$250/ton from 2040 on for the 2°C target.

For $Q \neq 0$, i.e. if the threshold doesn't have to be avoided with certainty, CCP as in Eq. 2 leads to conceptual problems. A decision maker performing CCP can be worse off with learning than without, and therefore reject to learn if possible. Most people would say this is unacceptable for a normative decision criterion, better information should be valuable. The benefits from learning can be measured by the expected value of information, $\text{EVOI} = \sum_{m \in M} q_m C(E_1^l, E_m^l) - C(E^{nl})$, where E_1^l, E_m^l and E^{nl} are optimal emissions before, after, and without learning, respectively. Hence, the EVOI is simply the difference in expected costs (or utility) between the case with and the case without learning. The possibility of a negative EVOI in CCP was first noted by Blau (1974) for a linear program and clarified in Hogan et al. (1981, 1984). Details of these papers were criticized by Charnes and Cooper (1975, 1983), but a rigorous analysis confirming the problem has been provided by LaValle (1986). In Section 2 of the Supplement, we show that CCP violates the independence axiom of von Neumann and Morgenstern, and we cite results that show that this necessarily leads to the possibility of a rejection of learning.

Here we construct a simple example for providing an intuition why the EVOI can be negative. We assume that climate sensitivity θ can take only three values with equal probability, $\theta = 2, 3, 4$°C. We also assume that if the threshold is avoided for a certain value of climate sensitivity, it is also avoided for all lower values. Finally, we assume $Q = 50\%$. We now compare the case without learning with the case of immediate perfect learning where the true value of θ is revealed at $t_l = 0$, i.e. before any decisions have to be made. The case of partial learning, where the posterior distributions are non-degenerate, is discussed in Section 3 of the Supplement. There are three policy options: Stay below the threshold for (I) only $\theta = 2$°C, (II) $\theta = 3$°C (and hence also $\theta = 2$°C), (III) $\theta = 4$°C. (I) is the cheapest and least stringent, (III) the most expensive and stringent alternative. Without learning, policy (II) is the cheapest alternative with admissible risk of 1/3. With learning, the choice depends on the true value of θ. If $\theta = 2$°C, (I) is the cheapest admissible alternative, if $\theta = 3$°C it is (II), and if $\theta = 4$°C it is (III). We have $\text{EVOI} = (1/3\, C(I) + 1/3\, C(II) + 1/3\, C(III)) - C(II)$. It is negative if abatement costs are sufficiently convex in emissions reductions so that $C(I) + C(III) > 2C(II)$.

We have argued that climate targets under uncertainty probably cannot or should not be met with certainty. A second conceptual problem is that if learning is taken into account, even the resulting probabilistic targets can generally not be met. This was first noted for a generic linear CCP problem by Eisner et al. (1971, they call Eq. 2 "conditional-go approach"). If, for instance, the threshold could not be avoided for $\theta = 4$°C in our simple example, it would be possible to limit the probability of crossing the threshold to 50% without learning but not in the "bad" learning case where $\theta = 4$°C is revealed as the true value. More generally, under perfect learning any probabilistic target with a threshold that cannot be avoided with certainty in the prior becomes infeasible. Perfect learning is not a bad approximation in the long run, and, as mentioned before, most thresholds such as 2°C arguably cannot be avoided with certainty given current information. If the probabilistic target is infeasible in some learning cases, it is unclear how to perform CCP. Infeasibility could be avoided by relaxing the target threshold from 2°C to 3 or 4°C, for instance. But the problem of a negative EVOI would persist as long as a chance constraint is applied. Besides, it would mean that the 2°C target can not be considered, which is problematic in itself.

Intuitively, what drives the results above is (i) the fact that the set of feasible (or target complying) emissions trajectories changes depending on what is learned and (ii) that the benefits of target compliance are not taken into account in the objective function. If the optimal policy without learning, i.e. (II) in our example, were feasible in all learning cases, neither infeasibility due to learning nor a negative EVOI would be possible. The latter is because choosing (II) in all learning cases would guarantee the same expected costs as without learning. And if sufficient benefits and not only the costs of choosing (III) instead of (II) if $\theta = 4°C$ is revealed were taken into account in the objective function, the EVOI would be positive despite a change in the set of feasible trajectories. In Section 3 we discuss how to include the benefits in the objective function.

The feasible emissions trajectories change because the probabilistic target is fixed and independent of what is learned and because the corresponding chance constraint was put on each individual posterior distribution. As mentioned before, CCP in an act-learn act framework could alternatively be formulated with a constraint only on the expected value of the probability of crossing the threshold across the different learning cases. Eisner et al. (1971) call this a "total probability constraint", and LaValle (1986) an "ex ante constraint". In this formulation the same trajectories are feasible with learning as without and the problems do not occur (see also LaValle 1986). But specifically this would mean that not reducing emissions at all if $\theta = 4°C$ is learned and staying below the threshold in the other two learning cases would be an admissible strategy. The expected probability of crossing the threshold would only be 1/3. I would also be the cheapest feasible strategy because it implies the least emissions reductions. It is a questionable recommendation, though, not to reduce emissions at all after learning $\theta = 4°C$ only because the probability of crossing the target would have been zero if something else had been learned. In decision theory it would be called a violation of consequentialism (e.g. Machina 1989).

The problems of CCP are known since the 1970s, and CCP is still widely used in many different areas from aquifer remediation design (Morgan et al. 1993) to air quality management (Watanabe and Ellis 1993). If learning about uncertainty and adjustment to new information can safely be neglected for a given problem, then CCP can be a satisfactory and intuitive decision criterion under uncertainty. This is the case if either little is learned, or if the EVOI is not of interest and the system is flexible enough so that anticipation of learning is not important. In the climate problem, though, learning and system inertia play an important role and should be taken into account in determining climate policy. Therefore, CCP, in our view, is not a suitable option.

3 Adjusting targets

In the preceding section we held the probabilistic target, i.e. the temperature threshold Z and the risk tolerance Q, fixed and independent of what is learned, and we did not include any benefits from target compliance in the objective function.

Current policy proposals, such as the 2°C target arguably assume that targets will be adjusted to new information in the future. The Copenhagen Accord explicitly mentions the "consideration of strengthening the long-term goal referencing various matters presented by the science" (UNFCCC 2009). In this section we discuss how to adjust targets and how to avoid the problems of CCP by including the benefits of target compliance in the objective function and by balancing costs and benefits in a broad sense.

One possibility is to assume that climate targets and optimal climate policy can be derived by a full-fledged CBA including a monetary climate damage function. As mentioned in the introduction, this kind of CBA has numerous critics. One of their main points is that by combining all damages in a monetary damage function, including loss of life, biodiversity, and the damages resulting from the highly uncertain disintegration of the West Antarctic Ice Sheet, for instance, CBA rather "conceals[s] ethical dilemmas" (Azar and Lindgren 2003) and difficult value, equity, and subjective probability judgments than highlighting them to decision makers (see the discussion in Azar and Lindgren 2003). Besides, it would be useful to have a decision criterion that is at least to some extend based on politically given climate targets.

As a consequence of the problems of CCP, Bordley and Pollock (2009) suggest in an engineering context to specify an additional target threshold for the costs and then to minimize the probability of crossing either threshold. Jagannathan (1985) uses a simple trade-off between costs and threshold exceedance probability in order to avoid a negative EVOI. Applied to the climate context, a linear form reads as

$$\min_{E_1}\Big\{ \textstyle\sum_{m\in M} q_m \min_{E_m}\big\{wC\,(E_1,\,E_m) + R_m\,(E_1,\,E_m)\,\big\}\Big\}, \tag{3}$$

The normative parameter w determines the trade-off between costs and risk. It equals the per centage points of risk increase that would be accepted in exchange for a unit decrease in costs. We will call Eq. 3 cost-risk analysis (CRA). CRA can be seen as a weighted multi-criterion decision analysis or also as a CBA in a broader sense. In contrast to CCP, the benefits, namely the reduction of risk, are now included in the objective function. The trade-off is assumed to be linear in order to have an equivalence to the expected utility maximization $\max_{E_1}\{\sum_{m\in M} q_m \max_{E_m}\{\int f(\theta|m)A\,(E_1,\,E_m,\theta)\}\}$ with $A(E_1,\,E_m,\theta) = -(wC(E_m) + \Theta(T_{\max}(E_1,\,E_m,\theta) - Z))$. The conceptual problems encountered for CCP therefore cannot occur (see also Section 2 of the Supplement). Jagannathan (1987) suggests to consider non-linear trade-offs as well, but we could not find a convincing non-linear form of the trade-off that is still equivalent to an expected utility maximization (see also LaValle 1987).

Mastrandrea and Schneider (2004, 2005) develop a risk management framework based on the probability of exceeding a threshold of "dangerous anthropogenic interference" (UNFCCC 1992) as risk metric. But they only report different risk levels for different stabilization targets and do not formalize the final trade-off between costs and risk, which becomes necessary if learning is included in the analysis. This could be done in CRA. Schneider and Mastrandrea (2005) also propose a more sophisticated risk metric that better represents the temperature path dependence of risk. It is based on the concept of maximum exceedance amplitude (MEA: by how many Kelvin the target threshold is exceeded) and the concept of degree years (DY: the area above the threshold between the temperature trajectory and the threshold). The expected value of some function Φ of MEA and DY could also be used as a risk metric in CRA, $R_m(E_m) = \int f(\theta|m)\Phi\,(\mathrm{MEA}(E_1, E_m, \theta), DY(E_1, E_m, \theta))$.

The main difference between CRA and standard CBA is that the former makes the necessary trade-offs between mitigation costs and impacts (risks) on a more aggregate level, directly in the objective function, and thereby more explicitly and to some extent based on given targets. Thus, the main difference is the framing of the decision. The main difficulty of CRA, as of most multi-criteria decision analyses, is that it is hard for decision makers to specify the value of the trade-off parameter w, i.e. to value a probability of crossing a threshold in terms of costs, for instance. But we would argue that at least for non-market and highly uncertain impacts, it might still be easier to specify and more practical than a monetary climate damage function.

More specifically, the following combination of standard CBA and CRA might better suit the climate problem than a pure CBA, CRA or CEA. Market-damages, whose value can be estimated by observing markets without significant externalities, are included over a damage function, which in turn is included in the cost metric $C(E)$. Non-market impacts like loss of life and public goods, impacts from highly uncertain climate tipping-points, as well as wider societal impacts like migration and conflict are included over an aggregate, climate target-based risk metric $R(E)$. As highlighted before, valuing these impacts is inherently difficult, and there is no way around some kind of multi criteria decision analysis. Instead of mixing the value judgments concerning these impacts with market impacts in a monetary damage function as in standard CBA, an aggregrate trade-off between a target-based risk and aggregate mitigation costs might be a more practical framing of the problem.

4 Conclusions

Climate targets such as the 2°C target probably cannot or are not supposed to be met with certainty. They should rather be interpreted as probabilistic targets. Cost-effectiveness analysis (CEA) for such targets constitutes a chance-constrained programming (CCP) problem. Transferring results from the literature to the climate context, we have highlighted that CCP can imply a negative expected value of information, which most people would consider normatively unsatisfactory. Furthermore, even a probabilistic interpretation of relevant targets, such as the 2° target, becomes infeasible if learning is taken into account, so that it is unclear how to perform CCP at all. Consequently, and because it is arguably part of the current target proposals, we have discussed how to avoid the problems by adjusting climate targets to new information and by balancing benefits and costs in a broad sense. A prominent way to do this is cost-benefit analysis (CBA) including a monetary climate damage function. But specifying such a damage function is notoriously difficult and controversial. We took the problems of both CBA and CEA as motivation for asking, whether there is a middle-ground between a full-fledged CBA and CEA. Partly based on previous suggestions in the literature, we discussed a combination of a damage function for market impacts and a more aggregate target-based risk metric for non-market and highly uncertain catastrophic impacts as a promising candidate.

Acknowledgements M.G.W.S. was supported by the EU project GEO-BENE (No. 037063). E.K. was supported by a Marie Curie International Fellowship (MOIF-CT-2005-008758) within the 6th European Community Framework Programme. Helpful comments by two anonymous reviewers are acknowledged.

References

Azar C, Lindgren K (2003) Editorial commentary: catastrophic events and stochastic cost-benefit analysis of climate change. Clim Change 56(3)

Blau RA (1974) Stochastic programming and decision analysis: an apparent dilemma. Manage Sci 21(3):271–276

Bordley RF, Pollock SM (2009) A decision-analytic approach to reliability-based design optimization. Oper Res 57(5):1262–1270

Bosetti V, Carraro C, Sgobbi A, Tavoni M (2009) Delayed action and uncertain stabilisation targets: how much will the delay cost? Clim Change 96(3):299–312

Charnes A, Cooper WW (1959) Chance constrained programming. Manage Sci 5:73–79

Charnes A, Cooper WW (1975) A comment on Blau's dilemma in stochastic programming and bayesian decision analysis. Manage Sci 22(4):498–500

Charnes A, Cooper WW (1983) Response to "Decision problems under risk and chance constrained programming: dilemmas in the transition". Manage Sci 29(6):750–753

den Elzen MGJ, Meinshausen M (2005) Meeting the EU 2°C climate target: global and regional emission implications. Report 728001031:2005

den Elzen MGJ, Meinshausen M, van Vuuren DP (2007) Multi-gas emission envelopes to meet greenhouse gas concentration targets: costs versus certainty of limiting temperature increase. Glob Environ Change Human Policy Dimensions 17(2):260–280

den Elzen MGJ, van Vuuren DP (2007) Peaking profiles for achieving long-term temperature targets with more likelihood at lower costs. Proc Natl Acad Sci 104(46):17931

Eisner MJ, Kaplan RS, Soden JV (1971) Admissible decision rules for the E-model of chance-constrained programming. Manage Sci 17(5):337–353

European Council (2005) Presidency conclusions. European Council, Brussels

Held H, Kriegler E, Lessmann K, Edenhofer O (2009) Efficient climate policies under technology and climate uncertainty. Energy Econ 31:S50–S61

Hogan AJ, Morris JG, Thompson HE (1981) Decision problems under risk and chance constrained programming: dilemmas in the transition. Manage Sci 27(6):698–716

Hogan AJ, Morris JG, Thompson HE (1984) Reply to professors charnes and cooper concerning their response response to "Decision problems under risk and chance constrained programming". Manage Sci 30(2):258–259

Jagannathan R (1985) Use of sample information in stochastic recourse and chance-constrained programming models. Manage Sci 31(1):96–108

Jagannathan R (1987) Response to 'On the "bayesability" of chance-constrained programming problems' by Lavalle Manage Sci 33:1229–1231

Johansson DJA, Persson UM, Azar C (2008) Uncertainty and learning: implications for the trade-off between short-lived and long-lived greenhouse gases. Clim Change 88(3–4):293–308. ISSN 0165-0009

Keppo K, O' BC, Riahi K (2007) Probabilistic temperature change projections and energy system implications of greenhouse gas emission scenarios. Technol Forecast Soc Choice 74(7):936–961. ISSN 0040-1625

Lange A, Treich N (2008) Uncertainty, learning and ambiguity in economic models on climate policy: some classical results and new directions. Clim Change 89(1):7–21

LaValle IH (1986) On information augmented chance-constrained programs. Oper Res Lett 4(5):225–230

LaValle IH (1987) Response to "Use of sample information in stochastic recourse and chance-constrained programming models": on the 'bayesability' of CCP's. Manage Sci 33(10):1224–1228

Machina MJ (1989). Dynamic consistency and non-expected utility models of choice under uncertainty. J Econ Lit 27(4):1622–1668

Mastrandrea MD, Schneider SH (2004) Probabilistic integrated assessment of "dangerous" climate change. Science 304:571–575

Meinshausen M, Hare B, Wigley TML, Van Vuuren D, Den Elzen MGJ, Swart R (2006) Multi-gas emissions pathways to meet climate targets. Clim Change 75(1–2):151–194

Meinshausen M, Meinshausen N, Hare W, Raper SCB, Frieler K, Knutti R, Frame DJ, Allen MR (2009) Greenhouse-gas emission targets for limiting global warming to 2°C. Nature 458(7242): 1158–1162

Morgan DR, Eheart JW, Valocchi AJ (1993) Aquifer remediation design under uncertainty using a new chance contrained programming technique. Water Resour Res 29(3):551–569

O'Neill B, Ermoliev Y, Ermolieva T (2006) Endogenous risks and learning in climate change decision analysis. Springer, Berlin, Germany, pp 283–300

Rive N, Torvanger A, Berntsen T, Kallbekken S (2007) To what extent can a long-term temperature target guide near-term climate change commitments? Clim Change 82(3):373–391

Schaeffer M, Kram T, Meinshausen M, van Vuuren DP, Hare WL (2008) Near-linear cost increase to reduce climate-change risk. Proc Natl Acad Sci USA 105(52):20621–20626

Schneider SH, Mastrandrea MD (2005) Probabilistic assessment of "dangerous" climate change and emissions pathways. Proc Natl Acad Sci USA 102(44):15728–15735

UNFCCC (1992) Article 2. http://unfccc.int/essential_background/convention/background/items/1353.php. Accessed 12 November 2010

UNFCCC (2009) Copenhagen accord. http://unfccc.int/documentation/documents/advanced_search/items/3594.php?rec=j&priref=600005735#beg. Accessed 12 November 2010

Watanabe T, Ellis H (1993) Stochastic programming models for air quality management. Comput Oper Res 20(6):651–663

Webster M, Jakobovits L, Norton J (2008) Learning about climate change and implications for near-term policy. Clim Change 89(1):67–85

Climate Targets under Uncertainty: Challenges and Remedies

Supplement

Matthias G.W. Schmidt · Alexander Lorenz · Hermann Held · Elmar Kriegler

1 Value-at-Risk

A probabilistic target is essentially equivalent to a limit on the Value-at-Risk (VaR) in finance. The $x\%$-VaR, or VaR at the $x\%$ confidence level, of a financial position equals the x-percentile of the distribution of the uncertain losses of the position. In other words, with $x\%$ certainty, losses will be smaller than the $x\%$-VaR. Hence, we can formulate the probabilistic target as a constraint on the VaR in the distribution of maximum temperature: The $(1-Q)$-VaR has to be smaller or equal than a given threshold, such as 2°C.

2 Violation of the Independence Axiom

We shortly introduce some basic decision theoretic terminology and formulate CCP as a preference relation on simple lotteries. Subsequently, we show that CCP does not fulfill the independence axiom by von Neumann and Morgenstern. There is an extensive literature on the consequences of relaxing the axioms of von Neumann and Morgenstern. We shortly review one result that shows that the possibility of a rejection of learning encountered in the main text follows from violation of the independence axiom.

A simple lottery describes an uncertain outcome. It is defined by the set of possible outcomes with their respective objective or subjective probability. For the climate example without learning every emissions path can be assigned a simple lottery. This lottery is defined by the vector of relevant outcomes, here maximum temperature and mitigation costs, and the probability (density) for these outcomes. So we denote lotteries by $L_{E,f} := \{(T_{\max}(E,\theta),\, C(E))\,,\, f(\theta)\}$. In a mixed lottery, the outcomes of a first stage lottery are again lotteries. We denote the mixture of two lotteries L_1 and L_2 with mixing probability β by $\beta L_1 + (1-\beta)L_2$.

The ordering of simple lotteries implied by CCP as in Eq. (1) in the main text is akin to a lexicographic ordering. Lexicographic orderings consist of a hierarchy of orderings like a lexicon: words with the same first letter are ordered according to the second letter and so on. The primary ordering ($\succ_1$) in CCP is according to whether the probabilistic target is met or not. It strictly prefers all emissions plans that meet the target over

Matthias G.W. Schmidt · Alexander Lorenz · Elmar Kriegler
Potsdam Institute for Climate Impact Research
Telegraphenberg A31, 14473 Potsdam, Germany
E-mail: schmidt@pik-potsdam.de, Tel. +49-311-2882566

Hermann Held
University of Hamburg - KlimaCampus
Bundesstr. 55, 20146 Hamburg, Germany

plans that do not ($L_1 \succ_1 L_2 \Leftrightarrow (R(L_1) \leq Q) \wedge (R(L_2) > Q)$, where $\wedge$ is the logical AND). But unlike for typical lexicographic orderings, the primary ordering in CCP does not lend itself to a definition of indifference as "none of the two lotteries is strictly preferred to the other" ($L_1 \simeq_1 L_2 \Leftrightarrow (\neg(L_1 \succ_1 L_2)) \wedge (\neg(L_2 \succ_1 L_1))$, where $\neg$ is the logical NOT). Such a definition would imply indifference between all emissions plans that meet the target and all of those that do not. When applying the secondary ordering in CCP, i.e. preference of the less costly plan over the more costly one ($L_1 \succ_2 L_2 \Leftrightarrow C(L_1) < C(L_2)$), to these two indifference classes, it will produce a sensible ordering of the plans that meet the target. But it would identify the business-as-usual case with zero emissions reductions as the preferred strategy among those that miss the target. This would be clearly unsatisfactory. In this sense CCP preferences can be regarded as incomplete and indifference in the primary ordering is limited to plans that meet the target ($L_1 \simeq_1 L_2 \Leftrightarrow (R(L_1) \leq Q) \wedge (R(L_2) \leq Q)$. The primary and secondary ordering in CCP allow differentiating between plans meeting and violating the target, and between plans that all meet the target, but not between plans that all miss the target. Alternatively to having an incomplete primary ordering, one could assume indifference in the primary ordering between plans that don't meet the target and apply a different, more satisfactory secondary ordering than cost minimization to these plans. However, the incompleteness is not necessarily problematic, because it still allows for the formulation of an overall preference relation

$$L_1 \succ L_2 \quad \Leftrightarrow (L_1 \succ_1 L_2) \vee ((L_1 \simeq_1 L_2) \wedge (L_1 \succ_2 L_2)) \tag{1}$$

that has the desirable properties of asymmetry ($L_1 \succ L_2 \Rightarrow \neg(L_2 \succ L_1)$) and negative transitivity ($\neg(L_1 \succ L_2) \wedge \neg(L_2 \succ L_3) \Rightarrow \neg(L_1 \succ L_3)$ (Kreps, 1988)). In particular, it allows identifying a choice set of most preferred strategies. However, the target infeasibility due to learning discussed in the main text has shown that the choice set of CCP can become useless, i.e. will indiscriminately include all available strategies, if no strategy can meet the target.

CCP as in Eq. (1) violates both the continuity and the independence axiom by von Neumann and Morgenstern. We only discuss the latter here. Independence is violated because the chance constraint cannot be formulated as a set of separate, or independent, constraints for each state of the world. The avoidance of the threshold in one state of the world, via the chance constraint, has an influence on the need to avoid the threshold in other states of the world. More formally, independence would be fulfilled if for any three lotteries L_1, L_2, L_3 and for all $\beta \in (0, 1]$ we had

$$L_1 \succ L_2 \Rightarrow \{\beta L_1 + (1-\beta)L_3 \succ \beta L_2 + (1-\beta)L_3\} \tag{2}$$

So independence means that the preferences are not changed by mixing the same lottery L_3 into two given lotteries L_1 and L_2. This is not the case for CCP because of the primary ordering according to the chance constraint. E.g., it is possible that $R(L_2) < R(L_1) < Q < R(L_3)$, $C(L_2) > C(L_1)$ and $(\beta R(L_2) + (1 - \beta)R(L_3)) < Q < (\beta R(L_1) + (1-\beta)R(L_3))$, i.e. both L_1 and L_2 fulfill the chance constraint but L_2 is less risky and gives higher costs than L_1. L_3 does not fulfill the constraint and β is chosen such that the mixed lottery of L_2 and L_3 fulfills the constraint, whereas the mixed lottery of L_1 and L_3 does not. We then have $L_1 \succ L_2$ and $\beta L_1 + (1-\beta)L_3 \prec \beta L_2 + (1-\beta)L_3$, which shows non-independence.

The possibility of a decision maker being worse off with learning than without that we encountered in the main text follows from violation of the independence axiom. Wakker (1988) proves the following consequence:

$$\neg\text{Independence} \wedge \text{Correct anticipation of future decisions} \wedge \text{Consequentialism}$$
$$\Rightarrow \text{Information can make decision maker worse off.} \tag{3}$$

So if the antecedents are fulfilled including violation of the independence axiom, then the receipt of additional information can make the decision maker worse off. We have already shown that CCP violates the independence axiom. Future decisions are also anticipated correctly. It is correctly anticipated that after receipt of a message, the target will have to be met based on the updated posterior information. More critical is the assumption of consequentialism. Consequentialism intuitively means that only current and future payoffs have an influence on current decisions. Past outcomes and foregone options have no influence on current decisions. CCP as in Eq. (2) of the main text is consequentialist because the chance constraint is applied to every single posterior, and forgone risk in other learning cases is not taken into account.

3 Partial Learning

One might object to the simple example in the main text that the EVOI only becomes negative because we consider perfect learning. Under perfect learning the posterior risk has to be reduced to zero and not only 50%. So the target stringency is effectively increased by learning. But firstly, perfect learning is probably not unrealistic in the long run, so the decision criterion should be able to handle it. Secondly, the same problems occur for partial learning, where the uncertainty is only reduced from a prior to a non-degenerate posterior distribution. Consider the prior and posterior distributions shown in Fig. 1. If maximum temperature is monotonic in climate sensitivity θ, i.e. if the target is met for θ_1 it is met for all $\theta \leq \theta_1$, then we can translate the risk tolerance into a maximum value of θ, for which the target threshold has to be avoided. This value, of course, depends on what is learned. It decreases from about 3°C to about 2°C in the "good" learning case (posterior 1) and increases to about 4°C in the bad case (posterior 3). These are the same values for θ as in the perfect learning example in the main text. Hence, we would get the same negative EVOI.

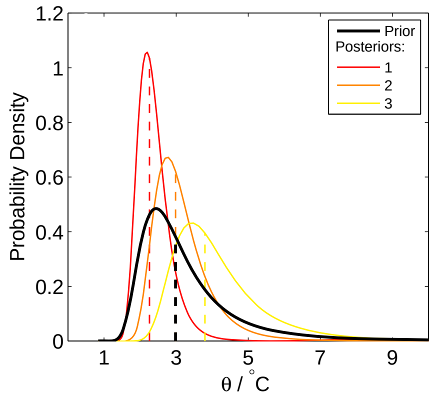

Fig. 1 Information structure with prior distribution and three posterior distributions. The vertical dashed lines indicate the medians.

References

D. Kreps. *Notes on the theory of choice*. Westview Press, 1988.

P. Wakker. Nonexpected utility as aversion of information. *Journal of Behavioral Decision Making*, 1:169–175, 1988.

Chapter 3

How important is Uncertainty for the Integrated Assessment of Climate Change? *

Alexander Lorenz
Elmar Kriegler
Hermann Held
Matthias G.W. Schmidt

*This chapter has been submitted to *Climate Change Economics* as Lorenz, A., E. Kriegler, H. Held, M.G.W. Schmidt. How important is Uncertainty for the Integrated Assessment of Climate Change?

How important is Uncertainty for the Integrated Assessment of Climate Change?

Alexander Lorenz*†, Elmar Kriegler*, Hermann Held‡, Matthias G.W. Schmidt*

Abstract

We investigate the importance of explicitly accounting for uncertainty in the assessment of optimal global climate policy. The benefit of accounting for uncertainty can be evaluated by comparing the outcome of a best guess policy in a deterministic setting with an expected value maximizing policy under uncertainty. We apply the approach to the case of uncertainty about the temperature response to greenhouse gas emissions and climate induced damage in the integrated assessment model MIND. In accordance with the literature we find that the welfare gain from adjusting climate policy to uncertainty is negligible. We use a decomposition of the uncertainty adjustment effect to explain its negligible magnitude in the standard setting as a result of compensating factors in the climate cause-effect chain from emissions to damage to welfare implications. We demonstrate several changes in the model structure, such as exponential climate damage and linear climate-carbon response that can lead to a significant welfare gain from explicitly accounting for uncertainty in climate policy analysis.

1 Introduction

Global Climate Change Analysis is surrounded by large uncertainties about key parameters in the socio-economic system and the climate system. The uncertainties arise from imperfect knowledge about the dynamics of the subsystems, from internal short term dynamics or stochasticity and from the long time lag between cause and effect within the climate system (Tol, 1999). Analysts use models that integrate the socio-economic system and the climate system to assess long term strategies to mitigate climate change (Nordhaus,

*A. Lorenz (corresponding author), Elmar Kriegler, Matthias G.W. Schmidt
Potsdam Institute for Climate Impact Research, 14412 Potsdam, Germany
E-mail: lorenz@pik-potsdam.de Tel.: +49-331-2882562
†now at
Environmental Change Institute, School of Geography and the Environment
University of Oxford, South Parks Road, Oxford, OX1 3QY
E-mail: alexander.lorenz@ouce.ox.ac.uk
Tel.:+44-1865-275894
‡H. Held
University of Hamburg & KlimaCampus Hamburg, Bundesstr. 55, 20146 Hamburg, Germany
also at
Potsdam Institute for Climate Impact Research, 14412 Potsdam, Germany

1994). A key question concerning uncertainty and climate change assessment is whether the analysts should explicitly account for uncertainty in their integrated assessment models to capture the effect of uncertainty on decisions? An initial intuition would be, that under uncertainty about the climate response to greenhouse gas emissions and resulting climate damage the risk of very high damage beyond the mean damage is introduced. Assuming a right-skewed damage distribution and some degree of risk aversion, one would be willing to undergo stronger emission reductions than in the best-guess case to better insure against high damage scenarios. Insuring against those high damage scenarios would also become more valuable than simply reducing average climate damage, thus we would expect climate policy to become more valuable under uncertainty compared to the best-guess case.

Over the last two decades following the pioneering work of Nordhaus [1994], many contributions have been made to check the intuition described above, and to answer the question how explicitly accounting for uncertainty changes the optimal climate policy (e.g. Keller et al., 2004; Wirl, 2007; Heal & Kristrï¿œm, 2002, and references therein) . Some studies (e.g. Pizer, 1997) also investigated the welfare effect from introducing uncertainty. A common result has emerged from cost-benefit analyzes with integrated assessment models that optimize the trade-off between mitigation costs of reducing greenhouse gas emissions and climate change induced damage costs (e.g. see Nas, 1996): Although the optimal climate policy might change significantly due to the introduction of uncertainty, the welfare gain associated with adjusting the policy to uncertainty is negligible (Pizer, 1997). One exception is the work of Weitzman [2010] who shows that the effect of uncertainty can become significant and even dominating if fat tailed probability distributions for the climate damage are considered. Under certain conditions, such fat tailed climate damage can lead to unbounded expected welfare losses.

Parts of the uncertainties about the climate system and climate change induced damage may be resolved by future learning, both passively via observing the state of the coupled socio-economic-climate system (Kelly & Kolstad, 1999) and actively via investing into new measurements and leveraging paleo-information (Lorenz et al., 2009). This would allow to adapt mitigation policy to new information. The welfare implications of the possibility to learn and adjust to new information has also gained much attention (e.g. see Nordhaus & Popp, 1997; Karp & Zhang, 2006; Ingham et al., 2007). Here a general finding is that adapting climate policy to new information can lead to significant welfare gains.

What is the explanation for the negligible effect of uncertainty in cost-benefit analyzes of climate policy? As pointed out in the literature, uncertainty has no effect on decisions and welfare if the model is (nearly) linear in the uncertain parameters (see e.g. Baker, 2006; Lange & Treich, 2008). Furthermore the marginal utility of the model also needs to be non-linear in the uncertain parameter for the optimal decisions to be effected, as they are determined by the trade-off between marginal benefits and costs. For the importance of future learning for short-term actions Webster [2002] has argued that so-called "cross-period interactions" in utility are a necessary condition for future learning to affect short-term decisions. However the existence of such a dependency of future utility on today's decisions is not sufficient for a strong influence of future learning, as different cross-period interactions could lead to contradicting influences from future learning on today's decisions that might cancel each other out. However these theoretical findings do not explain why uncertainty has nearly no effect on welfare within the more complex integrated assessment models, as those models do incorporate a number of (partly strong) nonlinearities as well as cross-period interactions: the welfare function, the temperature response to carbon emissions, the damage function, etc.

This study presents new insights into the source of the negligible uncertainty effect. Unlike previous studies, we go beyond general findings on sufficient conditions for a negligible uncertainty effect. Since such conditions are not fulfilled by integrated assessment models, the magnitude of the uncertainty adjustment effect becomes an empirical question. We present a method to analyze the importance of uncertainty by tracing it through the cause-effect chain from greenhouse gas emissions to temperature change to induced climate damage to welfare implications. This allows us to explain the negligible uncertainty effect as a result of compensating factors in the cause-effect chain.

More concretely, we introduce a decomposition of the overall benefit of climate policy into single components: The benefit of the best guess policy, the re-evaluation of the best guess policy under uncertainty, and the value of adapting the optimal policy to uncertainty. We extend the decomposition to the case of perfect learning and analyze the relative importance of all components in the MIND model. Additionally we project the complex integrated assessment model to an a-temporal marginal cost-benefit picture. This allows us to connect the different components of the overall benefit of climate policy to the functional form of the marginal benefits and costs.

Applying this methodology to the integrated assessment model MIND we can identify the following reasons for the negligible effect of accounting for uncertainty in the model: First, the overall benefit of climate policy is constrained by the saturation of the emissions to temperature change relationship compensating for the non-linearity in the climate damage function and by the consumption smoothing property of the welfare function. Second, this assessment does not change significantly under uncertainty because the additional marginal welfare benefit of reducing a unit of emissions becomes only significant for large mitigation effort, where mitigation costs are already high. Third, the welfare gains from adjusting the mitigation policy under uncertainty are limited because of the strongly increasing mitigation costs.

Using this understanding of the relationship between the model formulation, the shape of the marginal benefits and costs, and the components of the benefit of climate policy, we introduce several changes in the model structure to create a significant welfare effect from uncertainty. These changes include a sensitivity analysis with respect to the parameter of constant relative risk aversion, a switch towards exponential damage and an implementation of a linear climate response to cumulative carbon emissions as proposed by Metthews et al. [2009].

The paper is structured as follows: Section 2 introduces the general climate decision problem under uncertainty and (perfect) learning and describes the decomposition of the overall benefit of climate policy into the single components determining the benefit of climate policy under uncertainty. The decomposition is illustrated within a simple analytical model of climate change with quadratic costs and benefits of mitigation. In Section 3 the framework is applied to the **M**odel of **In**vestment and Technological **D**evelopment (MIND). The importance of introducing uncertainty is compared to the overall benefit of acting against climate change and to the benefits from perfect learning. The sensitivity of the results towards changes in normative parameters is evaluated. The marginal cost-benefit picture of the model as well as the functional dependency along the climate cause-effect chain is presented to investigate the origin of the negligible welfare effect from uncertainty. In Section 4 several changes to the model structure are investigated with respect to their influence on the importance of uncertainty: changes in the parameter of constant relative risk aversion, exponential damage, and linear climate carbon response. Section 5 concludes.

2 How to measure the importance of Uncertainty and Perfect Learning?

2.1 The decision problem

First, we formulate the general decision problem incorporating uncertainty and (perfect) learning in its simplest version. We only consider one decision period. The principle agent (DM) decides upon a set of decision variables x, like investments, emission control rates, etc. The decisions might also represent a whole time path of single decisions $x(t)$. Depending upon the decisions and upon the state of the world (SOW) the DM derives an overall welfare $U(x, \theta)$. Uncertainty about the SOW is represented by a probability distribution $\pi(\theta)$.

Technically speaking we model an open loop optimal control problem and thus neglect the effect of changes in available information over time. The problem now is to maximize the overall expected utility $V(x, \pi)$:

$$\max_x V(x,\pi) = \max_x \sum_j \pi(\theta_j) U(x, \theta_j) \ . \tag{1}$$

Second, we introduce the cases of the DM's information structure relevant for the climate change example. The random variable θ represents the uncertain magnitude of the temperature response to greenhouse gas emissions and of climate change induced damage. The DM's knowledge, or belief, about the values of the uncertain climate response and damage is represented by the probability distribution function $\pi(\theta)$. The general case of uncertainty is simply denoted by π. The degenerate case, where the DM is certain about

θ taking the value θ_j is defined via:

$$\pi_j \equiv \pi \begin{cases} 1 & \theta = \theta_j \\ 0 & \text{else} \end{cases} . \quad (2)$$

Two special cases of the degenerate distribution are the case of no climate damage at all, denoted by π_0, and the case of certainty about θ taking its expected value, denoted by $\bar{\pi}$:

$$\pi_0 \equiv \pi \begin{cases} 1 & \theta = 0 \\ 0 & \text{else} \end{cases} , \qquad \bar{\pi} \equiv \pi \begin{cases} 1 & \theta = \sum_j \pi(\theta_j)\theta_j \\ 0 & \text{else} \end{cases} . \quad (3)$$

Third, utilizing the decision framework and the special instances of information structure, we define the following four policy scenarios, relevant for measuring the importance of uncertainty and perfect learning:

The No-Control Case (NC): We define the policy case of "no-control" as the optimal policy in the absence of any climate damage: $\hat{x}_0 \equiv \arg\max_x V(x, \pi_0)$. The rationale behind this definition stems from the difference between the non cooperative and the cooperative solution of a decentralized market economy. In such an economy a no-control behavior is caused by the imperfect cooperation of a large number of decision makers. In a competitive setting each decision maker only anticipates her own small share of global climate induced economic damage leading to an almost total neglect of the climate problem in individual decisions. Thus the climate problem is called an externality to the market. In a fully cooperative setting the single actors would optimize their combined welfare and thus correctly anticipate global warming. Within the model, we simulate the lack of cooperation by making the DM ignorant towards climate change. Within our setting the no-control case is not only suboptimal due to the lack of mitigation efforts, but additionally the savings rate cannot be adjusted to the observed climate damage. Thus the benefit from internalizing climate damage is slightly exaggerated. However, this error is small, as the savings rate adjustment due to climate damage only becomes significant for very high levels of climate damage ($D \gg 50\%$).

The Best-Guess Case: We define the optimal policy under certainty about climate response and damage by $\hat{x}_1 \equiv \arg\max_x V(x, \bar{\pi})$. This is the common approach to take the expected value of the uncertain parameters as best guess values.

The Uncertainty Case: We define the optimal decision under uncertainty about climate response and damage by $\hat{x}_2 \equiv \arg\max_x V(x, \pi)$.

The Perfect Information Case: For comparing the importance of uncertainty and learning we consider the case of immediate perfect learning. The information about the true state of the world is revealed before any decision has been made. Any other implementation of partial or later learning leads to less benefits of information thus we consider the limiting case of what can be gained by learning about the SOW. We define the optimal decision under perfect information by $\hat{x}_3(\theta_j) \equiv \arg\max_x V(x, \pi_j)$. From the ex ante perspective there is no longer a single optimal policy, but a set of policies conditional on what has been learned before deciding. The expected utility over all possible messages that might have been revealed is given by $W \equiv \sum_k q(k)\, V(\hat{x}_3(\theta_j), \pi_{kj})$, whereby in the case of consistent learning the probability q_k of getting message π_{kj} is simply equal to the probability π_j of the revealed state θ_j.

2.2 Metrics for measuring the Importance of Uncertainty and Perfect Learning

In this section we use the nomenclature defined above to introduce metrics that measure the different components of the overall benefit of climate policy separately. Combining the above defined policy scenarios $[\hat{x}_0, \hat{x}_1, \hat{x}_2, \hat{x}_3(\theta)]$ with the possible assumptions of how the world reacts to the policy decisions, represented by information structures $[\pi_0, \bar{\pi}, \pi]$, leads to 12 possible outcomes in terms of expected utility V (or W in case of perfect learning), from which only 7 outcomes are of further interest. Those are depicted schematically in Fig.1.

The welfare differences between those cases, measured as changes in certainty and balanced growth equivalents (ΔCBGE, see Appendix 5), can be used as metrics for the importance of the different effects of uncertainty in welfare terms:

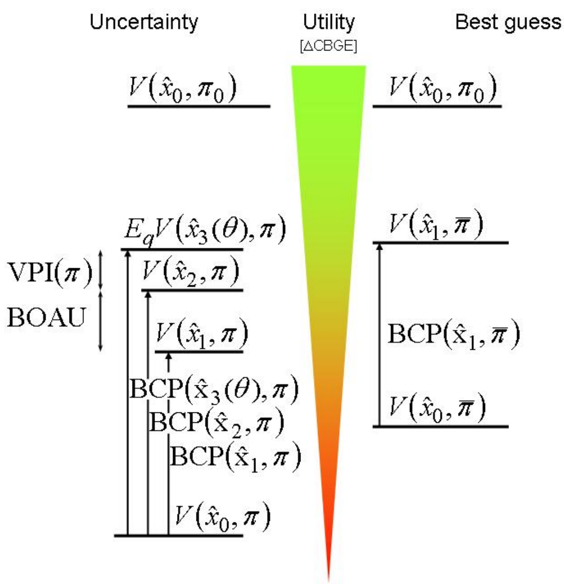

Figure 1: Welfare levels for the combinations of policy scenarios $[\hat{x}_0, \hat{x}_1, \hat{x}_2, \hat{x}_3(\theta)]$ and information structures $[\pi_0, \bar{\pi}, \pi]$ that are relevant for defining importance metrics. Also shown are relevant welfare differences, measured in ΔCBGE.

The relevant measure for the importance of climate policy in a best-guess world is the net benefit of reacting to climate change, i.e. changing from $\hat{x}_0$ to $\hat{x}_1$, $\text{BCP}(\hat{x}_1, \bar{\pi}) = \Delta\text{CBGE}\,[V(\hat{x}_1, \bar{\pi}), V(\hat{x}_0, \bar{\pi})]$. This benefit of climate policy is usually small compared to the mitigation costs within a cost-efficiency framework, as it already combines the benefits from reducing climate damage with the costs of mitigation.

Introducing uncertainty has two effects: First, the valuation of policies $\hat{x}_0$ and $\hat{x}_1$ changes. If, for all x, the (expected) utility $V(x, \bar{\pi})$ is concave (convex) in the uncertain parameter θ, the expected utility for $\hat{x}_0$ and $\hat{x}_1$ in the uncertain case (π) will be smaller (larger) than in the best-guess case ($\bar{\pi}$) (Fig. 1 shows the case of concave utility). As the benefit of climate policy $\mathrm{BCP}(x,\bar{\pi})$ is defined as difference between two levels of (expected) utility the behavior of $\mathrm{BCP}(x,\bar{\pi})$ for a switch between the best-guess and the uncertainty case depends on the curvature of the marginal (expected) utility in the uncertain parameter θ. If, for all x, the marginal (expected) utility is convex (concave) in θ, then the increase (decrease) of (expected) utility due to uncertainty in θ compared to the best guess $\bar{\theta}$ is smaller (larger) for $\hat{x}_1$ than for $\hat{x}_0$. Hence the benefit of adopting the optimal climate policy from the best guess world increases (decreases) when evaluated in the uncertain world: $\mathrm{BCP}(\hat{x}_1,\pi) > (<)\,\mathrm{BCP}(\hat{x}_1,\bar{\pi})$. The difference between the benefit from the best guess policy in the uncertain world and the certain world is denoted by $\Delta\mathrm{BCP}(\hat{x}_1) = \mathrm{BCP}(\hat{x}_1,\pi) - \mathrm{BCP}(\hat{x}_1,\bar{\pi})$.

Second, when not only evaluating the solution of the best guess world under uncertainty but explicitly maximizing the expected utility, the optimal climate policy will change (from $\hat{x}_1$ to $\hat{x}_2$). This possibility of adjusting climate policy to uncertainty leads to an increase in overall expected utility, denoted as Benefit of Adjusting to Uncertainty, $\mathrm{BOAU} \equiv \mathrm{BCP}(\hat{x}_2,\pi) - \mathrm{BCP}(\hat{x}_1,\pi)$.

Taking both effects of uncertainty into account, the overall benefit from optimally responding to climate change under uncertainty $\mathrm{BCP}(\hat{x}_2,\pi)$ can be divided into three parts:

$$\mathrm{BCP}(\hat{x}_2,\pi) \equiv \mathrm{BCP}(\hat{x}_1,\bar{\pi}) + \Delta\mathrm{BCP}(\hat{x}_1) + \mathrm{BOAU}\ . \tag{4}$$

A common measure for the "strength" of this adjustment effect is the absolute or relative change in optimal decisions itself, i.e..$\Delta\hat{x} \equiv (\hat{x}_2 - \hat{x}_1)/\hat{x}_1$ (e.g. see Tol et al., 1999). We argue that the comparison of optimal policies is insufficient to decide upon the importance of including uncertainty as even a large Δx does not necessarily has to correspond to a large BOAU. To assess the importance of uncertainty and of optimizing expected utility, we compare the single contributions of Eq. 4 normalized to their sum.[1].

Introducing (perfect) learning allows the DM to adjust her policy to the received signal. As a limiting case of early learning, the optimal policy can be chosen conditional on the perfect knowledge about the respective state of the world, $\hat{x}_3 = \hat{x}_3(\theta)$. The expectation for the overall welfare is now not taken over uncertain states of the world, but over the possible messages leading to certain states of the world. To be consistent with the ex ante knowledge, the distribution over the messages has to be identical to the distribution over the SOW in the uncertain case. The benefit of perfect learning is measured by comparing the expected utility with and without learning. The Value of Perfect Information is defined via: $\mathrm{VPI}(\pi) \equiv \mathrm{BCP}(\hat{x}_3(\theta),\pi) - \mathrm{BCP}(\hat{x}_2,\pi)$. The relative importance of perfect learning can be compared to the importance of maximizing under uncertainty and the importance of reevaluating the optimal best guess policy by dividing the overall benefit of acting upon climate change under perfect learning into the components:

$$\mathrm{BCP}(\hat{x}_3(\theta),\pi) \equiv \mathrm{VPI}(\pi) + \mathrm{BCP}(\hat{x}_1,\bar{\pi}) + \Delta\mathrm{BCP}(\hat{x}_1) + \mathrm{BOAU}\ , \tag{5}$$

[1] As done by Pizer [1997], who uses the relative measure $\mathrm{BOAU}/\mathrm{BCP}(\hat{x}_2,\pi)$ to assess the importance of optimizing under uncertainty.

and then comparing the normalized contributions of the single effects.

2.3 A simple example: Quadratic Benefits and Costs of Mitigation

We apply a simple analytical model of costs and benefits for climate change mitigation to illustrate the three components of the overall benefit to act upon climate change in the case of uncertainty. We will also use the simple setting to illustrate the connection between the different welfare effects from including uncertainty and the marginal expected benefits and costs of mitigation. Later this "marginal cost-benefit" picture will be used to understand the finding of small effects from uncertainty in the more complex integrated assessment model MIND.

Let x denote the level of abatement of greenhouse gas emissions relative to the no-control case in a simple one period framework. The abatement leads to benefits $B(x,\theta)$ due to reduced climate damage and comes with costs of mitigation $C(x)$. The dependence of B on the state of the world (SOW) θ denotes the uncertainty in climate damage and benefits of mitigation. The decision problem is to maximize the net benefits of mitigation:

$$\max_x U = E_\theta B(x,\theta) - C(x)\ . \tag{6}$$

Trivially, the optimal abatement level in the no-control case is $\hat{x}_0 = 0$. To be able to derive the optimal policies in the best guess and uncertainty case, $\hat{x}_1$,$\hat{x}_2$ and resulting welfare changes analytically, we choose quadratic benefits and costs of mitigation, i.e. $B(x,\theta) = a_1x^2 + f(\theta)x + c_1$ and $C(x) = a_2x^2 + b_2x + c_2$, similar to the model used by Karp & Zhang [2006].We assume the benefits and damage to vanish for zero mitigation: $c_1 = c_2 = 0$. The SOW is a normally distributed random variable: $\theta = \mathcal{N}(\mu,\sigma^2)$. For simplicity we assume $f(\theta)$ to take the form $f(\theta) = \theta^2$. By solving the first order condition for x the optimal solutions for the best guess and the general case of uncertainty read:

$$\hat{x}_1 = \frac{1}{2}\frac{\mu^2 - b_2}{a_2 - a_1}\ , \quad \hat{x}_2 = \frac{1}{2}\frac{\mu^2 + \sigma^2 - b_2}{a_2 - a_1}\ . \tag{7}$$

Fig. 2 shows the marginal benefits and costs for the best guess and the uncertain case. Presented this way, the benefits of climate policy can be illustrated simply as areas between the curves for marginal costs and benefits. For our functional setup, the three components from Eq. 4, normalized to the overall benefit of climate policy under uncertainty, are fully determined by the "strength" of the uncertainty effect on decisions, i.e. by $\Delta\hat{x} := \sigma^2/(\mu^2 - b_2)$. The relative contributions of the components are shown in Fig. 3. For amplitudes of the relative change in optimal policies due to uncertainty $\Delta\hat{x}$ that would be considered significant, e.g. 10%, the welfare gain from adjusting the optimal policy is much lower (1%), nearly the whole benefit of acting upon climate change is already realized by choosing, and reevaluating, the best-guess climate policy. Only very large adjustments of optimal policy $\Delta\hat{x} > 45\%$ lead to a significant ($> 10\%$) contribution to the overall benefit of climate policy.

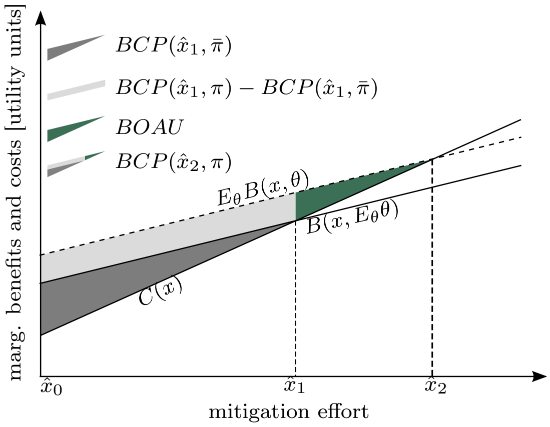

Figure 2: Marginal benefits and costs of mitigation for quadratic benefits and costs.

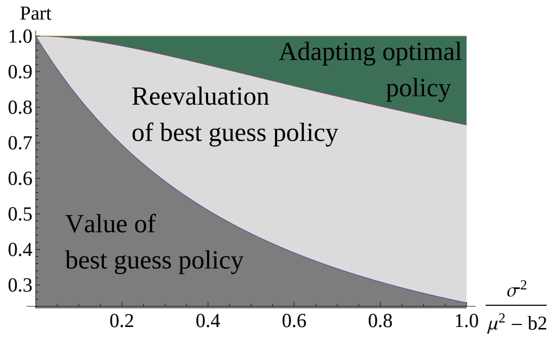

Figure 3: Three components of the overall benefit (value of best-guess policy in darker Grey, reevaluation of best-guess policy under uncertainty in lighter Grey, value of adapting optimal policy to uncertainty in darker green) of acting upon climate change under uncertainty depending on the amplitude of uncertainty $\sigma^2/(\mu^2 - b_2)$.

The optimal policy can also be derived graphically within the (marginal) cost benefit picture. The optimal best guess policy is determined by the intersection between the marginal costs and the marginal benefits for the expected value of the uncertain parameter. The optimal policy under uncertainty is determined by the intersection between the marginal costs and the expected marginal benefits. The size of the three components of the overall benefit of climate policy under uncertainty is determined by the slope and curvature of the marginal benefits and costs and by the slope and curvature of the difference between the expected marginal benefits in the uncertainty case and the marginal benefits in the best guess case, called the marginal risk premium MRP.

Assuming simple polynomial costs $C(x) = x^\alpha$ and benefits $B(x) = x^\beta$, the slope and curvature of the marginals is determined by α and β. For the costs of mitigation the common assumption is to take $\alpha > 2$, thus the marginal costs are convex increasing in the mitigation effort. For the benefits of mitigation the situation is not clear. Depending on β, three regimes would be possible: For $0 < \beta < 1$ the marginal benefits are decreasing, for $1 < \beta < 2$ they are concave increasing and for $\beta > 2$ they are convex increasing. Fig. 4 shows examples from the different regimes and the resulting benefits of best guess climate policy, here the marginal benefits are normalized by $B'(x, \beta) = \beta x^{(\beta-1)}/(\beta \cdot 0.5^{(\beta-1)})$ such that the marginal benefits coincide at $x = 0.5$ for all β. Clearly the benefit of climate policy decreases with increasing β.

Introducing uncertainty changes the marginal benefits. Commonly one assumes the benefits to be convex in the uncertain parameter θ, thus the marginal benefits increase due to uncertainty. However the slope and curvature of the marginal risk premium MRP $\equiv \sum_i \pi(\theta_i)B(x, \theta_i) - B(x, \sum_i \pi(\theta_i)\theta_i)$ are not clearly determined. The benefit from adapting

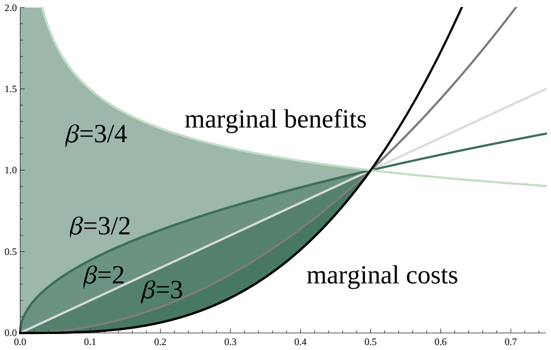

Figure 4: Marginal costs and benefits of mitigation for different slopes and curvatures of marginal benefits. The cost and benefit functions are normalized to reach the same value at $x = 0.5$..

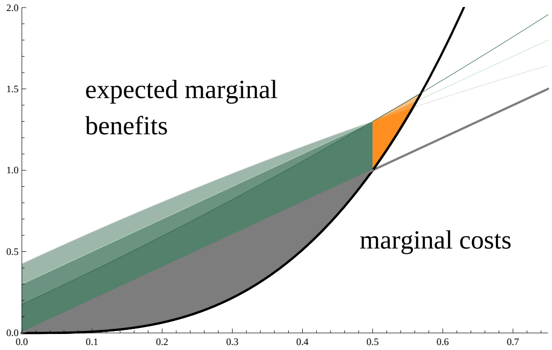

Figure 5: Expected marginal costs and benefits of mitigation for different slopes and curvatures of the marginal risk premium (MRP). The benefit of the best guess policy is marked as darker Grey area, the benefit of reevaluating the best guess policy is marked as green area and the benefit of adapting the policy to uncertainty is marked as orange area. The different shadings result from the different MRP realizations.

the policy to uncertainty (BOAU) is strongly influenced by the MRP. Fig. 5 shows a case of convex increasing marginal benefits with different possible MRP functions. The BOAU increases the more convex increasing the MRP. If the curvature of the MRP itself does not depend on the mitigation level x, then the increasing BOAU with increasing convexity of the MRP is contrasted with a decrease in the reevaluation of the benefit of climate policy.

3 Importance of Uncertainty in MIND

Why does accounting for uncertainty about the climate response and the climate damage change the results in standard applications of integrated assessment models of climate change only to a small degree? In this section we investigate this question by replacing the simple analytic model by the more complex integrated assessment model MIND, and applying the decomposition of the benefit of climate policy presented above. For this purpose, we introduce uncertainty and immediate perfect learning about climate sensitivity and the severity of climate change induced damage into the MIND model (Section 3.1). First, we reproduce the findings in the literature (e.g. Pizer, 1997; Manne, 1995) that explicitly including uncertainty has a small influence on the benefit of climate policy (Section 3.2). We then interpret the MIND model in the "cost-benefit" picture from Section 2.3 (Section 3.3) and resolve the functional dependencies between the decision variables and the resulting marginal benefits and costs of climate change mitigation to shed light on the prerequisites for a significant effect of uncertainty (Section 3.4).

3.1 The *M*odel of *In*vestment and Technological *D*evelopment (MIND)

We employ the *M*odel of *In*vestment and Technological *D*evelopment (MIND) (Edenhofer et al., 2005) in its stochastic version presented by Held et al. [2009]. Additionally we include learning as introduced to the model by Lorenz et al. [2011]. MIND is a model in the tradition of the Ramsey growth model and similar to the well-known DICE model (Nordhaus, 1993). The version we use differs from the classical Ramsey model in two major aspects: Firstly, the production sector depends explicitly on energy as production factor, that is provided by a crudely resolved energy sector. The energy sector contains (**i**) fossil fuel extraction, (**ii**) secondary energy production from fossil fuels, and (**iii**) renewable energy production. The macroeconomic constant-elasticity-of-substitution (CES) production function depends on labor, capital and energy as input factors. Secondly, technological change is modeled endogenously in two ways. The DM can invest into research & development activities to enhance labor and energy efficiency. Additionally, productivity of renewable and fossil energy producing capital increases with cumulative installed capacities (learning-by-doing). We assume welfare to be an inter-temporally separable isoelastic utility function of per capita consumption with a constant relative risk aversion $\eta = 1.5$ that is changed for the sensitivity study later on. It takes the form:

$$U(c(I,s)) = \sum_{t_0}^{t_e} L(t) \cdot \frac{1}{1-\eta} \left[\left(\frac{[c(I,s)](t)}{L(t)} \right)^{1-\eta} - 1 \right] e^{-\rho t} dt \ , \tag{8}$$

where $I = (I_{\mathrm{K}}, I_{\mathrm{R\&D}}, I_{\mathrm{Fossil}}, I_{\mathrm{Renewables}})$ is the vector of investment flows in the different sectors over time, s is the unknown state of the world, ρ is the pure rate of social time preference taken to be $0.01/yr$, and $L(t)$ is an exogenously given population scenario. Investments are related to the global consumption $[c(I,s)](t)$ via the budget constraint:

$$Y_{\mathrm{net}}(t,s) = [c(I,s)](t) + \sum_n I_n(t,s) \ , \ c(I,s) \geq 0 \ , \tag{9}$$

with the Gross World Product (GWP) Y_{net} net of climate related damage. Y_{net} is related to gross GWP over $Y_{net} = Y_{gross} \cdot DF$, where DF is a multiplicative damage factor defined by the damage function (see Roughgarden & Schneider, 1999):

$$DF(T) = \frac{1}{1 + a \cdot T^b} \ . \tag{10}$$

3.2 Importance of Uncertainty and Perfect Learning in MIND

The uncertainties about climate sensitivity and climate damage are described by probability distribution functions. The information about climate sensitivity CS is modeled by a log-normal distribution by Wigley & Raper [2001]: $\bar{\pi}(\mathrm{CS}) = \mathcal{LN}(0.973, 0.4748)$. The uncertainty about climate damage is taken to influence the amplitude a of the damage factor, but not the exponent b, which is taken as constant $b = 2$. The distribution over a is derived from a normal distribution over the parameter a' in $DF(T)^* = 1/[1+(T/a')^2]$, with

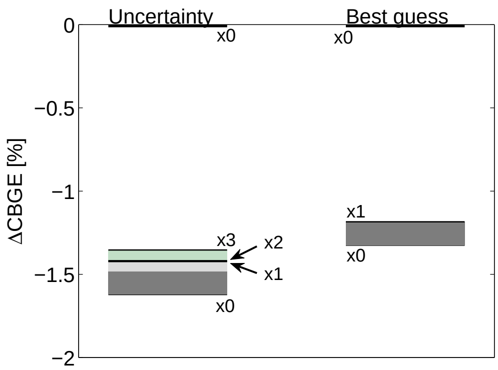

Figure 6: Welfare levels, measured as changes in CBGE relative to no-control, for the different scenarios with and without uncertainty. Shown are the results for $\eta = 1.5$.

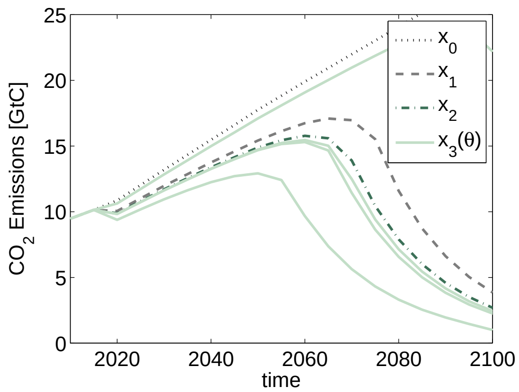

Figure 7: Optimal emission pathways for the different information settings. Shown are the solutions for risk aversion $\eta = 1.5$.

$a' = \mathcal{N}(18, 5)$. This choice of the mean is near to the best guess case by Nordhaus [2008] ($a = 0.0028$ vs. our $a = 0.0030$). The uncertainty range is inspired by the distribution by Gerst et al. [2010], who chose $a = \mathcal{N}(0.0028, 0.0013)$, but due to the inverse distribution, higher damage values are favored by our distribution. For the numerical implementation we draw samples of size n from the distributions according to a scheme related to descriptive sampling (see Saliby, 1997). The uncertainty space is divided into n hypercubes. Each hypercube i carries a chosen probability weight w_i and is represented by the expected value of the parameters on this hypercube. For simultaneous uncertainty about both climate sensitivity and damage, each dimension is sampled with four equi-probable points which are combined to only four learning paths for the perfect learning comparison according to the descriptive sampling scheme (instead of 16 learning paths with a full factorial design). We tested the influence of the low sampling size on the results by complementing the analysis with a sampling of 10×10 samples in CS and a for the case of uncertainty vs. best guess. The results did not change significantly.

Fig. 6 shows the welfare changes, relative to the no-control case, for the different scenarios with and without uncertainty within the MIND model. First, the benefit from acting upon climate change is small relative to the net costs due to the existence of climate change. In other words only a small part of the climate change induced welfare losses can be countered by mitigation policy. This observation stays the same in the uncertain setting, although the best guess climate policy leads to higher benefits against the no-control policy with uncertain damage. The welfare benefit from adapting the optimal policy is nearly invisible, whereas the welfare gain from perfect learning is significant. This finding compares to only small changes in optimal emission pathways from the best guess to the uncertain policy against strong changes in case of perfect learning (see Fig. 7)..

In a study close to this one, Pizer [1997] investigated the effect from explicitly including uncertainty into the DICE model by Nordhaus [1994]. He not only considered uncertainty about the socio-economic and the climate system but also about the normative parameters of risk aversion and pure rate of time preference. He found that the uncertainty about the normative parameters by far dominates the uncertainties about the socio-economic system. We perform a sensitivity study of the uncertainty components towards the parameter of constant relative risk aversion. The necessary scenarios for optimal climate policies under best guess, uncertainty and perfect learning have been evaluated for 8 different values of η. The resulting changes in the partition of the benefit of optimal climate policy are depicted in Fig. 8 for the case of uncertainty and in Fig. 10 for the case of perfect learning. The changes in optimal decisions between the best guess and the uncertainty case are depicted in Fig. 9.

From Fig. 8 a clear ordering of the different components of the overall benefit of climate policy emerges: The main part of the overall benefit of climate action can be realized by simply taking the optimal best guess policy. However, reevaluating this best guess policy in an uncertain information setting significantly increases the benefit. The changes in optimal decisions between the best guess and the uncertainty setting (see Fig. 9) are at least partly significant, e.g. a $> 5\%$ change in cumulative carbon emissions for the next two centuries. But the resulting welfare effect from this adjustments (BOAU) is insignificant for the whole range of η, thus the explicit incorporation of uncertainty into the optimization only plays a minor role. From Fig. 10 one can see that the contribution from perfect learning by far dominates the contribution from adapting to uncertainty. With increasing η, the value of perfect learning even dominates all other components. Summarizing the numbers indicates that within MIND uncertainty about the climate response to anthropogenic carbon emissions and about climate induced damage is not important for the assessment of optimal climate change mitigation whereas perfect learning clearly changes the picture.

An additional interesting feature of the sensitivity study with respect to relative risk aversion η is the rapidly decreasing overall benefit of climate policy for increasing η. For values of $\eta > 2$ the benefit from acting upon climate change gets lower than 0.01% of change in CBGE consumption. This can be explained with the dual role of the parameter η. Within the expected utility framework employed in most studies of optimal global mitigation assessment the parameter η represents both, the DM's constant relative risk aversion and her aversion to fluctuations of consumption over time. With increasing risk aversion, the DM reacts with stricter policies to minimize the uncertainty in climate impacts. But with increasing aversion to consumption fluctuations over time, within an overall growing economy the DM's incentive to shift consumption from the future towards the present becomes stronger. Within MIND, obviously the second effect is stronger, as the mitigation effort decreases with η, and therefore the benefit from acting upon climate change also decreases. A separation of both roles of η can be achieved within a normative satisfying setting by Treager (2009, and references therein). This is left for future studies.

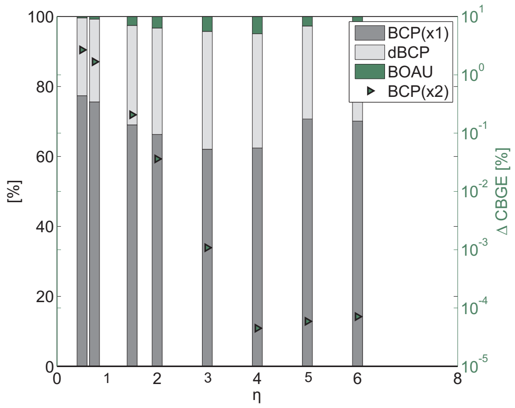

Figure 8: The three contributions (benefit of best guess policy in darker green, reevaluation of best guess policy in lighter green, benefit of adapting policy in yellow) normalized to the overall benefit of climate policy under uncertainty, calculated for different values of constant relative risk aversion η (the three small values for $\eta \leq 2$ are $\eta = .5; .75; 1.5$). Also shown is the overall benefit of climate policy itself on the right axes.

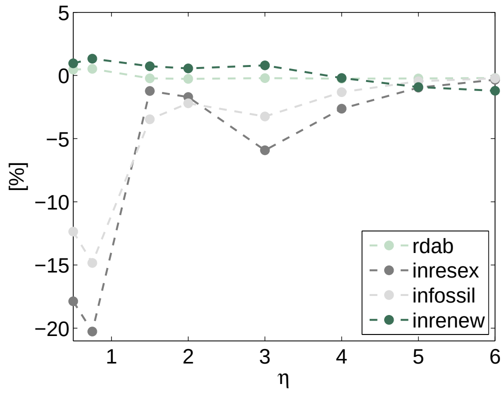

Figure 9: Relative changes in decision variables $\Delta\hat{x}$ (investments in R&D to increase energy efficiency (rdab), investments in extraction of fossil energy resources (inresex), investments in the capital stock of fossil energy carriers (infossil), and investments in the capital stock of renewable energy carriers (inrenew)), cumulated over the full time horizon (2010-2200), from the optimal best guess strategy $\hat{x}_1$ to the optimal strategy under explicit inclusion of uncertainty $\hat{x}_2$.

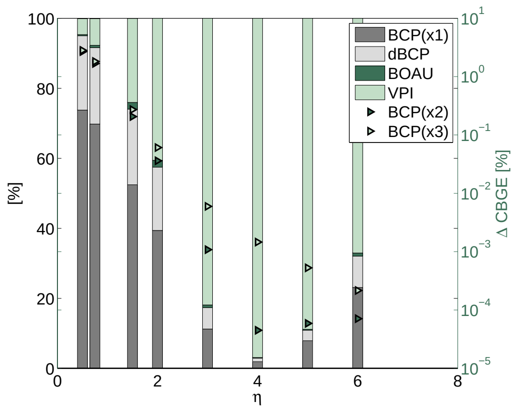

Figure 10: The four contributions (benefit of best guess policy in darker green, reevaluation of best guess policy in lighter green, benefit of adapting policy in yellow, value of perfect information in purple) normalized to the overall benefit of climate policy under perfect learning, calculated for different values of constant relative risk aversion η (the three small values for $\eta \leq 2$ are $\eta = .5; .75; 1.5$). Also shown are the benefits of climate policy for the case of uncertainty (lighter green triangles) and perfect learning (darker green triangles).

3.3 The Marginal Cost - Benefit picture of MIND

In the following we apply the marginal cost-benefit picture from Section 2.3 to the MIND model to understand the reasons for the small welfare effect from explicitly including uncertainty about the climate response and climate induced damage. Therefore we interpret the welfare benefit from choosing the optimal policies $\hat{x}_1, \hat{x}_2$ instead of the no-control policy $\hat{x}_0$ as composition of mitigation benefits $B(x, \theta)$ and mitigation costs $C(x)$, analogue to Sec. 2.3. For any climate policy x we define $B(x, \pi) \equiv \{[U(x, \pi) - U(x, \pi_0)] - [U(\hat{x}_0, \pi) - U(\hat{x}_0, \pi_0)]\}$ as the difference in the welfare impacts due to the existence of climate induced damage between the policies x and $\hat{x}_0$. Thereby π indicates a world with climate damage, and π_0 indicates a world without climate damage. We define mitigation costs $C(x)$ of policy x as: $C(x) \equiv [U(\hat{x}_0, \pi_0) - U(x, \pi_0)]$. This is the loss in welfare for choosing a suboptimal policy x instead of the optimal policy $\hat{x}_0$ in a world without climate damage (π_0). Simple calculus shows, that this choice actually delivers the desired composition for any policy x:

$$
\begin{aligned}
B(x, \pi) - C(x) &= \{[U(x, \pi) - U(x, \pi_0)] - [U(\hat{x}_0, \pi) - U(\hat{x}_0, \pi_0)]\} \\
&\quad - [U(\hat{x}_0, \pi_0) - U(x, \pi_0)] \\
&= U(x, \pi) - U(\hat{x}_0, \pi) \ .
\end{aligned}
$$

Using this composition, the problem of finding the optimal climate policy $\hat{x}$ for a given information setting (Eq. 1) can be rewritten as maximizing the difference between mitigation benefits and costs. This can be recast in an a-temporal cost-benefit picture by identifying the intersection of the marginal benefits ($dB(x, \pi)/dx$) and marginal costs ($dC(x)/dx$). The intersection point on the x-axis corresponds to the optimal policies $\hat{x}_1$ (using $B(x, \bar{\pi})$ or $\hat{x}_2$ (using $B(x, \pi)$). Numerically the benefits and costs for a given policy x are calculated by evaluating welfare differences as relative changes in CBGEs (see Sec. 5).

To be able to inspect the cost-benefit picture visually we additionally need to project the multi-dimensional decision variable x on a single-dimensional quantity. Thereby we loose the exact equivalence between the welfare picture and the cost-benefit picture. The goal is to choose a projection $x \to \tilde{x}$ that approximates the welfare effect of uncertainty with high accuracy and allows an interpretation of the small amplitude. We achieve the one-dimensional projection by introducing a constraint on cumulative emissions in a setting without climate damage. For a constraint above 3165GtC the no-control policy emerges. With decreasing levels of admissible cumulative emissions, the DM reacts by adjusting the investments into the different energy technologies (see Fig. 11). With increasing stringency of the constraint on cumulative emissions the investments into R&D in energy efficiency increase, as well as the investments into carbon free renewable energy. Contrary, the investments in carbon intensive fossil energy carriers and the corresponding resource extraction sector decrease.

Introducing a dense sampling in the cumulative emissions constraint and evaluating the resulting policies in settings with and without uncertainty allows to construct the cost-

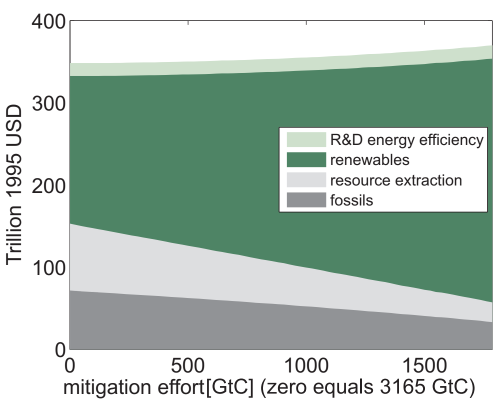

Figure 11: Aggregated investments (NPV, discounted with 5%) into the energy system depending on the required mitigation effort (in GtC cumulative reduction relative to the no-control in which 3165 GtC are emitted over the period 2010-2200).

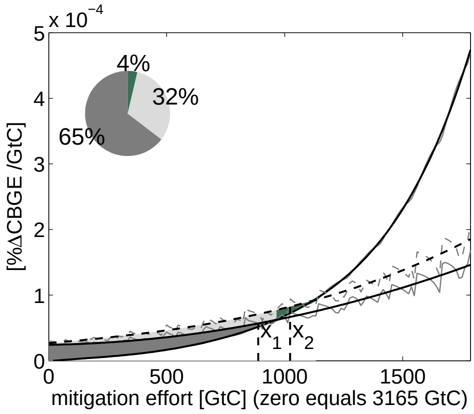

Figure 12: Marginal costs (lower black line) and (expected) marginal benefits (upper black and black dashed line) of mitigation within the MIND model. The Grey curves are the raw data from MIND, the black curves are polynomial fits. Mitigation effort is parametrized by a decreasing constraint on cumulative emissions. The optimal policies from the "correct" optimization are shown as vertical dashed lines. The constant relative risk aversion η is set to 2, this can be compared to the other values investigated in Sec. 4.1.

benefit picture for the MIND model (see Fig. 12). The marginal benefits for the best guess case are derived by fixing all uncertain parameters to their expected value while the expected marginal benefits are derived by applying the cost-benefit decomposition to the expected utility. The fluctuations in the Grey curves, that represent the raw data from the model can be explained by the limited temporal resolution of the model (5 yrs). When optimizing under a binding constraint with increasing stringency (such as the constraint on cumulative emissions), the timing of the mitigation effort to stay below the constraint can only be adjusted within this limited temporal resolution. This leads to small jumps in the overall welfare and thus also in marginal welfare and in marginal benefits and costs. The bold black lines are polynomial fits to the raw data.

Analogue to Section 2.3 the optimal mitigation effort for the best guess and the uncertainty setting can be obtained as intersections between the marginal costs and the (expected) marginal benefits of mitigation. The different contributions to the overall benefit of acting upon climate change can be visualized as areas between the benefit and cost curves. The pie chart in the upper left corner shows their relative contributions to the overall benefit of climate policy. A comparison between the optimal values of cumulative emissions derived from the marginal picture and those derived from the "correct" welfare optimization shows the "error" of the approximation. The optimal level of mitigation in cumulative emissions is represented within an 4% error, while the welfare effects of uncertainty are overestimated by up to 5%.

Nevertheless the cost-benefit picture allows to identify reasons for the negligible uncertainty effect. First, the overall value of acting upon climate change is constraint due to the **convex increasing** functional form of both (expected) marginal benefits and marginal costs. The combination of these functional forms lead to a very small area between both curves and thus, a small overall benefit of climate policy. This result is somewhat counter intuitive as one would assume climate damage to be convex increasing in temperature and temperature more or less linearly connected with mitigation effort and thus would expect **decreasing** marginal benefits. The reason for the counter-intuitive result from the MIND model is discussed further below. If marginal benefits were concave increasing in the mitigation effort or even decreasing, for fixed intersection points $\hat{x}_1$ and $\hat{x}_2$, the overall benefit of climate policy would increase. However, the same is not true for the value of adapting to uncertainty (BOAU) that is represented by the darker green "triangle" within Fig. 12. For fixed $\hat{x}_1$ and $\hat{x}_2$ more concave marginal benefits would increase the BOAU, but even more so the other components of the BCP, thus the relative importance of explicitly including uncertainty would even decrease.

Second, the marginal risk premium, that is the difference between expected marginal benefits under uncertainty and marginal benefits for the expected parameter values, only increases linearly in the mitigation effort and with a small slope. Together with the strongly convex increasing marginal costs, this leads to a relatively small difference between the two optimal policies $\hat{x}_1$ and $\hat{x}_2$. To increase the BOAU, one would need to either increase the slope of the MRP or even better, increase the convexity of the MRP in the mitigation effort. Both measures would lead to a larger difference between the marginal benefits in the best guess and the uncertainty case and thus to a larger difference between $\hat{x}_1$ and $\hat{x}_2$. However, increasing only the slope of the MRP would also increase the reevaluation effect of the best guess policy (lighter Grey area) and thereby limiting the relative importance of the BOAU.

3.4 Functional Dependencies within MIND

To find an explanation for both results, the slope and curvature of marginal benefits and the small MRP, we apply the marginal representation to the single steps in the climate cause-effect chain. The absolute and marginal functional form of the individual elements of the chain are shown in Fig. 13. This allows us to investigate in detail how the slope and curvature are determined in the integrated assessment model MIND.

The (cumulative) emissions lead to a rising concentration of greenhouse gases in the atmosphere and increasing radiative forcing. The maximum forcing reached for different levels of mitigation effort is shown in panel **a**. The maximum total forcing is concave increasing in cumulative emissions and concave decreasing in mitigation effort respectively. This can be explained by the saturation effect represented by the logarithmic relation from concentration to forcing. With increasing atmospheric concentration of carbon dioxide, the frequency band in which CO_2 absorbs the outgoing radiation saturates, thus a further

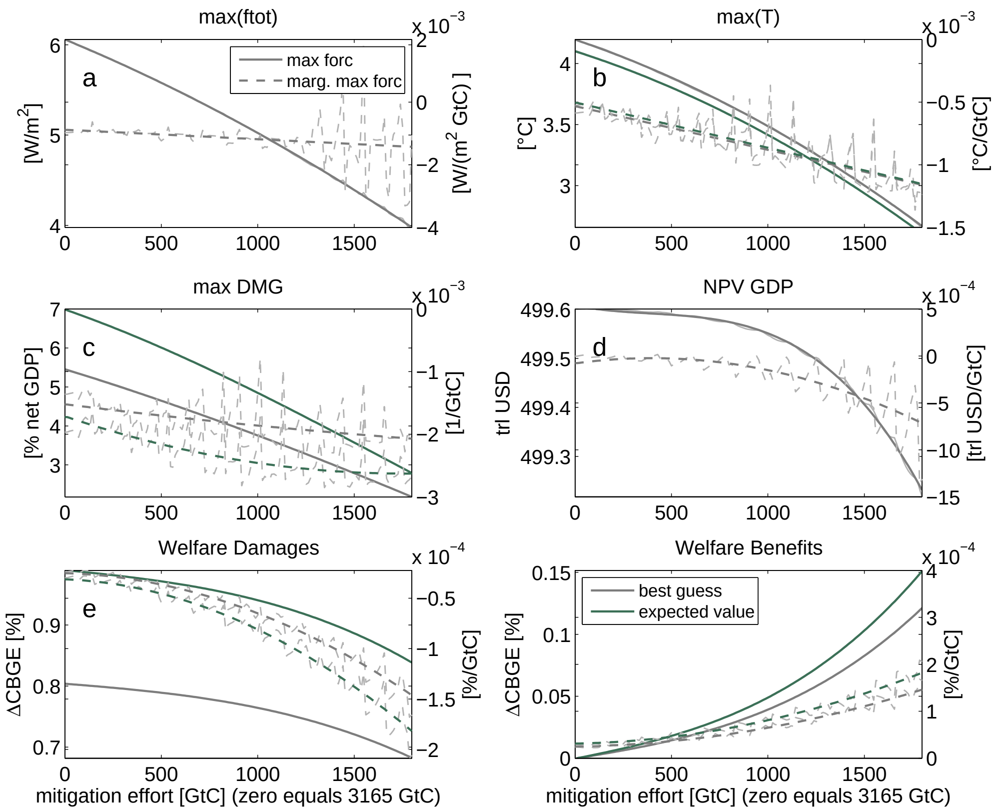

Figure 13: Functional dependencies of individual components in the cause-effect chain of climate change on mitigation effort measured in terms of cumulative emissions reductions from the BAU emissions of 3165 GtC in the period $2010-2200$: **(a)** maximum radiative forcing, **(b)** maximum temperature change, **(c)** damage in % of net GDP for the maximum temperature change, **(d)** net present value (NPV) of gross output including mitigation costs, but excluding climate damage, **(e)** welfare equivalent damage measured in %ΔCBGE, and **(f)** welfare benefits measured in %ΔCBGE. Shown are the functions (continuous lines) and the marginal functions (dashed lines). For those quantities that depend on the uncertain SOW the best guess value is shown in darker Grey and the expected value of the uncertain setting is shown in darker green. The original model data are shown in lighter Grey, where as the functional dependencies are polynomial fits of the data.

increase in concentration leads to less and less additional radiative forcing. The same concave behavior, increasing in cumulative emissions and decreasing in mitigation effort, occurs for the maximum temperature increase $\max(T)$, shown in panel **b**. The global mean temperature reacts to changes in radiative forcing. The overall temperature response is determined by the amplitude (climate sensitivity) and time scale (ocean diffusivity) of a simple impulse response model. The climate damage that is incurred by the maximum temperature change, measured in % of net GDP, is shown in panel **c**. Panel **d** shows the net present value of gross output excluding climate damage, aggregated over time by an endogenous discount rate $\rho_t \equiv \delta + \eta \cdot g_t$, where δ is the pure rate of time preference, η is the rate of constant relative risk aversion and g_t is the endogenously determined growth rate of consumption. The gross economic output is concave decreasing in the mitigation effort, leading to convex increasing mitigation costs in GDP terms, which are derived by subtracting the gross output curve from the gross output of the no-control case $x = 0$.

Multiplying the damage factor, $\mathrm{DF} = 1/(1 + D)$, where D are the net GDP damage from panel **c**, with the gross GDP in each time step gives the time series of net GDP that constrains the investment decisions and consumption level via a budget equation. Thus both the costs from mitigation (as seen in the gross GDP) and the climate damage lower the consumption level and thus the welfare. The welfare equivalent damage for the different mitigation scenarios, shown in panel **e**, is derived by evaluating the difference in CBGE between a case with the damage factor DF as above and a case without damage (but with mitigation costs), where $\mathrm{DF} = 1$. Formally the welfare damage is given as $\Delta\mathrm{CBGE}(V(x, \pi_0), V(x, \pi/\bar{\pi}))$ or in loose notation as $U(x, \pi_0) - U(x, \pi/\bar{\pi})$. Normalizing the welfare damage to the no-control case delivers welfare benefits from mitigation, shown in panel **f**. Formally this normalization is done by subtracting the welfare damage for the no-control case, leaving us with the definition of benefits from Section 3.3 $U(x_0, \pi_0) - U(x_0, \pi) - (U(x, \pi_0) - U(x, \pi)) = B$.

The two most interesting features in the cost-benefit picture of MIND are the positive slope and the convex curvature of the marginal benefits of mitigation. Concerning the slope of the marginal benefits in welfare, the explanation can already be found in panels **b** and **c**. The marginals of maximum (Panel c) and welfare equivalent damage (Panel e) are decreasing in the mitigation effort implying increasing marginal benefits. As can be seen from a comparison of panel c and e, this behavior is not a result of the welfare evaluation of climate damage (although it is strengthened by it), but already present in the marginal of the maximum damage. Since the maximum damage is convex increasing with rising temperature, their marginal is increasing with temperature as well. The fact that their marginal is decreasing when plotted against increasing mitigation effort instead of temperature, due to concave instead of convex decreasing maximum damage, points to the fact that the concavity found in the temperature response to mitigation dominates the convexity of damage in rising temperature. Thus, we find that the saturation of the emissions to temperature change relationship over-compensates the non-linearity in the climate damage function, leading to increasing instead of decreasing marginal benefits of mitigation, and thus limiting the overall benefit of the best guess climate policy (the dark Grey area in Fig. 2).

The convexity of marginal welfare benefits in mitigation effort however does not originate from the combination of maximum temperature with the damage function, but emerges from the welfare valuation of climate damage. This can be seen by comparing the convex decreasing marginal of maximum damage (panel **c**) to the concave decreasing marginal of welfare equivalent damage (panel **e**). Hence, the influence of the welfare function, i.e. of the normative parameters of constant relative risk aversion η and pure rate of time preference ρ determines the curvature of the marginal benefits. Comparing the marginal benefits for the best guess case and the case of uncertainty, it can be seen that the convexity increases when accounting for uncertainty, implying a convex increasing MRP. Thus, the additional marginal welfare benefit of reducing a unit of emissions under uncertainty grows with increasing mitigation effort. This works against a large contribution of re-evaluating the best guess climate policy under uncertainty (DBCP; light Grey area in Fig. 2), and favors a larger relative contribution of adjusting the mitigation policy under uncertainty (BOAU; orange area in Figure 2). However, due to the strongly increasing marginal mitigation costs, the welfare gains from adjusting the mitigation policy remain small.

Another important observation in the cause-effect chain is the small influence of uncertainty about climate sensitivity and the correlated time lag of the climate response (panel **b**). The expected maximum temperature increase for uncertain cs is slightly shifted towards lower levels, which in itself is counter intuitive, as uncertainty about cs has an asymmetric upper tail. The explanation for this could be the correlation between climate sensitivity and the time scale of climate response due to the observations of 20^{th}century global mean temperature. Higher values of climate sensitivity are connected to longer time scales of temperature response, thus the temperature increase will only show later. In combination with the limited time horizon due to discounting this limits the influence of high cs values. In addition, the marginal maximum temperature increase shows nearly no change from best guess to the expected case. Thus the uncertainty about cs can not lead to a change in marginal benefits due to uncertainty.

4 Changes in the Model Structure

Which assumptions about the climate cause-effect chain would lead to a significant welfare gain from adapting the optimal policy to uncertainty? In this section we investigate several changes in the model structure and their influence on the cost-benefit picture and the BOAU.

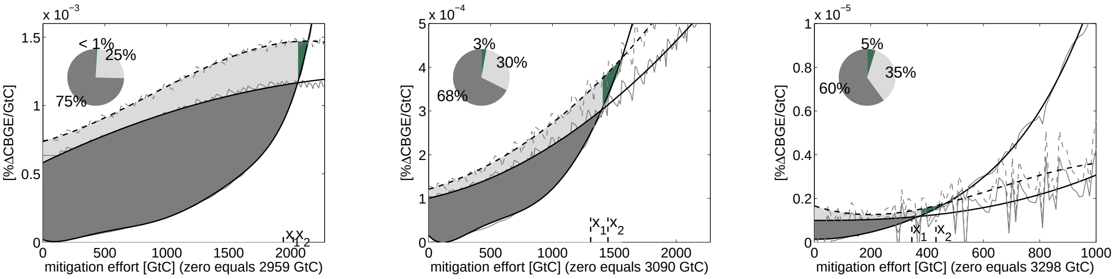

Figure 14: Sensitivity of the marginal cost-benefit picture of MIND with respect to changes in the parameter of constant relative risk aversion η. Shown are the pictures for $\eta = 0.75$(left), $\eta = 1.5$(middle), $\eta = 3$(right). The legend is equivalent to Fig. 12.

4.1 Constant relative risk aversion η

We have shown in Section 3.4 that the curvature of the welfare function, represented by the parameter of constant relative risk aversion η, strongly influences the curvature of the marginal benefits of mitigation. We have also shown (Section 2.3) that the curvature of the marginal benefits strongly influence the overall benefit of climate policy. We use these dependencies and investigate the relative importance of adjusting the optimal policy to uncertainty depending upon the parameter of constant relative risk aversion. The changing cost-benefit pictures of MIND are shown in Fig. 14 for values of η between 0.75 and 3.

The effects of the curvature of the welfare function are manifold: The first, and most important one, is a scaling effect. As already shown in Fig. 8, the overall net benefit of climate policy strongly decreases with increasing η. This effect can be explained by the dual role of η. It does not only represent risk aversion, but also the DM's aversion towards inter temporal fluctuations in consumption. If this aversion is high, the DM prefers a smooth, constant consumption stream over a fluctuating, increasing one. In a growing economy with consumption growth in the future the decision maker prefers to delay mitigation as it would require to divert consumption into early investments in carbon free energy technologies. This effect can already be seen in the baseline cumulative emissions (the number given as label below the figures). Hence the net benefits from mitigation, i.e. reduced climate damage minus mitigation costs, are lower for a high η as mitigation reduces early consumption. The second effect concerns the curvature of the marginal benefits. This effect is directly evident: as discussed in Sec. 2.3, the exponent η of the welfare function obviously directly determines the curvature of the marginal welfare depending on consumption.

In combination both effects result in increasing absolute overall benefits of climate policy and increasing absolute welfare gains from adjusting climate policy to uncertainty with decreasing η . However, the relative contribution of the BOAU to the benefit of climate policy decreases with η.

4.2 Exponential Damage

Unlike Weitzman [2010] who focused on the potential fat tails of the distribution on climate sensitivity and climate damage we are searching for a setting in which a strong impact of uncertainty on optimal mitigation efforts also occurs for thin tailed distributions. As stated before a stronger increase, and convexity, in the marginal risk premium in welfare terms for rising mitigation effort would lead to a stronger BOAU.

First , we replace the standard quadratic formulation of the damage function by an exponential formulation:

$$\mathrm{DF}_e = \frac{1}{1 + k \cdot \exp(\frac{T}{l}) - k} \ . \tag{11}$$

We choose the parameters k and l such that the exponential damage in net GDP, $k \cdot \exp(T/l) - k$ equal the standard formulation at $T = 3°$ for the best guess case. We assume a normally distributed l with $l = \mathcal{N}(2.2571, 0.61)$. Together with $k = 0.01$, this choice leads to a best guess marginal damage function which is nearly identical to the standard best guess marginal damage function used in the previous section. However, the expected marginal damage function is far more convex in temperature than in the quadratic case. Thus the marginal risk premium in net GDP damage increases more strongly. The resulting difference in the marginal of the maximum damage is shown in rows **a** and **b** of Fig. 15 together with the identical maximum temperature functions and the resulting cost benefit pictures. The change towards exponential damage shows several interesting effects: First, as we have chosen identical best guess marginal damage, the optimal policy in the best guess case also does not change, but the higher marginal risk premium leads to an increased $\hat{x}_2$. The welfare benefit from adapting the policy to uncertainty increases significantly and now contributes 13% to the overall benefit of climate policy (instead of 4% in the quadratic case). The increase in the relative contribution of the BOAU is dampened by the fact that the effect of reevaluating the best guess policy under uncertainty is also strongly increasing. This is due to the fact, that the expected marginal benefits do not only increase more strongly than before but are also shifted upwards over the whole domain. Second, the change towards exponential damage leads to at least partly increasing expected marginal damage in net GDP. However this shift in the slope of the marginal damage is not strong enough to be reflected in the expected marginal benefits, it "gets lost" through the convolution with the welfare function. Finally, compared to the case of quadratic damage, the overall value of climate policy more than doubles for the assumption of exponential damage, with more than two thirds of the benefit due to taking uncertainty into account.

4.3 Linear Carbon Climate Response

Finally, we replace the climate module by a linear relationship between cumulative carbon emissions and increase in global mean temperature that has been found by Matthews et al. [2009] within an ensemble of state of the art climate models. By introducing the

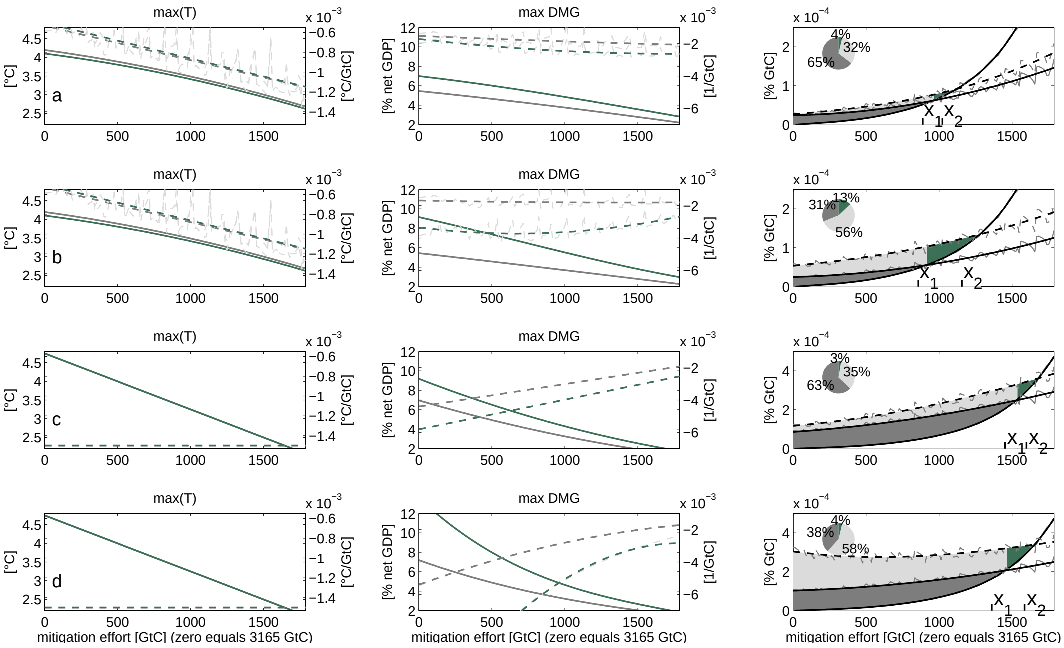

Figure 15: Maximum temperature change, maximum climate damage in net GDP, and marginal cost benefit picture for four different structural model settings: **(a)** standard climate module and quadratic damage function, **(b)** standard climate module and exponential damage function, **(c)** linear climate carbon response and quadratic damage function, and **(d)** linear climate carbon response and exponential damage function. The functional relations are shown in darker Grey for the best guess case and in darker green for the expected value of the uncertainty case. The dashed lines represent the marginal functions. The fluctuating lines in lighter Grey are the original model data. The smooth lines are polynomial fits to the data. The legend for the cost benefit pictures is analogous to Fig. 12.

so called carbon climate response (CCR) parameter, the relationship between global mean temperature change (relative to pre industrial) ΔT and cumulative carbon emissions reads:

$$\Delta T(t) = \mathrm{CCR} \cdot \sum_{t_0}^{t'} e(t') \ , \tag{12}$$

where e are globally aggregated carbon emissions. Within the model ensemble, Matthews et al. [2009] found a carbon climate response of $\mathrm{CCR} = 1.5[1.0 - 2.1]$°C/TtC. The values in square brackets mark the 5 and 95 quantile. We choose a log-normal distribution for the CCR with $\mathrm{CCR} = \mathcal{LN}(\log(1.461), 0.23)$, which gives the best fit to the quantiles and 1.5°C as expected value. The resulting maximum temperature, maximum net GDP damage and the cost benefit pictures are shown in rows c and d of Fig. 15. In row c the linear climate carbon response is combined with quadratic damage and in row d with exponential damage from Sec. 11. Considering the maximum temperature response, the difference between the best guess case and the expected case under uncertainty nearly vanishes. This is clear, as the uncertain parameter CCR now enters linearly into the function, thus the expectation operator only acts on the parameter itself. Thus the uncertainty in the climate system is now irrelevant for the mitigation problem. But even more interesting is the change in the best guess maximum temperature function itself. It declines more strongly in the mitigation effort than before. This leads to a stronger difference between the best guess and expected damage. The changed curvature of the climate response leads to convex decreasing net GDP damage. However, the increasing slope of the marginals is changed back to decreasing marginal welfare equivalent damage further downstream by the welfare function. Hence the marginals in the welfare benefit are still increasing, but less convex than before. Compared to the standard climate module, the overall benefit of climate policy increases strongly (by a factor of X), but the individual contributions of the three components remain largely unchanged. In particular, the welfare contribution from adapting the optimal policy to uncertainty is still negligible.

Combining the linear climate carbon response with exponential damage amplifies the distinct features of the two cases. The convexity in expected net GDP damage gets strong enough to “survive” the convolution with the welfare function leading to initially decreasing expected marginal benefits in welfare terms. This further increases the benefit from reevaluating the best guess policy. The BOAU stays small.

Summarizing the results from this Section, the functional formulation of the building blocks of the climate cause-effect chain (temperature response, climate induced damage and the aggregated welfare function) and especially their marginals determine the strength of the effect of including uncertainty. Thereby the non-linearities in the temperature response and the damage function partly compensate each other. Everything else equal, the importance of adjusting policies to uncertainty becomes the more important the more convex the damage function, the lower the risk aversion parameter and the lower the concavity of the maximum temperature in cumulative emissions.

5 Conclusion

We applied a decomposition of the overall benefit of acting upon climate change into its single components to measure the importance of uncertainty and perfect learning within the integrated assessment model MIND.
Uncertainty influences both, the optimal mitigation policy and the expected utility of different policies. Including uncertainty explicitly is important, if it leads to a significant change in the optimal policy (before any potential future learning) that in turn leads to a significant change in the benefit gained from acting upon climate change. Uncertainty might also be considered relevant, but would not have to be included explicitly into the optimization framework, if it significantly changed the assessment of the benefit of climate policy compared to the best guess case (reevaluation effect), even though the optimal climate policy (before learning) did not change significantly.
Within the MIND model the reevaluation effect is dominating the welfare gain from adjusting the policy under uncertainty, while perfect learning is dominating both of these effects. Overall, the welfare effect of accounting for uncertainty is rather small, which further corroborates the findings in the literature.
To understand the origin of these findings, we projected the complex MIND model to an a-temporal marginal cost-benefit picture and resolved the functional relationship between the single steps of the climate cause-effect chain. Thereby we located the origin of the negligible welfare gain from adapting the optimal policy to uncertainty. This benefit of anticipating uncertainty (BOAU) is only of significant size if uncertainty leads to non-linear shifts in the marginal benefits of mitigation, that would lead to a convex strongly increasing marginal risk premium with increasing mitigation effort.
In the standard model setting with a quadratic damage function and a zero-dimensional climate-carbon response box model, this behavior was constrained by the saturation of the emissions to temperature change relationship compensating for the non-linearity in the climate damage function and by the consumption smoothing property of the welfare function. Thus for seeing a significant influence from including uncertainty one has to consider alternative model settings that induce a strongly convex increasing marginal risk premium (MRP).
Two such changes in the model setup, an exponential climate damage function and a linear climate carbon response have been implemented. We showed that those changes to the model structure indeed can lead to a strongly convex increasing MRP and a significant uncertainty effect.
The other feature that constrains the importance of including uncertainty is the strongly increasing marginal mitigation cost curve in the model. Thus a change in the convexity of marginal mitigation costs, especially reducing the strong increase for higher levels of mitigation would also lead to more significant uncertainty effects. This is of special interest as it emphasizes the combined importance of the modeling of mitigation options and the impact and damage formulation for the overall importance of uncertainty for the integrated

assessment of climate change.

These results come with the usual caveats. The employed integrated assessment model MIND, although more complex than quasi-analytical cost-benefit models and the commonly used DICE model, still includes a strongly simplified representation of the cause-effect chain of climate change.The representation of uncertainty and learning had to be constrained to a few sample points and to the limiting case of perfect learning, and we only investigated the effect from a single information setup.
Thus this study should not be seen as an attempt to find a conclusive answer to the question whether accounting for uncertainty and learning is important for the assessment of climate policy. Rather, we present an approach to decompose and trace the uncertainty effect in complex integrated assessment models, which we believe will prove useful to improve our understanding about the effect of structural model assumptions on the significance of the uncertainty effect.

A Comparing Welfare across different scenarios

As (expected) utility is only defined up to an affine transformation, we use differences in the certainty and balanced growth equivalents (CBGE), as presented by Anthoff & Tol [2009], to compare different scenarios. The certainty equivalent of an uncertain consumption outcome is an amount of consumption the DM would demand instead of a distribution of outcomes to get the same expected utility. The same principle works for the balanced growth equivalent: here the consumption path, that possibly varies over time, is replaced by a path consisting of an initial consumption level that growth over time with a constant growth rate α and gives the same utility. If one is only interested in relative changes in the CBGE between different scenarios, the measure is independent of the growth rate α. Thus the relative change in CBGE, denoted by ΔCBGE, can be interpreted as fraction of consumption the DM would be willing to pay, now and forever, to switch from a scenario with lower CBGE to the other scenario. Formally the ΔCBGE for isoelastic utility reads:

$$\Delta\text{CBGE}\left[EU_1, EU_2\right] \equiv \begin{cases} \left(\frac{EU_1}{EU_2}\right)^{1-\eta} - 1 & \text{for } \eta \neq 1 \\ \exp\left(\frac{EU_1 - EU_2}{\sum_{t_0}^{T} P_t (1+\rho)^t}\right) - 1 & \text{for } \eta = 1 \end{cases} . \tag{13}$$

References

D. Anthoff and R.S.J. Tol. The impact of climate change on the balanced growth equivalent: An application of FUND. *Environmental & Resource Economics*, 43(3):351–367, 2009.

Erin Baker. Increasing risk and increasing informativeness: Equivalence theorems. *Operations Research*, 54(1):26–36, 2006.

Ottmar Edenhofer, Nico Bauer, and Elmar Kriegler. The impact of technological change on climate protection and welfare: Insights from the model MIND. *Ecological Economics*, 54:277–292, 2005.

Michael D. Gerst, Richard B. Howarth, and Mark E. Borsuk. Accounting for the risk of extreme outcomes in an integrated assessment of climate change. *Energy Policy*, 38(8): 4540–4548, August 2010. ISSN 0301-4215.

Geoffrey Heal and Bengt Kriström. Uncertainty and climate change. *Environmental & Resource Economics*, 22(1):3–39, 2002.

H. Held, E. Kriegler, K. Lessmann, and O. Edenhofer. Efficient climate policies under technology and climate uncertainty. *Energy Economics*, 31:S50–61, 2009.

Alan Ingham, Jie Ma, and Alistair Ulph. Climate change, mitigation and adaptation with uncertainty and learning. *Energy Policy*, 35(11):5354–5369, November 2007.

Larry Karp and Jiangfeng Zhang. Regulation with anticipated learning about environmental damages. *Journal of Environmental Economics and Management*, 51(3):259–279, May 2006.

K. Keller, B. M. Bolker, and D. F. Bradford. Uncertain climate thresholds and optimal economic growth. *Journal of Environmental Economics and Management*, 48(1):723–741, July 2004.

D. L. Kelly and C. D. Kolstad. Bayesian learning, growth, and pollution. *Journal of Economic Dynamics & Control*, 23(4):491–518, February 1999.

A. Lange and N. Treich. Uncertainty, learning and ambiguity in economic models on climate policy: some classical results and new directions. *Climatic Change*, 89(1):7–21, July 2008.

A. Lorenz, M.G.W. Schmidt, E. Kriegler, and H. Held. Anticipating climate threshold damages. *Environmental Modeling and Assessment*, 2011.

Alexander Lorenz, Hermann Held, Eva Bauer, and Thomas Schneider Deimling. Constraining ocean diffusivity from the 8.2 ka event. *Climate Dynamics*, 34(5):719–734, 2009. ISSN 0930-7575.

A. S. Manne. *EMF-14 Subgroup - Analysis for Decisions under Uncertainty*, chapter Hedging Strategies for Global Carbon Dioxide Abatement: A Summary of Poll Results. Stanford University, Stanford, CA., 1995.

H. Damon Matthews, Nathan P. Gillett, Peter A. Stott, and Kirsten Zickfeld. The proportionality of global warming to cumulative carbon emissions. *Nature*, 459(7248):829–832, June 2009. ISSN 0028-0836.

Tevfik F. Nas. *Cost-benefit analysis: theory and application.* SAGE, 1996. ISBN 9780803971332.

W. D. Nordhaus. Rolling the DICE - an optimal transition path for controlling greenhouse gases. *Resource and Energy Economics*, 15(1):27–50, March 1993.

W. D. Nordhaus. *Managing the global commons.* MIT Press Cambridge, MA, 1994.

W. D. Nordhaus and D. Popp. What is the value of scientific knowledge? an application to global warming using the PRICE model. *Energy Journal*, 18(1):1–45, 1997.

W.D. Nordhaus. *A Question of Balance: Weighing the Options on Global Warming Policies*. Yale University Press, illustrated edition edition, June 2008. ISBN 0300137486.

William A. Pizer. Optimal choice of policy instrument and stringency under uncertainty: The case of climate change. *Working Paper: Resources for the Future*, 97:17, 1997.

Tim Roughgarden and Stephen H. Schneider. Climate change policy: quantifying uncertainties for damages and optimal carbon taxes. *Energy Policy*, 27(7):415–429, 1999.

E. Saliby. *Descriptive Sampling: An Improvement Over Latin Hypercube Sampling*. December 1997.

Richard S. J. Tol. Safe policies in an uncertain climate: an application of FUND. *Global Environmental Change*, 9(3):221–232, October 1999. ISSN 0959-3780.

Christian P. Traeger. Recent developments in the intertemporal modeling of uncertainty. *Annual Review of Resource Economics*, 1:261–286, 2009.

M. Webster. The curious role of "learning" in climate policy: Should we wait for more data? *Energy Journal*, 23(2):97–119, 2002.

Martin L. Weitzman. GHG targets as insurance against catastrophic climate damages. *National Bureau of Economic Research Working Paper Series*, No. 16136, June 2010.

T. M. L. Wigley and S. C. B. Raper. Interpretation of high projections for Global-Mean warming. *Science*, 293(5529):451–454, July 2001.

Franz Wirl. Energy prices and carbon taxes under uncertainty about global warming. *Environmental & Resource Economics*, 36(3):313–340, 2007.

Chapter 4

Anticipating Climate Threshold Damages *

Alexander Lorenz
Matthias G.W. Schmidt
Elmar Kriegler
Hermann Held

*This chapter has been published as: Lorenz, A., Schmidt, M.G.W., Kriegler, E., Held, H 2011. Anticipating Climate Threshold Damages. *Environmental Modeling and Assessment*. DOI:10.1007/s10666-011-9282-2.

Environ Model Assess
DOI 10.1007/s10666-011-9282-2

Anticipating Climate Threshold Damages

Alexander Lorenz · Matthias G. W. Schmidt · Elmar Kriegler · Hermann Held

Received: 31 July 2010 / Accepted: 1 April 2011

Abstract Several integrated assessment studies have concluded that future learning about the uncertainties involved in climate change has a considerable effect on welfare but only a small effect on optimal short-term emissions. In other words, learning is important but anticipation of learning is not. We confirm this result in the integrated assessment model "model of investment and technological development" for learning about climate sensitivity and climate damages. If learning about an irreversible threshold is included, though, we show that anticipation can become crucial both in terms of necessary adjustments of pre-learning emissions and resulting welfare gains. We specify conditions on the time of learning and the threshold characteristic, for which this is the case. They can be summarized as a narrow "anticipation window."

Keywords Epistemic uncertainty · Learning · Anticipation · Value of information · Value of anticipation · Threshold damages

A. Lorenz (✉) · M. G. W. Schmidt · E. Kriegler · H. Held
Potsdam Institute for Climate Impact Research,
14412 Potsdam, Germany
e-mail: lorenz@pik-potsdam.de

A. Lorenz
Environmental Change Institute,
School of Geography and the Environment,
University of Oxford, South Parks Road,
Oxford, OX1 3QY, UK

H. Held
University of Hamburg & Klima Campus Hamburg,
Bundesstr. 55, 20146 Hamburg, Germany

1 Introduction

Climate change poses a formidable global problem. Climate impacts may occur over a wide range of sectors, countries and time. Moreover, the regions most vulnerable to the impacts differ from those responsible for the largest parts of emissions. Although climate science has gained a profound understanding of the elementary processes underlying climate change, big uncertainties about its magnitude and implications remain. These scientific uncertainties will be reduced in the future, and it will be possible to adjust climate policy accordingly.[1] Investments in mitigation of greenhouse gas emissions are at least partially sunk or irreversible, respectively. The combination of uncertainty, learning about uncertainty and irreversibility makes it interesting to study the effect of anticipation of future learning on optimal near-term climate policy. Important questions in this context are: Should society wait for better information about the climate system and climate damages before committing to mitigation measures or should it mitigate preemptively? Does anticipation of future learning yield significant welfare increases?

A theoretical literature has established theorems about the sign of the anticipation effect, i.e., the effect of anticipation of future learning on optimal short-term decisions. In very simple two-period models, a Bayesian decision maker (DM) is characterized by a goal

[1] We will assume that learning eventually reveals the true values of parameters. For interesting examples, where new information might narrow the uncertainty around a false value, see Oppenheimer et al. [30] and Kriegler [19].

Published online: 28 July 2011

function $U(x_1, x_2, s)$, where s is the state of the world, and the decision variables x_t, $t \in \{1, 2\}$ denote direct consumption of a generic good, emissions of a pollutant, or investment decisions. The DM first chooses x_1, then gets some message y containing information about the uncertain s, and finally chooses x_2. The question under consideration is: In which direction does the optimal first period decision x_1 change depending on the informativeness of y? The most general answer to this question has been given by Epstein [9], who showed that it depends on the properties of the 2^{nd}-period value function $j(x_1, \pi) \equiv \max_{x_2} \sum_s \pi_s U(x_1, x_2, s)$, where π_s is the probability of s. More information (in the sense of Blackwell [4]) unambiguously, i.e., independent of the specific form of the information structure (in the sense of Marschak and Miyasawa [23]), leads to a lower optimal level of x_1 if and only if $\partial j/\partial x_1$ is convex in π_s. One strand of the literature applies Epstein's condition in simple analytically solvable models (see, e.g., Kolstad [18]; Gollier et al. [12]). In more complex models, though, Epstein's condition is of limited value for two reasons: Firstly, it is hard to apply because it is difficult to determine the convexity of the marginal value function in π_s. Therefore, Baker [2] and and Salanie and Treich [34] have recently provided necessary and sufficient conditions for the primitives of the model, i.e., $U(x_1, x_2, s)$ instead of $j(x_1, \pi)$, for being able to decide upon the anticipation effect unambiguously: U has to be separable in s, which means that U has to be linear in some function $g(s)$. Unfortunately, most integrated assessment models do not belong to this class; thus, further investigation and imposition of more structure on the model and information setup will be necessary to come to a satisfactory answer.

The integrated assessment literature has therefore focused on explicitly calculating optimal short-term decisions under learning in more complex numerical models. A few studies have investigated the effect of learning under a climate target O'Neill et al. [28], Bosetti et al. [5], Johansson et al. [14], and parts of Webster et al. [40]. The latter, e.g., find that anticipation of learning about climate sensitivity leads to significantly stronger short-term emission reductions under a strict targets. However, Schmidt et al. [36] argue that this effect results from a disputable interpretation of climate targets as targets that have to be met with certainty. Investigations of the anticipation effect in cost-benefit analysis include Peck and Teisberg [31], Yohe and Wallace [42], Kelly and Kolstad [17], Leach [22], and parts of Webster et al. [40]. See Lange and Treich [21] for a review. These studies have shown that learning has generally a small effect on optimal short-term decisions, whereas the question of the welfare gain due to anticipatory changes in pre-learning decisions was not addressed.

Here, we confirm this result in the integrated assessment model "model of investment and technological development" (MIND) for two key uncertainties of the climate problem, namely climate sensitivity and climate damages. We find considerable values of information but insignificant gains from anticipating learning. We then focus on the question whether the anticipation of learning about a tipping point-like irreversible threshold damage is important. This was already done with a different model and somewhat different focus by Keller et al. [16]. We advance on this analysis by investigating the welfare gain from anticipation, by using a different integrated assessment model, and by performing additional sensitivity analysis. We find that the anticipation of learning about threshold damages can lead to significant welfare gains if learning takes place in a specific "anticipation window," which depends on the threshold under consideration and the flexibility of the decision maker to reduce emissions. Thereby, the largest welfare gain due to anticipation does in general not result from the largest anticipatory change of near-term emissions.

The paper is structured as follows: Section 2 shortly introduces the problem formulation, the terminology of the expected value of anticipation, and the integrated assessment model MIND. The results from learning about climate sensitivity and smooth climate damages are presented in Section 3.1. Section 3.2 focuses on learning about irreversible, tipping point-like threshold damages and includes the main results. Section 4 concludes with potential implications for climate policy. A table of the nomenclature we will use is shown on the right.

Nomenclature

BAU	Business as usual
BOCP	Benefit of climate policy
(C)BGE	(Certainty) and balanced growth equivalents
CEVOI	Conditional expected value of information
DM	Decision maker
EVOA	Expected value of anticipation
EVOI	Expected value of information
(E)VPI	(Expected) value of perfect information
MIND	Model of investment and technological development
RnD	Research and development

2 Model and Methodology

2.1 Problem Formulation

We introduce learning, i.e., the change of information available to the DM over time, in its simplest possible form. The overall time horizon is split into a first period before and a second period after a one-time updating of information at learning point t_{lp}. A strategy consists of first period decisions (investments) $x_1 = I(t)$, $t_0 < t \leq t_{\text{lp}}$ and second period decisions $x_2(y) = I(y)(t)$, $t_{\text{lp}} < t \leq T$, which are conditional on messages y. The problem of the decision maker is now to maximize the outcomes of the chosen strategy in terms of an intertemporally separable, aggregated expected utility.

The learning between the two periods can formally be described by the concept of an information structure. The terminology follows Marschak and Miyasawa [23] as presented in Jones and Ostroy [15]. We denote states of the world and messages, or observations, by $s \in S$ and $y \in Y$, respectively. Let π and q be prior probability vectors on S and Y, respectively. Let π^y be a posterior probability vectors on S after receipt of message y and Π the matrix whose columns are the π^y. If the learning is consistent, which is ensured by applying Bayes' rule to update the prior probabilities, it holds

$$\pi_s = \sum_y q_y \pi_s^y \,. \tag{1}$$

Therefore, we will shortly denote the information structure by the tuple (Π, q).

Using this notation, the recursive optimization problem reads:

$$\max_{x_1} \sum_s \pi_s u_{1,s}(x_1) + \sum_y q_y \max_{x_2} \sum_s \pi_s^y u_{2,s}(x_1, x_2, y) \\ =: EU(\Pi, q) \,, \tag{2}$$

where $u_{1,s}(\cdot)$ and $u_{2,s}(\cdot)$ are the vectors of utility in period 1 and 2, respectively, with elements equal to utility for a specific state of the world s. We solve the problem numerically in the equivalent, but more convenient, sequential form

$$\max_{x_1^y, x_2^y} \sum_y q_y \sum_s \pi_s^y \left(u_{1,s}\left(x_1^y\right) + u_{2,s}\left(x_1^y, x_2^y\right)\right) ,$$
$$s.t. \qquad x_1^j = x_1^k,\ \forall j \neq k. \tag{3}$$

Here, the constraint ensures that only second period decisions can be tailored to the messages.

2.2 Terminology

We will distinguish between a "no learning" case, represented by an information structure with posterior distributions equal to the prior distribution, and a "learning" case in which the probability distribution narrows between the two time periods due to the received messages y. We will further distinguish two learning cases: Either the DM anticipates future learning before it happens or not. Learning has both an effect on optimal pre- and post-learning decisions, i.e., x_1 and x_2, both of which have a positive effect on welfare. The pre-learning adjustments are due to the anticipation of future learning, whereas post-learning adjustments can be made even if the learning is not anticipated. This is shown schematically in Fig. 1.

We now introduce several concepts that separate the effect of anticipated and non-anticipated learning. The benefits from adjusting post-learning decisions to new information for given first period decisions can be measured by the *conditional expected value of information* (CEVOI). Formally

$$\text{CEVOI}(x_1, \Pi, q) \equiv V(x_1; \Pi, q) - V(x_1; \pi, 1) \,, \tag{4}$$

where $V(x_1; \Pi, q)$ is the so-called value function, namely the optimal second period utility for given first period decisions and information structure (Π, q), $V(x_1; \Pi, q) = \sum_y q_y \max_{x_2} \sum_s \pi_s^y u_{2,s}(x_1, x_2)$. $V(x_1; \pi, 1)$ is the value function without learning.

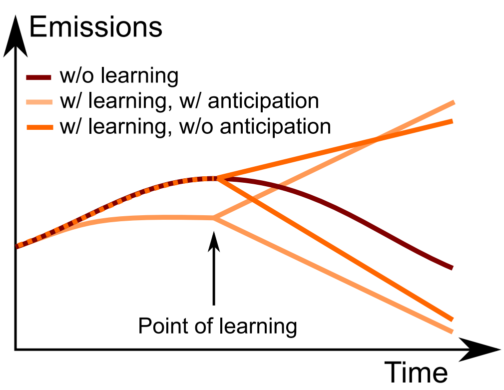

Fig. 1 Schematic plot of optimal emissions over time under different information scenarios and for two learning paths

The anticipatory adjustment of first period decisions to future learning can be measured by the *expected value of anticipation* (EVOA):

$$\mathrm{EVOA}(\Pi, q) \equiv \sum_s \pi_s u_{1,s}(x_1^*) + V(x_1^*, \Pi, q) - \left(\sum_s \pi_s u_{1,s}(x_1') + V(x_1', \Pi, q)\right), \quad (5)$$

where x_1^* and x_1' denote the optimal first period decisions with and without learning, respectively.

The overall wealth benefits from future learning can be measured by the *expected value of information* (EVOI). It is defined as the difference between expected utility with and without learning

$$\begin{aligned}\mathrm{EVOI}(\Pi, q) &\equiv EU(\Pi, q) - EU(\pi, 1)\\ &= \sum_s \pi_s u_{1,s}(x_1^*) + V(x_1^*, \Pi, q)\\ &\quad - \left(\sum_s \pi_s u_{1,s}(x_1') + V(x_1', \pi, 1))\right)\\ &= \mathrm{CEVOI}(x_1', \Pi, q) + \mathrm{EVOA}(\Pi, q)\,, \quad (6)\end{aligned}$$

The EVOI could be used to decide about the implementation of a certain observation campaign or scientific program providing certain information. The EVOI would therefore be compared to the implementation costs. The relevance of anticipatory changes in short-term policy as part of the overall benefits from information can be measured by the ratio EVOA/EVOI.

CEVOI, EVOA, and EVOI are defined as differences in expected utility, which are not invariant with respect to linear affine transformations of utility. To obtain this invariance, we use the concept of balanced growth equivalents (BGE) due to Mirrlees and Stern [24]. The BGE is defined as an initial level of consumption γ such that the balanced growth path $c(t) = \gamma \cdot \exp(\alpha t)$ yields the same expected utility as the original consumption path. Since we consider uncertainty and learning, we use the certainty equivalent BGE (CBGE) defined by Anthoff and Tol [1], where the certainty equivalent is with respect to the uncertain state of the world and the learning paths. For constant relative risk aversion η, the relative change in CBGE is:

$$\Delta\mathrm{CBGE} = \frac{\gamma(\mathrm{EU}) - \gamma(\mathrm{EU}')}{\gamma(\mathrm{EU}')} = \begin{cases} \left[\frac{\mathrm{EU}}{\mathrm{EU}'}\right]^{\frac{1}{1-\eta}} - 1 & \eta \neq 1\\ \exp\left(\frac{\mathrm{EU}-\mathrm{EU}'}{\sum_{t=0}^{T} L_t (1+\rho)^{-t}}\right) - 1 & \eta = 1\,, \end{cases} \quad (7)$$

where EU and EU$'$ are expected utility with and without learning, respectively, and the other denominations are population L_t and a discount factor due to impatience $(1+\rho)^{-t}$. It can easily be shown that relative changes in CBGE are independent of the growth rate α (Anthoff and Tol [1]). Intuitively, a 1% reduction inCBGE, for instance, can be interpreted as a permanent loss of consumption of 1%.

2.3 The Integrated Assessment Model MIND

We use the MIND (Edenhofer et al. [7]).[2] We use the version from Held et al. [13] and add anticipated learning about uncertainty (see Section 2.1), but we leave out carbon capturing and sequestration (CCS) for tractability. Edenhofer et al. [7] and Held et al. [13] perform cost-effectiveness analysis for a given climate target. We have shown elsewhere (Schmidt et al. [36]) that cost-effectiveness leads to conceptual problems if learning about uncertainty is taken into account. Therefore, we perform cost-benefit analysis.

MIND is a model in the tradition of the Ramsey growth model and similar to the well-known DICE model (Nordhaus [26]). The version we use differs from the classical Ramsey model in three major respects: Firstly, the production sector depends explicitly on energy as production factor that is provided by a crudely resolved energy sector. The energy sector contains (a) fossil fuel extraction, (b) secondary energy production from fossil fuels, and (c) renewable energy production. The macroeconomic constant-elasticity-of-substitution (CES) production function depends on labor, capital, and energy as input factors. Secondly, technological change is modeled endogenously in two ways. The social planner can invest into research and development activities to enhance labor and energy efficiency. Additionally, productivity of renewable and fossil energy producing capital increases with cumulative installed capacities (learning by doing). Thirdly, a simple energy balance model is used to translate global CO_2 and SO_2 emissions[3] to radiative forcing and changes in global mean temperature (Petschel-Held et al. [32]; Kriegler et al. [20]). SO_2 emissions are coupled to CO_2 emissions with an exogenously declining ratio of sulfur per unit CO_2 representing desulfurization. Radiative

[2]Modified model versions feature an endogenous carbon capturing and sequestration (CCS) module (Bauer [3]), a more elaborate carbon cycle and atmospheric chemistry module (Edenhofer et al. [8]), and parametric uncertainty (Held et al. [13]).

[3]The emissions are induced by (a) endogenous consumption of fossil fuels and (b) exogenous CO_2 emissions from land-use change (SRES A1T).

forcing from other greenhouse gases and aerosols is included as exogenous scenario (see Held et al. [13]).

We assume welfare to be an inter-temporally separable isoelastic utility function of per capita consumption with a constant relative risk aversion of $\eta = 2$. It takes the form:

$$U(c(I,s)) = \sum_{t_0}^{t_e} L(t) \cdot \frac{1}{1-\eta} \times \left[\left(\frac{[c(I,s)](t)}{L(t)} \right)^{1-\eta} - 1 \right] e^{-\rho t} \mathrm{d}t \,, \tag{8}$$

where $I = (I_{\mathrm{K}}, I_{\mathrm{R\&D}}, I_{\mathrm{Fossil}}, I_{\mathrm{Renewables}})$ is the vector of investment flows in the different sectors over time, s is the unknown state of the world, ρ is the pure rate of social time preference taken to be 0.01/year, and $L(t)$ is an exogenously given population scenario. Investments are related to the global consumption $[c(I,s)](t)$ via the budget constraint:

$$Y_{\mathrm{net}}(t,s) = [c(I,s)](t) + \sum_n I_n(t,s) \,, \; c(I,s) \geq 0 \,, \tag{9}$$

with the gross world product (GWP) Y_{net} net of climate related damages. Y_{net} is related to gross GWP over $Y_{\mathrm{net}} = Y_{\mathrm{gross}} \cdot \mathrm{DF}$, where DF is a multiplicative damage factor defined by the damage function (see Roughgarden and Schneider [33]):

$$\mathrm{DF}(T) = \frac{1}{1 + a \cdot T^b} \,. \tag{10}$$

For some of the results, we will limit the flexibility of the decision maker in MIND in one of two ways. First, we introduce a maximum flexibility in emissions changes $\Delta E_{\max}$/year as the maximum possible relative emissions change in one year both upward and downward. This inflexibility is assumed to originate from processes that are not included in the model MIND, such as political or societal constraints. Second, we limit the use of different mitigation options in MIND and particularly renewable energy and investments in energy efficiency. This increases the costs for emission reductions and thus lowers the flexibility in emission reductions. The influence of these two different kinds of inflexibility on the value of learning and anticipation is investigated.

2.4 Implementation of Learning About Climate Sensitivity and Damage Amplitude

We now consider a perfect learning case, i.e., messages y reveal the true state of the world. We focus on uncertainty about climate sensitivity CS, defined as equilibrium temperature change for a doubling of atmospheric CO_2 concentration from pre-industrial level, and on uncertainty about the climate damage parameters a and b in Eq. 10. We consider learning about climate sensitivity and damages separately as well as the combined effect of learning about both uncertainties simultaneously. The time of arrival of new information is varied between early ($t_{\mathrm{lp}} = 2030$), intermediate ($t_{\mathrm{lp}} = 2050$), and late learning ($t_{\mathrm{lp}} = 2070$). The uncertainties are described by probability distribution functions that are given explicitly in Appendixes 1 and 2. For the numerical implementation, we draw samples of size n from the distributions according to a scheme related to descriptive sampling (see Saliby [35]). The uncertainty space is divided into n hypercubes. Each hypercube i carries a chosen probability weight w_i and is represented by the expected value of the parameters on this hypercube. Thereby we do not choose an equiprobable spacing but choose a few central sampling points that carry the main part of probability and complement them by some points at the outer margin of probability. This technique of explicitly sampling the 1st and 99th percentile allows us to account for the low-frequency high-impact events in the tails of the distributions. For the implementation of learning about single uncertainties, we choose a sampling size $n = 5$. For the simultaneous learning about both uncertainties, each dimension is sampled with four equiprobable points which are combined to only four learning paths according to the descriptive sampling scheme (instead of 16 learning paths with a fully factorial design).

2.5 Implementation of Learning About Threshold Damages

Keller et al. [16] have found significant changes in emissions due to anticipation of learning if a highly nonlinear irreversible threshold is included in the analysis. More specifically, they considered a possible shutdown of the North Atlantic thermohaline (THC) circulation (Broecker [6]). We add to this study by focusing on the welfare benefits from anticipation, i.e., the EVOA, by using MIND as a model featuring endogenous technical change, and by performing a sensitivity analysis with respect to learning time, flexibility in emissions reductions, threshold temperature, and damages.

Hence, in addition to the damage function in Eq. 10 by Nordhaus [25], we consider explicit tipping point-like threshold damages. Similar to Keller et al. [16], who considered a threshold in atmospheric CO_2 concentration depending on climate sensitivity, we assume that the temperature T_0, at which the threshold occurs, is known, but the resulting damages $\mathrm{DF}_{\mathrm{thresh}}$ are uncertain. The damages are added to Nordhaus's damage factor DF leading to output net of damages, $Y_{\text{net}} = Y_{\text{gross}} \cdot \text{DF}_{\text{thresh}}$. We assume that the threshold is irreversible, i.e., if it has been crossed, the threshold damages continue to be incurred even if temperature returns to values below the threshold. This can be expressed formally as

$$\text{DF}_{\text{thresh}}(t, I_{n,t}, s) = \frac{1}{1 + a \cdot T^b + D_{\text{thresh}}(s) \cdot \xi(t, I_{n,t}, s)}, \tag{11}$$

where $D_{\text{thresh}}(s)$ is the amount of damages in the uncertain state of the wold s and $\xi(t, x_t, s)$ indicates whether the threshold was crossed before time t in the state s for given decisions up to time t, $I_{n,t}$. ξ is defined as

$$\xi(t, I_{n,t}, s) = 1 - \prod_{t'=t_0}^{t} [1 - \Theta(T(t', I_{n,t}, s) - T_0)], \tag{12}$$

and equals one if the threshold was crossed in the past and zero if not. Here, Θ is Heavyside's step function.

For simplicity, we only consider perfect learning about the threshold-damage amplitude D_{thresh}, which can only take two values, $D_{\text{thresh}} = [D_x, 0]$. Damage $D_{\text{thresh}} = D_x$ occurs with probability p and damage $D_{\text{thresh}} = 0$ with $1 - p$, such that the expected damage $\text{ED}_{\text{thresh}} = 1.5\%$ of net GDP is in accordance with empirical estimates for the expected impact of a THC shutdown by Tol [37]. We calculate the EVOI and the EVOA for different threshold temperatures T_0, threshold damages D_x (where p is adjusted such that expected net damages are unchanged, whereas the expected gross damage factor $\text{DF}_{\text{thresh}}$ changes), and learning points t_{lp}.

3 Results

3.1 Learning About Climate Sensitivity and Damage Amplitude

The welfare benefits from learning about climate sensitivity and standard climate damages, measured by the EVOI, are listed in Table 1. Learning about damages leads to an increase in CBGE of about 0.1% for early learning. When asking for the importance of including learning into the analysis of optimal climate policy, this value might best be compared to the overall benefit of climate policy (BOCP). The BOCP is the welfare difference between BAU and optimal policy measured in CBGE. It amounts to 0.12% CBGE in case of uncertain climate sensitivity and 0.14% CBGE in case of uncertain damages. Including learning about damages increases the BOCP by (21.8–64.5)% for late and early learning, but learning about climate sensitivity by only 1.75-4.95%. Hence, learning about damages can substantially increase the benefits from climate policy. Learning about climate sensitivity is less valuable by roughly an order of magnitude.

Simultaneous learning about both uncertainties strongly increases the EVOI, e.g., up to 0.45% for early learning. That relates to an increase of the BOCP by up to 347%. Hence, learning multiplies the benefits from climate policy if both parameters are uncertain. States of the world characterized by extreme values in both parameters imply very high damages. These can be mitigated after learning without having to spend the associated costs in all states of the world.

Also shown in Table 1 is the proportion of the EVOI that is obtained by anticipatory changes in pre-learning decision, i.e., the ratio EVOA/EVOI (see Section 2.2). We see that it is generally small ($< 2\%$). The welfare benefits from anticipating future learning about damages or climate sensitivity is negligible.

The result that learning implies only very small anticipatory changes in optimal pre-learning decisions in cost-benefit analysis was already found in other integrated assessment models (see, e.g., Ulph and Ulph [38]; Nordhaus and Popp [27]; Webster [39]; O'Neill and Melnikov [29]; Webster [40]). Why could we have expected an effect in the model MIND? As shortly discussed in Section 1, optimal first period decisions change, if the derivative of the second period, ex post value function $V_2(x_1, \pi_s^y) = \max_{x_2} \sum_s \pi_s^y u_{2,s}(x_1, x_2)$ with respect to the first-period decision x_1 is non-linear in the vector of posterior probabilities π_s^y (Epstein [9]), $\alpha\partial/\partial x_1 V_2(x_1, \pi_s^i) + (1-\alpha)\partial/\partial x_1 V_2(x_1, \pi_s^j) \neq \partial/\partial x_1 V_2$

Table 1 The EVOI measured in %CBGE of the no-learning case and the EVOA/EVOI ratio for different scenarios: perfect learning about CS and damages separately as well as jointly and for early, intermediate, and late learning

	CS		Damages		CS and Damages	
t_{lp}	EVOI (%)	EVOA/EVOI (%)	EVOI (%)	EVOA/EVOI (%)	EVOI (%)	EVOA/EVOI (%)
2030	0.006	0.004	0.09	0.29	0.45	0.022
2050	0.004	0.15	0.06	0.53	0.33	0.112
2070	0.002	0.20	0.03	1.77	0.22	0.287

CS climate sensitivity

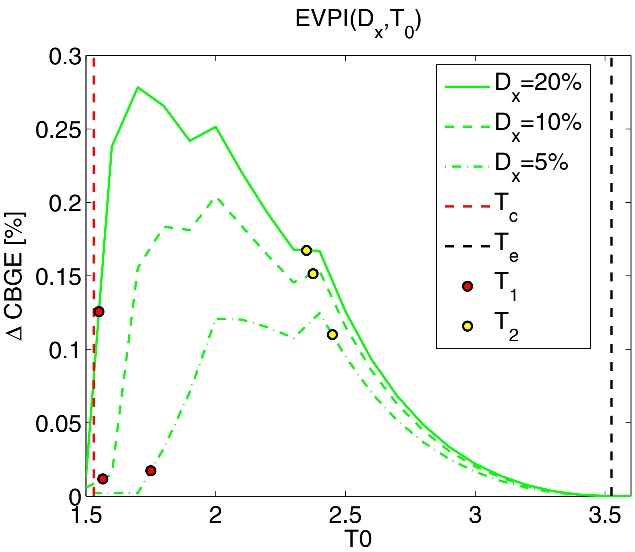

Fig. 2 The EVPI for different values of D_x and T_0. The EVPI is measured in %CBGE of the no-learning case. T_c denotes the temperature the decision maker is already committed to cross. For $T_0 > T_1(D_x)$, avoiding the threshold is optimal for perfect information that $D_{\text{thresh}} = D_x$. For $T_0 > T_2(D_x)$, avoiding the threshold is optimal even in the no-learning case. T_e is never reached for any information setup

$(x_1, \alpha\pi_s^i + (1-\alpha)\pi_s^j)$. Obviously, a necessary precondition for this is that the optimal second period utility V_2 actually depends on the first period decision x_1 and the derivative is non-zero. MIND includes several such cross-period interactions that are not present in other integrated assessment models. More specifically, it features multiple capital stocks, a knowledge stock, and learning by doing in technologies. However, the numerical results above clearly show that the effect of anticipation is negligible in this setting.

3.2 Learning About Threshold Damages

3.2.1 The Expected Value of Perfect Information

We start by considering two extreme cases: Either the decision maker has perfect information, i.e., learning occurs before any decision is to be taken, or she does not learn at all. Figure 2 shows the associated expected value of perfect information (EVPI)[4] for different values of the threshold specific damages D_x occurring with mean-adjusted probability $p(D_x)$ (see Section 2.3) and different threshold temperatures T_0. Also shown is the critical temperature $T_2(D_x)$ that divides the parameter space into two regimes: (A) For all threshold temperatures $T_0 < T_2$, it is optimal without learning to cross the threshold, and (B) for all $T_0 \geqq T_2$, it is optimal without learning to stay below the threshold. A further separation occurs within regime A: For threshold temperatures $T_0 < T_1(D_x)$, it is optimal to cross the threshold even in case of perfect information as the mitigation costs more than outweigh the threshold damages.

The EVPI is zero for high values of $T_0 > T_e$ because information about a threshold that is not crossed for the optimal policy without threshold damages is useless. However, the same is not true for very low values of $T_0 < T_c$, when the decision maker is committed to cross the threshold. The information about the received threshold damages is still valuable as it is used to adjust the savings rate. At a certain T_0, the EVPI reaches a maximum. For lower T_0, the emissions reductions that are necessary to avoid the threshold are too costly. For higher T_0, the avoided threshold damages decrease because higher T_0 are only reached later in time, and thus, the corresponding damages are discounted.

Since the EVOA is bounded from above by the EVOI and the EVOI is bounded from above by the EVPI, the potential benefits from anticipation are larger in regime A than in regime B. We also note from Fig. 2 that the EVPI is increasing in D_x, although expected damages are held constant by reducing the probability of the threshold when increasing D_x. This is due to the risk aversion of the decision maker, which makes her prefer a low D_x with a higher probability to a higher D_x with a low probability.

3.2.2 The Value of Anticipation

Now we investigate the dependence of the EVOI and the EVOA on the time of learning t_{lp}. Figure 3 shows the EVOA and EVOI for learning points between the year $t_{\text{lp}} = 2010$ and $t_{\text{lp}} = 2080$ in steps of 5 years. It also shows the cumulative anticipatory changes in emissions (ΔE) before learning relative to the no-learning case. The EVOI decreases from the EVPI obtained in 2010 to zero for $t_{\text{lp}} = 2200$. The latter is essentially the no-learning case. The EVOA has to be zero for $t_{\text{lp}} = 2010$ because there are simply no pre-learning decisions to be made. It is also zero for $t_{\text{lp}} = 2200$ because the discounted utility after this time is too small to justify anticipation.

Within regime A, where the threshold is crossed in the case of no learning, three different regimes of

[4] The EVPI is defined as the difference in welfare between the case of perfect information and the no-learning case. It is measured in %CBGE of the no-learning case.

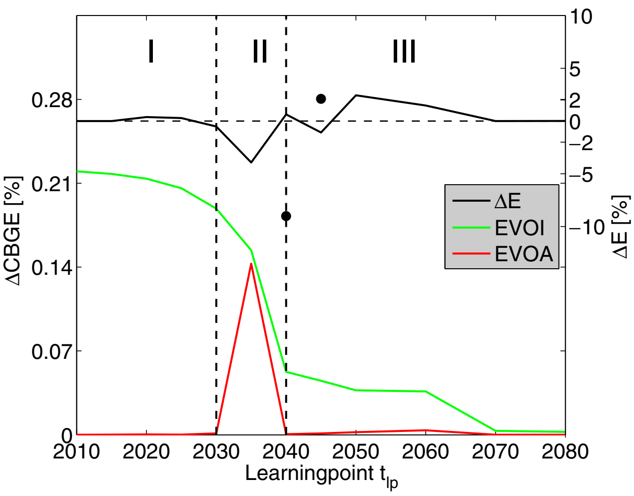

Fig. 3 Expected value of information (EVOI), expected value of anticipation (EVOA), and relative changes in cumulative pre-learning emissions in anticipation of learning (ΔE) are shown depending on the time of learning t_{lp}. The *dashed lines* mark three distinct regimes of anticipation (I–III). The *two black points* in 2040 and 2045 mark local optima that are only slightly worse compared to the shown "optimal" path

anticipative behavior can be identified. They are indicated in Fig. 3: (I) For early learning, it is possible to avoid the threshold easily by adjusting the post-learning decisions. Doing so in case $D_{\mathrm{thresh}} = D_x$ is learned leads to a substantial EVOI without the need for downward anticipation. Not having to anticipate downward benefits the case where $D_{\mathrm{thresh}} = 0$ is learned. There is even some upward anticipation to come closer toward the solution that would be optimal for perfect information about $D_{\mathrm{thresh}} = 0$.

(II) For increasing t_{lp}, there is less time between learning and crossing the threshold (without adjustments). Since mitigation costs are convex, this increases the costs of avoiding the threshold in the "bad case" ($D_{\mathrm{thresh}} = D_x$) by post-learning adjustments alone. Therefore, in regime II, the DM lowers pre-learning emissions compared to the no-learning case. The benefits of doing so experienced in the "bad" case outweigh its costs in the "good" case. For further increasing t_{lp}, avoiding the threshold with post-learn adjustments alone becomes physically infeasible. The motive for anticipation is then to keep the option open to avoid the threshold in the bad case in the first place. The associated costs increase with t_{lp}.

(III) At the border between regimes II and III, these costs reach a point, at which the decision maker is indifferent between keeping the option open and not keeping the option open, i.e., crossing the threshold also in the "bad" case. This leads to local optima with identical expected utility. Two of them are indicated by black dots in the upper panel of Fig. 3. Although the threshold is crossed for both learning paths in regime III, learning about the damages has a value, as witnessed by the significant EVOI for $t_{\mathrm{lp}} > 2040$ in Fig. 3. The reason is that learning still enables the DM to adjust her savings rate to damages and thus to perform consumption smoothing. More specifically, savings are decreased after crossing the threshold if the threshold is "bad". Finally, regime III shows a positive anticipation effect in emissions. However, the benefits from this anticipation are negligible.

In conclusion, downward anticipation for being able to avoid the threshold at all, or at low costs, in the bad case is the dominant effect. Anticipation of learning about threshold damages leads to a significant welfare gain only if the learning occurs within a specific time window $t_1 < t_{\mathrm{lp}} < t_2$. This "anticipation window" is narrow, and it spans at most one decade. Due to the 5-year time steps in MIND, it is not possible to determine its exact extent. The fact that the anticipation window is narrow is explained by the relatively high flexibility of the model in increasing or decreasing emissions. We will discuss this further in Section 3.2.4.

3.2.3 Availability of Renewable Energy

We investigate the origin of the anticipation window by focusing on the anticipation effect in the decision variables. These are investments in renewable energy, fossil energy, RnD aimed at improving labor or energy efficiency, and investments in the aggregate macro-economic capital stock. The cumulative anticipatory changes of the decision variables relative to the case without learning are shown in the left panel of Fig. 4. The right panel shows the cumulative post-learning adjustments up to 2200 separately for $D_{\mathrm{thresh}} = 0$ and $D_{\mathrm{thresh}} = D_x$. The resulting EVOI and EVOA are shown in Fig. 5.

The main option for reducing emissions used by the model is substituting fossil energy by renewable energy. Renewables are used to avoid the threshold after learning in regime I and for anticipatory emission reductions in regime II. The latter can be seen by comparing the "all options" case in Fig. 5 with the case, where the usage of renewables is restricted to be lower than in business-as-usual ("no renewables"), which is not zero but very little. The EVOA vanishes in the latter case. Apparently, anticipatory emissions reductions via reductions in energy demand or increased energy efficiency would be too costly. Hence, the existence of the anticipation window rests on the availability of a sufficiently cheap and flexible, carbon free, substitute

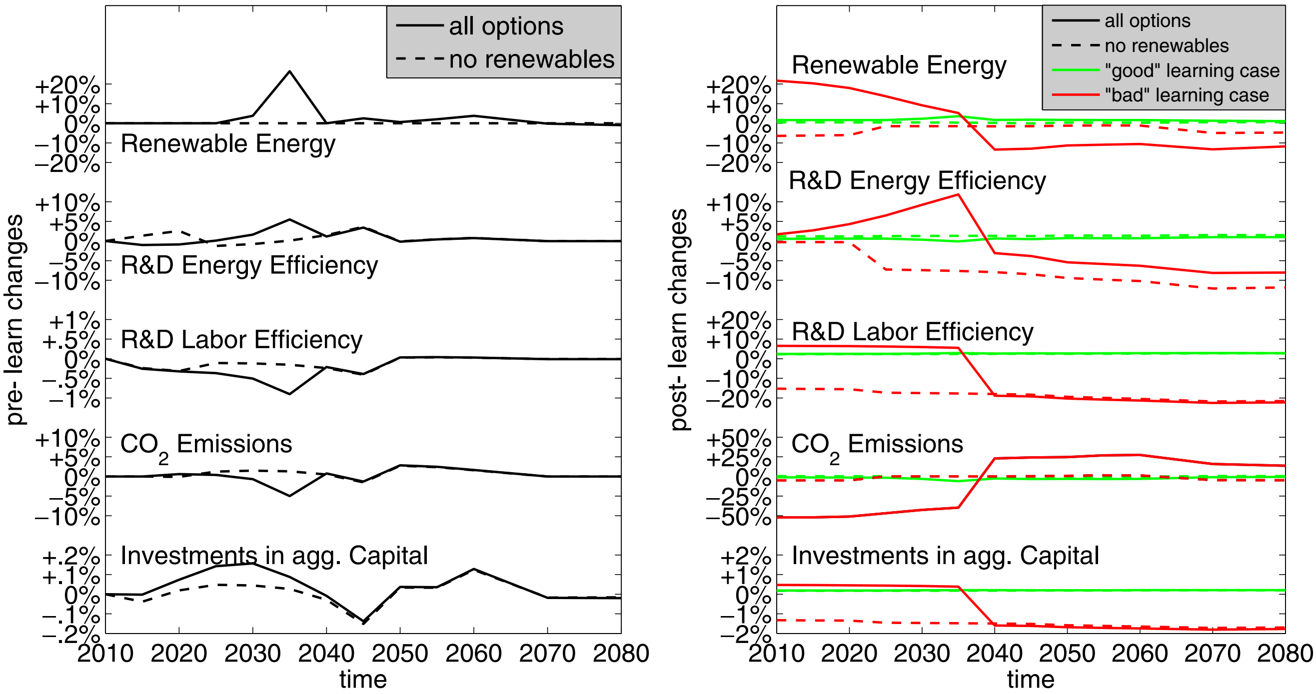

Fig. 4 The anticipation effect (*left*) and post-learning decisions (*right*) both in cumulative decision variables and with and without the availability of renewable energy

for fossil energy. However, too much flexibility would again diminish the EVOA because adjustments could be made entirely after learning. This suggest that an intermediate flexibility generates anticipation.

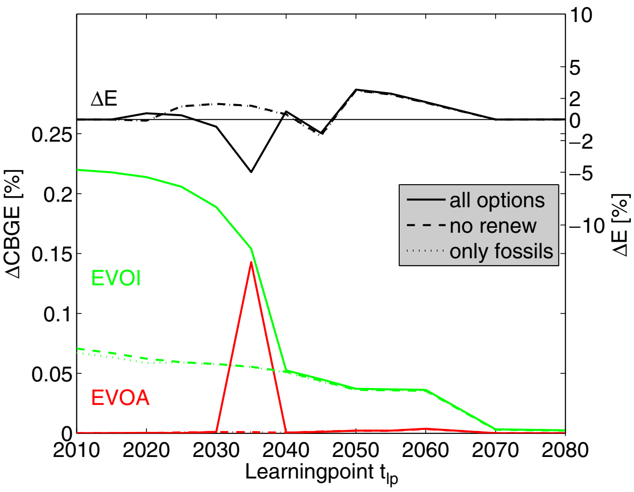

Fig. 5 Expected value of information (EVOI), expected value of anticipation (EVOA), and relative changes in cumulative pre-learning emissions in anticipation of learning (ΔE) are shown depending on the time of learning t_{lp}. Shown are three scenarios differing in the availability of mitigation options. In the "no renew" case, the usage of renewable energy is restricted to be lower than in the business-as-usual case where renewables are only used in the twenty-second century to counter the scarcity of fossil energy. In the "only-fossils" case, other options, like investments into "R&D" in energy and labor efficiency, are also not available

3.2.4 Sensitivity of the "Anticipation Window"

Now we investigate the sensitivity of the anticipation window with respect to T_0, D_x and the flexibility of the decision maker to change emissions over time. The results are shown in Fig. 6a–c.

Dependence on Threshold Position T_0 With rising threshold specific temperature T_0, the maximum of the EVOI decreases because the threshold is crossed later in time and less mitigation efforts are needed to stay below the threshold. For the same reason, the anticipation window is pushed toward later learning points. As already discussed above, for $T_1 < T_0 < T_2$, which is the case for $T_0 \in [2, 2.3]$°C, anticipation occurs to stay below the threshold in the high-damage case. Now we compare this result with one for $T_0 > T_2$, where the threshold is avoided even in the no-learning case. In the latter case, there is no incentive for downward anticipation, but the before mentioned incentive for upward anticipation in order to optimize the good learning case occurs. This leads to an EVOA that slowly increases with t_{lp} up to a maximum beyond which a higher pre-learning deviation from the optimal no-learning path

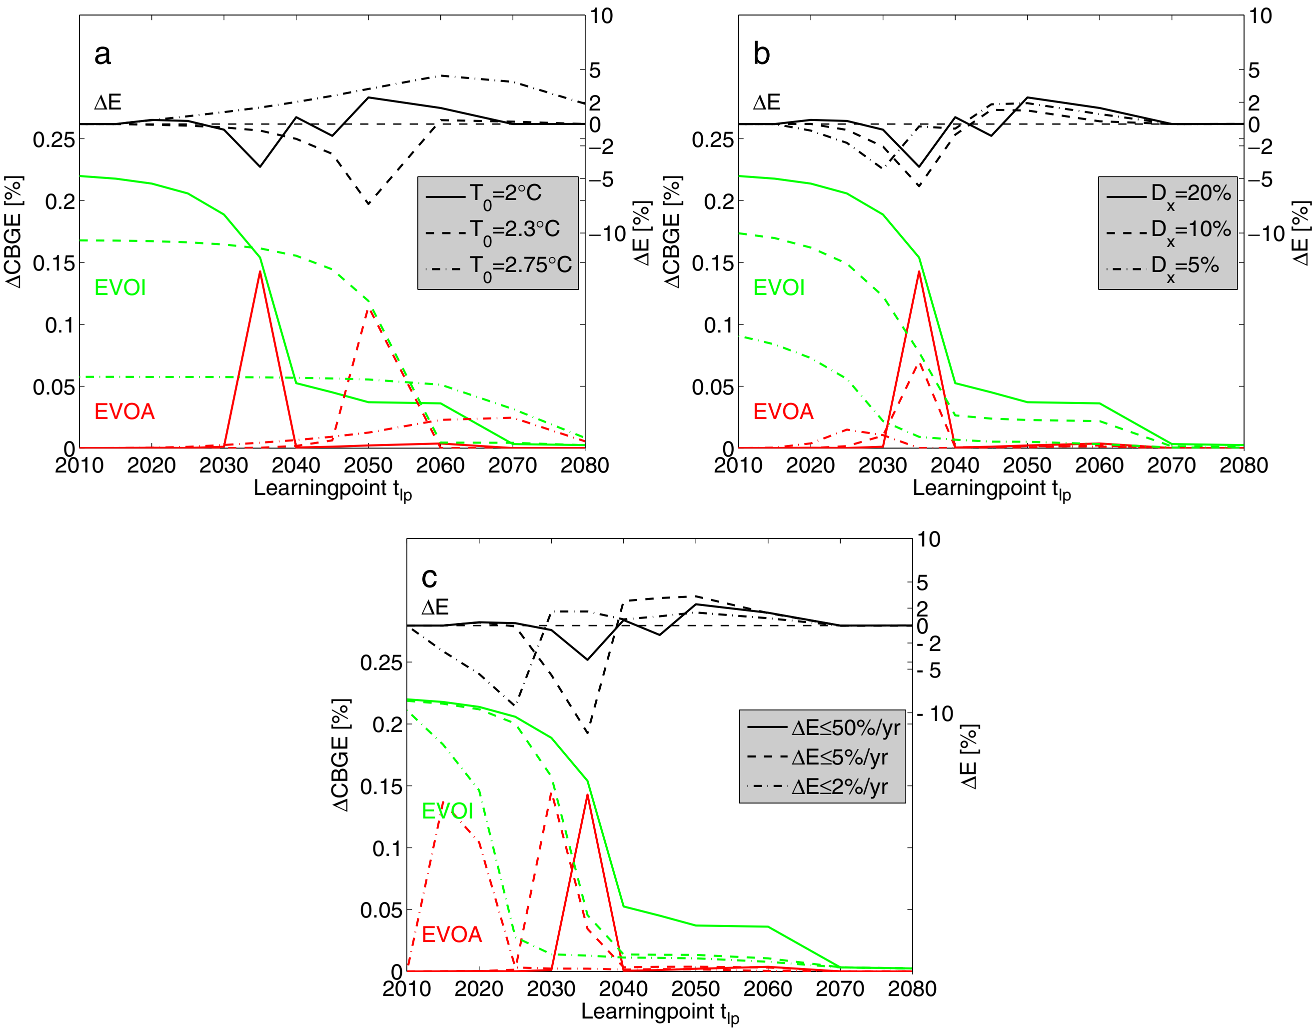

Fig. 6 Sensitivity of the anticipation effect: EVOI, EVOA, and the anticipatory relative changes in cumulative pre-learning CO_2 emissions are shown as a function learning time t_{lp}. **a** The dependence on different threshold temperatures T_0, **b** the dependence on the damage amplitude D_x, **c** the dependence on different exogenous inflexibilities ΔE_{max} of the decision maker in reducing or increasing emissions

leads to too high costs in the bad case. Although the absolute values of the EVOA and EVOI are smaller for high T_0, anticipation remains important in relative terms (EVOA/EVOI ratio).

Dependence on Mean Threshold Damages D_x Fig. 6 shows that both the EVOI and the EVOA are increasing in the threshold damage D_x. The anticipation window is slightly shifted toward earlier learning points for small threshold damages. This is due to the fact that the equilibrium between the mitigation costs to keep the threshold open in the bad case and the threshold damages is shifted toward lower values by decreasing D_x. The relative importance of anticipation remains large.

Dependence on an Artificial Emissions Flexibility ΔE_{max} The limited maximum emission flexibility ΔE_{max} is assumed to originate in processes that are not represented in the model, such as political and socioeconomic inertia. The first effect of limited flexibility is to move the curve toward lower values of t_{lp}. Since the ability to react to new information is now limited, anticipation becomes necessary for earlier learning times. In the limit of very low flexibility ($\Delta E < 1\%$/year) (not shown), the EVOA vanishes and even the EVOI for perfect learning in 2010 decreases as the decision maker cannot avoid crossing the threshold. In this case of low flexibility, the information can only be used to postpone the crossing of the threshold to later times by reducing emissions, but not to avoid the threshold.

4 Conclusions

We first introduced and clarified some terminology that can be used to assess the importance of anticipation of future learning. In particular, we introduced the concept of an expected value of anticipation.

We then investigated future learning about two key parameters of the climate problem, climate sensitivity and climate damages. We used the integrated assessment model MIND to calculate the welfare benefits from learning and the implications of anticipation of future learning for optimal near-term climate policy in terms of changes in the cumulative pre-learning emissions. The welfare benefits from learning were significant but benefits due to anticipation of this learning were not. This confirmed previous results in the literature.

We then investigated anticipated learning about uncertain threshold damages. The anticipation of learning leads to both higher and lower pre-learning emissions depending on the severity and position of the threshold. The welfare gains from this anticipation were in general considerably higher for downward anticipation (lower pre-learning emissions) than for upward anticipation (higher pre-learning emissions).

However, anticipation was only important if learning occurred within a specific, narrow time window, which depended on the flexibility of the decision maker to reduce and increase emissions. Inside this window, the welfare benefits due to anticipation can contribute almost the entire value of information ($\approx 95\%$). The strongest anticipation effect on pre-learning emissions did in general not lead to the strongest welfare gain. There was even one point in time such that learning at this point leads to two equally preferred solutions whereof one avoids the threshold and the other one does not.

The existence of a significant anticipation effect rested on the assumption of highly nonlinear damages and the availability of a flexible, scalable, and relatively cheap substitute for fossil energy. However, the anticipation effect was increased if the flexibility of adjusting emissions was reduced by other means than the availability of renewable energy. We showed this by introducing exogenous constraints on emissions changes motivated as political constraints or processes not represented in the model.

The analysis we have performed is only semi-quantitative and conclusions come with some caveats. The known limitations of all integrated assessment models with their highly simplified representation of the socioeconomic and physical processes apply. The representation of the threshold, the resulting damages, flexibility, uncertainty, and the learning process (as one-time perfect learning) could certainly be improved. More complex learning processes could be studied by changing toward a dynamic programming framework. Studying multiple and partly reversible thresholds occurring at uncertain temperatures could lead to more complex pattern of anticipation. All this, of course, would make the numerical solution more difficult.

Beside these limitations, a clear implication for real world climate policy can be drawn from our study: Although we are actually uncertain about both the position of potential thresholds as well as about their economic impacts, anticipating uncertain thresholds can be an important argument for lower emissions but not higher emissions.

Acknowledgements We are grateful for the helpful comments of two anonymous reviewers. A.L. acknowledges support by the German National Science Foundation and M.G.W.S. acknowledges funding by the BMBF project PROGRESS (03IS2191B).

Appendix 1: Climate Sensitivity

The climate module of MIND calculates the temperature response to anthropogenic forcing induced by CO_2 and SO_2 (which are coupled to CO_2 emissions) and exogenous forcing from other greenhouse gases:

$$\dot{T} = \mu \left(\ln\left(C/C_{\text{pi}}\right) + f_{\text{SO}_2} + f_{\text{OGHG}}\right) - \alpha T\,, \tag{13}$$

where C is current and C_{pi} pre-industrial atmospheric CO_2 concentration, T denotes global mean temperature anomaly, and μ the radiative forcing for a doubling of pre-industrial atmospheric CO_2 content divided by the heat capacity of the ocean (dominating the inertia of the climate system) and $\ln 2$. The parameter α is the response rate of the climate to changes in radiative forcing. It is linked to climate sensitivity CS via:

$$\text{CS} = \frac{\mu}{\alpha}\ln 2\,. \tag{14}$$

Actually, both μ and α in the temperature equation are uncertain and correlated via the global mean temperature record of the last two centuries (e.g., see Forest et al. [10]; Frame et al. [11]). For simplicity, we assume a perfect correlation and $\frac{1}{\mu} = \frac{1}{\bar{\mu}} - 10 \cdot \exp(-0.5\,\text{CS})$. The acceptability of this assumption can be assessed in Fig. 7.

The temperature response is now fully determined by CS. As prior information about CS we take a log-normal distribution from Wigley and Raper [41]: $\bar{\pi}(\text{CS}) = \mathcal{LN}(0.973, 0.4748)$.

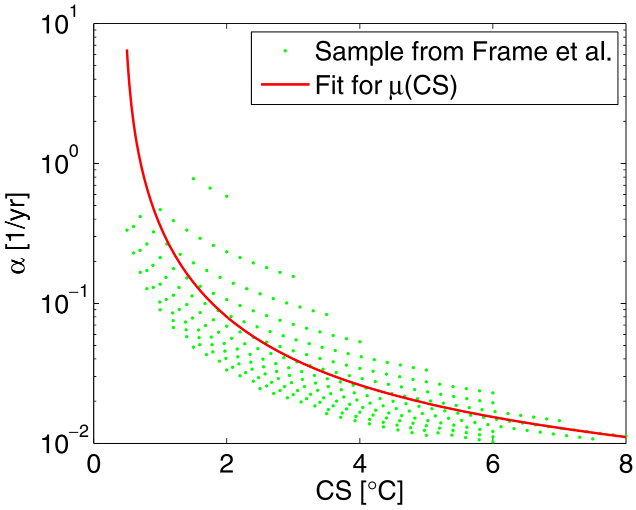

Fig. 7 Correlation of α and CS from the temperature record of the last two centuries (see Frame et al. [11]) [*green dots*]. By assuming a strict relationship between $\frac{1}{\mu}$ and CS as $\frac{1}{\mu} = \frac{1}{\bar{\mu}} - 10 \cdot \exp(-0.5\,\mathrm{CS})$ the correlation narrows to the [*red curve*]

Appendix 2: Climate Damages

The uncertain parameters a and b in the exponential damage function $\mathrm{DF}(T) = \frac{1}{1+a\,T^b}$ are determined from an expert-based assessment done by Roughgarden and Schneider [33]. They provide a joint probability distribution for both parameters. We use their methodology to derive the damage functions that are representative for the quantiles described by the sampling probability weights ω_i. Figure 8 shows the damage functions that represent the quantiles chosen for our experimental setup: $\omega_i = [1, 20, 60, 18, 1]\%$.

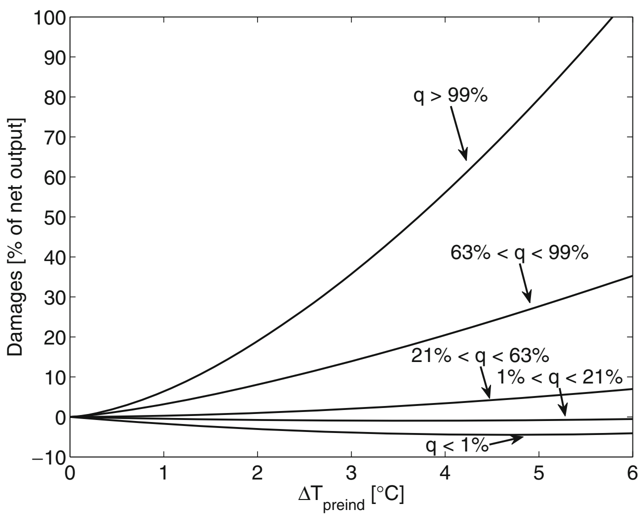

Fig. 8 Samples taken according to a descriptive sampling scheme from a joint probability distribution of the damage function parameters a and b from Roughgarden and Schneider [33]. Shown are the damage functions representative for the quantiles q with probability weights $\omega_i = [1, 20, 60, 18, 1]\%$ that have been used within the experiments

References

1. Anthoff, D., & Tol, R. (2009). The impact of climate change on the balanced growth equivalent: An application of FUND. *Environmental & Resource Economics, 43*(3), 351–367.
2. Baker, E. (2006). Increasing risk and increasing informativeness: Equivalence theorems. *Operations Research, 54*(1), 26–36.
3. Bauer, N. (2005). *Carbon capture and sequestration: An option to buy time?* Ph.D. thesis, Potsdam University, Germany.
4. Blackwell, D. (1951). Comparison of experiments. In Neyman, J. (Ed.), *2nd Berkeley symposium on mathematical statistics and probability*. Berkeley: University of California Press.
5. Bosetti, V., Carraro, C., Sgobbi, A., & Tavoni, M. (2008). *Delayed action and uncertain targets. How much will climate policy cost?* Technical report, CESifo Group Munich.
6. Broecker, W. S. (1997). Thermohaline circulation, the Achilles heel of our climate system: Will man-made CO2 upset the current balance? *Science, 278*(5343), 1582–1588.
7. Edenhofer, O., Bauer, N., & Kriegler, E. (2005). The impact of technological change on climate protection and welfare: Insights from the model MIND. *Ecological Economics, 54*, 277–292.
8. Edenhofer, O., Lessmann, K., & Bauer, N. (2006). Mitigation strategies and costs of climate protection: The effects of ETC in the hybrid model MIND. *Energy Journal, 27*, 207–222.
9. Epstein, L. G. (1980). Decision making and the temporal resolution of uncertainty. *International Economic Review, 21*(2), 269–83.
10. Forest, C. E., Stone, P. H., Sokolov, A. P., Allen, M. R., & Webster, M. D. (2002). Quantifying uncertainties in climate system properties with the use of recent climate observations. *Science, 295*(5552), 113–117.
11. Frame, D. J., Booth, B. B. B., Kettleborough, J. A., Stainforth, D. A., Gregory, J. M., Collins, M., et al. (2005). Constraining climate forecasts: The role of prior assumptions. *Geophysical Research Letters, 32*(9), L09702.
12. Gollier, C., Jullien, B., & Treich, N. (2000). Scientific progress and irreversibility: An economic interpretation of the 'Precautionary principle'. *Journal of Public Economics, 75*(2), 229–253.
13. Held, H., Kriegler, E., Lessmann, K., & Edenhofer, O. (2009). Efficient climate policies under technology and climate uncertainty. *Energy Economics, 31*, S50–61.
14. Johansson, D. J. A., Persson, U. M., & Azar, C. (2008). Uncertainty and learning: Implications for the trade-off between short-lived and long-lived greenhouse gases. *Climatic Change, 88*(3–4), 293–308.
15. Jones, R. A., & Ostroy, J. M. (1984). Flexibility and uncertainty. *Review of Economic Studies, 51*(1), 13–32.
16. Keller, K., Bolker, B. M., & Bradford, D. F. (2004). Uncertain climate thresholds and optimal economic growth. *Journal of Environmental Economics and Management, 48*(1), 723–741.
17. Kelly, D. L., & Kolstad, C. D. (1999). Bayesian learning, growth, and pollution. *Journal of Economic Dynamics & Control, 23*(4), 491–518.
18. Kolstad, C. D. (1996). Learning and stock effects in environmental regulation: The case of greenhouse gas emissions.

Journal of Environmental Economics and Management, 31(1), 1–18.
19. Kriegler, E. (2009). Updating under unknown unknowns: An extension of Bayes' rule. *International Journal of Approximate Reasoning, 50*, 583–596.
20. Kriegler, E., Held, H., & Bruckner, T. (2007). Climate protection strategies under ambiguity about catastrophic consequences. In J. Kropp, J. Scheffran (Eds), *Advanced methods for decision making* (pp. 3–42). New York: Nova Science.
21. Lange, A., & Treich, N. (2008). Uncertainty, learning and ambiguity in economic models on climate policy: Some classical results and new directions. *Climatic Change, 89*(1), 7–21.
22. Leach, A. J. (2007). The climate change learning curve. *Journal of Economic Dynamics & Control, 31*(5), 1728–1752.
23. Marschak, J., & Miyasawa, K. (1968). Economic comparability of information systems. *International Economic Review, 9*(2), 137–174.
24. Mirrlees, J. A., & Stern, N. H. (1972). Fairly good plans. *Journal of Economic Theory, 4*(2), 268–288.
25. Nordhaus, W. (2007). *The challenge of global warming: Economic models and environmental policy.* New Haven: Yale University.
26. Nordhaus, W. D. (1993). Rolling the DICE—an optimal transition path for controlling greenhouse gases. *Resource and Energy Economics, 15*(1), 27–50.
27. Nordhaus, W. D., & Popp, D. (1997). What is the value of scientific knowledge? an application to global warming using the PRICE model. *Energy Journal, 18*(1), 1–45.
28. O'Neill, B., Ermoliev, Y., & Ermolieva, T. (2006). Endogenous risks and learning in climate change decision analysis. In *Coping with uncertainty, modeling and policy issues* (pp. 283–300). Berlin: Springer.
29. O'Neill, B. C., & Melnikov, N. B. (2008). Learning about parameter and structural uncertainty in carbon cycle models. *Climatic Change, 89*(1–2), 23–44.
30. Oppenheimer, M., O'Neill, B., & Webster, M. (2008). Negative learning. *Climatic Change, 89*(1), 155–172.
31. Peck, S. C., & Teisberg, T. J. (1993). Global warming, uncertainties and the value of information—an analysis using CETA. *Resource and Energy Economics, 15*(1), 71–97.
32. Petschel-Held, G., Schellnhuber, H. J., Bruckner, T., Toth, F. L., & Hasselmann, K. (1999). The tolerable windows approach: Theoretical and methodological foundations. *Climatic Change, 41*(3–4), 303–331.
33. Roughgarden, T., & Schneider, S. H. (1999). Climate change policy: Quantifying uncertainties for damages and optimal carbon taxes. *Energy Policy, 27*(7), 415–429.
34. Salanie, F., & Treich, N. (2009). *Option value and flexibility: A general theorem with applications.* Technical report, LERNA, University of Toulouse.
35. Saliby, E. (1997). Descriptive sampling: An improvement over Latin hypercube sampling. In S. Andradóttir, K. J. Healy, D. H. Withers, & B. L. Nelson (Eds.), *Proceedings of the 1997 winter simulation conference*.
36. Schmidt, M. G. W., Lorenz, A., Held, H., & Kriegler, E. (2011). Climate targets under uncertainty: Challenges and remedies. *Climatic Change: Letters, 104*(3–4), 783–791.
37. Tol, R. (1998). *Potential slowdown of the thermohaline circulation and climate policy.* Institute for Environmental Studies Vrije Universiteit Amsterdam, Discussion Paper DS98/06.
38. Ulph, A., & Ulph, D. (1997). Global warming, irreversibility and learning. *Economic Journal, 107*(442), 636–650.
39. Webster, M. (2002). The curious role of "learning" in climate policy: Should we wait for more data? *Energy Journal, 23*(2), 97–119.
40. Webster, M., Jakobovits, L., & Norton, J. (2008). Learning about climate change and implications for near-term policy. *Climatic Change, 89*(1–2), 67–85.
41. Wigley, T. M. L., & Raper, S. C. B. (2001). Interpretation of high projections for global-mean warming. *Science, 293*(5529), 451–454.
42. Yohe, G., & Wallace, R. (1996). Near term mitigation policy for global change under uncertainty: Minimizing the expected cost of meeting unknown concentration thresholds. *Environmental Modeling and Assessment, 1*(1), 47–57.

Chapter 5

Synthesis and Outlook

The focus of this thesis is the assessment of the role of accounting for uncertainty and learning in the integrated assessment of optimal global climate change mitigation policy. Within the integrated assessment of climate change, the climate problem is framed as a decision problem for a representative agent (decision maker; DM) who aims at maximizing global social welfare by optimally chosing a level of greenhouse gase mitigation that balances the economic costs of mitigation and the climate change induced damages. In this framework the decision maker is facing substantial uncertainties in all parts of the underlying cause-effect chain. The importance of accounting for uncertainty can be assessed by comparing a situation in which the decision maker (DM) faces multiple possible outcomes of his decisions, each with a known probability of occurrence, to the situation in which the DM is certain to face the expected outcome. Some of the DM's uncertainty about the socio-economic and the climate system will be resolved by future observations and scientific progress. The importance of anticipating these future learning possibilities about uncertainty can be assessed by comparing two situations: The case of uncertainty is compared to a situation in which the DM initially is uncertain about the outcomes, but over time receives new information about the probabilities of the outcomes and thus can adjust his or her decision towards the new situation. In this context, the three main research questions are: how important is uncertainty for optimal climate mitigation policy? How important is the prospect of future learning for optimal climate mitigation policy? And which decision frameworks are applicable to analyzing optimal climate policy under uncertainty and learning? This thesis interprets the questions for the importance of accounting for uncertainty and learning along the line that uncertainty and learning are important for the analysis if their inclusion significantly changes both, the optimal emission policy itself and the net welfare gain from adopting an optimal climate policy instead of following the business as usual approach. As the answers to the first two questions strongly depend on the decision framework the climate problem is stated in, the third question underlies both of the others.

Contributions to the answers of all three questions have been made within this thesis. They are summarized in the following. The final section concludes with a general outlook and future research requirements.

5.1 Formulating the Climate Problem under Uncertainty and Learning

The most commonly used framework within the integrated assessment modeling of climate change under uncertainty is the expected utility (EU) maximization. This is due to the fact, that the von Neumann-Morgenstern axioms underlying the EU framework represent widely accepted norms of rational decision making under uncertainty (Machina, 1987). The straight forward implementation of the EU framework is given by the so-called cost-benefit analysis (CBA), that incorporates the costs of mitigation as well as a monetized representation of climate induced damages and thus the benefits of mitigation. Two of the main criticisms of CBA are the pure inability to monetize certain non-market damages, and more broadly, the incommensurability of different kinds of damages, e.g. the comparison of losses of biodiversity to damages to infrastructure.

Another EU implementation, the also often used cost-efficiency analysis (CEA) refrains from monetizing climate induced damage directly and thereby avoids the criticism. The trade-off between mitigation and benfits of avoiding certain levels of climate change is conducted implicitly by choosing a certain limit (or target) in a variable describing the amplitude of climate change which then constrains the EU maximization.

In the more recent past, many studies have calculated the implications of a certain temperature target, e.g. the 2°C target adopted by the UN in Cancun (UNFCCC, 2010). Chapter 2 confirms findings from the older literature in decision theory that this cost-efficiency approach runs into axiomatic problems if uncertainty and learning are included in the analysis. As argued by Held et al. (2009), for applying a climate target under uncertainty the information about the maximum temperature to reach is not sufficient, but it has to be accompanied by a measure of probability with which the target is to be kept, as it is in principle not possible to keep the "whole distribution" of temperatures below a given threshold. This method, of constraining not temperature but rather probability for crossing a temperature, is called chance constraint programming. Chapter 2 shows that this chance constraint programming framework can lead to normatively unappealing consequences when future learning is included. These consequences are the possibility of negative expected value of information and the possibility of infeasibility of the decision problem due to learning, both are derived from an axiomatic analysis and are demonstrated within the context of an intuitive example. The possibility of negative value of future information stems from the fact that the framework does not include an explicit tradeoff between the risk of crossing the temperature threshold and the associated costs. Keeping the target fixed irrespective of what is learned is more costly in a convex, nonlinear way if "extreme" messages are received. Thus a decision maker complying with this framework would actually pay for not having to receive new information, which is not desirable from the common normative perspective. The potential infeasibility of the decision problem due to uncertainty results from the fact that in some future learning scenarios it is not only more costly, but impossible to maintain the given probabilistic climate target. The point here is that the decision criterion is incomplete as it does not give advice on how to proceed in this situation.

Chapter 2 continues by proposing an alternative decision criterion that does not use a fixed probabilistic climate target but allows for, and requires, a trade-off between mitigation costs and the risk of exceeding the target. This so-called cost-risk analysis (CRA) pays tribute to the fact that some kind of trade-off between the "bad" and "good" consequences of a decision needs to be included for each state of the world, to make the EU maximization feasible under uncertainty and learning. However, this trade-off does not necessarily include a complete monetization of climate damage but can be done on a more aggregated level. It may also still include the notion of a climate target. The information requirements for assessing the DM's preferences are equally high in a CBA and a CRA. The question of which formulation is more practical, in that real decision makers really behave as if following one of the problem formulations and the factual estimation of the necessary preference parameters is left for further studies.

There are other critiques of the standard EU framework, two of which are the inability to handle so-called structural, or deep, uncertainty, and the representation of different kinds of aversion against large differences in consumption (across states of the world, time, people, etc.) with only one parameter. The former is a feature in observed decision-making of people being averse to not knowing the probabilities of the outcomes they have to decide about. Several frameworks have been proposed to handle this aspect of deep uncertainty from which the model of ambiguity due to Klibanoff et al. (2009) seems to be the most promising one. However, due to computational complexity, up to now, the model still awaits its application in the more complex integrated assessment models of climate change. The latter problem of disentangling the parameters of relative risk aversion from the parameter of inter-temporal substitution elasticity in consumption has been on the agenda at least since the 1970s. The necessity to do so arises from the observation that real world decision makers reveal different levels of aversion to both effects. Thus, a consistent disentangling of both parameters could deliver an explanation of several observed paradoxes. Kreps & Porteus (1978) first proposed such a framework. It allows the desired disentanglement but it introduces an intrinsic preference of the DM for early or late resolution of uncertainty. This is again normatively undesirable as it means that the DM would even pay money for useless information, i.e. if she could not adjust her actions according to the new information. This deficiency has currently been overcome by a framework that simultaneously solves both problems: disentanglement of aversion parameters, in combination with neutrality towards the timing of uncertainty resolution (see Traeger, 2009, and references therein). As with the Klibanoff model for ambiguity, the widespread application of the Traeger model, and especially the combination of both approaches, proposes very exciting challenges for further research.

But even if the correct formulation of the decision problem for the integrated assessment of climate change has been found, and one could argue that with the models mentioned above we are close to getting there, the challenge remains of finding, or estimating, or choosing, the normative parameters. The outcomes of the integrated assessment and the role of uncertainty and learning rest upon the choice of those parameters. The philosophical debate is still open as to whether one should measure those parameters in the markets, as done by Nordhaus (2008b), or set them normatively as done by Stern (2007).

5.2 Importance of Uncertainty for Global Climate Analysis

There is a widespread intuition, at least amongst scientists, that uncertainty surrounding climate change and its potential impacts is a crucial element of the problem of climate change. Hence, intuitively, the explicit inclusion of uncertainty about, say, climate sensitivity and climate damage amplitude into an integrated assessment model should have a strong impact on both the optimal global mitigation policy and the resulting net welfare benefits of acting upon climate change. However, from the standard IAMs, which formally implement a CBA, this result is not supported (see e.g. Nordhaus & Popp, 1997; Saphores, 2004). The standard solution of a low carbon price that moderately increases over time, the so-called policy ramp, does not change significantly when uncertainty is included.

Several changes to the model formulation have been proposed that lead to considerable effects from uncertainty, such as the consideration of fat-tailed probability distributions by Weitzman (2010), or including heterogeneous damage, as in Schmidt et al. (2011a). However, only a few other studies have investigated the influence of including uncertainty (and perfect learning) on the welfare gain from acting upon climate change, e.g. Pizer (1997). Chapter 3 investigates the origin of the negligible welfare effect in a "standard" IAM and proposes several changes to the structure of the IAM itself that would lead to considerably higher welfare effects from the inclusion of uncertainty.

The overall net welfare benefit of acting optimally with respect to climate change under uncertainty is decomposed into three components: The benefit of optimally acting on climate change under certainty, the change of the benefit of this action due to the inclusion of uncertainty, and the benefit of adjusting the action from the optimal action under certainty to the optimal action under uncertainty. It is proposed to use the last term, the welfare benefit of adjusting the mitigation action to uncertainty, relative to the overall welfare benefit of acting on climate change as a metric for the importance of including uncertainty into the integrated assessment of climate change.

Furthermore, the IAM MIND is projected onto an a-temporal marginal cost-benefit picture of cumulative emission reductions that lead to marginal mitigation costs and marginal benefits from reducing climate damages. This picture allows linking the elementary components of the benefit of climate policy (BCP) to the functional structure (temperature response, consumption losses due to damages and mitigation effort, and welfare effects of the former, all depending on the level of cumulative emission reductions) of the climate cause and effect chain within MIND.

The key finding is that the welfare gain from explicitly including uncertainty is only significant if it leads to nonlinear shifts in the marginal functions within the cause-effect chain, e.g. a nonlinear increase in marginal damages (marginal w.r.t. cumulative emission reductions). Such a change in turn leads to a difference in the expected marginal net welfare benefits of mitigation under uncertainty compared to certainty, the so-called marginal risk premium, that is convex increasing. In this situation the additional benefit from adjusting the mitigation level to the situation of uncertainty can take up a significant part of the overall net benefit of mitigation, thus uncertainty is important.

These necessary shifts in the marginal functions within the climate cause-effect chain originate from the nonlinear damage function and the nonlinear utility function. The function of maximum

temperature increase in cumulated mitigation efforts is only weakly nonlinear, and uncertainty does not change the form of the marginal function, only shifts it slightly upward. This weak nonlinearity explains why uncertainty about climate sensitivity, however prominently discussed, has only a very small influence on optimal policy decisions.
Another important finding is that for the importance of uncertainty the *curvature* of the marginal functions in the climate cause-effect chain matters. Here the curvature of the maximum temperature function in cumulative mitigation efforts and the curvature of the corresponding damage function compensate each other to some extent and thereby weaken the overall impact of uncertainty.
Generalising this result, it shows how the structure of one part of the cause-effect chain can influence the importance of uncertainty within another part. The effect of including uncertainty in climate sensitivity on the choice of optimal climate policy and the resulting increase in welfare is small. However, if changing the damage function from a quadratic to a more nonlinear shape, the importance of CS uncertainty increases.

The utility function and thus the normative parameters of pure time preference and of constant relative risk aversion also have a very strong impact on the importance of uncertainty, as well as on the overall net welfare benefit from acting upon climate change. This also corroborates the findings in the literature, e.g. by Pizer (1997), that uncertainty about normative parameters would by far dominate uncertainty about the climate system.

Due to the limitations of the model employed within our analysis, Chapter 3 can provide no general answer to the question of the importance of uncertainty for the integrated assessment, but what it does offer is a comprehensive set of measures to assess the importance of the different effects of uncertainty and learning separately within a fairly complex integrated assessment model. Future research will investigate the applicability of the approach to even more complex models which comprise more realistic representations of the climate system. The study is also to be expanded to the impact of other uncertainties, as it already has been shown that a combination of two single parameter uncertainties is in no way additive in the aggregated outcome. A similar model decomposition could be tested for the assessment of anticipative learning (as partly done in Chapter 4) and for the inclusion of other changes to the model, like the introduction of regional inequity, new industry sectors, changes in the climate representation, etc.

5.3 Importance of Anticipating Future Learning

As already demonstrated with respect to the importance of uncertainty, this thesis is not only concerned with the question of the change of near-term decisions due to anticipation of future learning, but also with the resulting changes in welfare gains from climate policy. To investigate these changes, the notion of the so-called expected value of anticipation is introduced by decomposing the overall expected value of future information into the welfare gain from pre-learning and post-learning decision adjustments.

This decomposition is then applied to the study of future learning about climate sensitivity and amplitude of climate induced damage within a simple exogenous one time learning framework with the integrated assessment model MIND. Thereby we confirm the findings in the literature that although future learning about those quantities has a significant overall expected value of information the value of anticipating this future learning by changing near-term decisions is negligible. Chapter 4 continues by introducing uncertainty and future learning about the amplitude of additional damage stemming from the crossing of a tipping point like threshold in temperature, representing e.g. a breakdown of the north Atlantic thermohaline circulation as introduced by Keller et al. (2004), or the partial destabilization of the west Antarctic ice sheet. We show that anticipatory changes in near-term emissions towards more mitigation can become crucial to harvest the value of learning. This is the case, if the value of new information cannot be used after the learning occurred, due to the fact, that the DM would already be committed to cross the threshold by pre-learning actions. Thus the expectation of future learning alone, without knowing what will be learned, leads to stronger pre-learning mitigation action, thereby keeping the option to mitigate the threshold in case the threshold damages would be severe. In this case almost the complete value of information is due to anticipation. However, this significant value of anticipation only occurs if the learning takes place in a specific “anticipation window”. If learning happens earlier, the whole value of information can be simply harvested by post-learning adjustments. If learning happens later, anticipatory changes necessary to maintain the option to avoid the threshold become too expensive and thus the DM already commits herself to crossing the threshold irrespectively of what is learned. This significantly lowers the overall value of information and thus also the value of anticipation. Within the standard model setting of MIND the “anticipation window” is quite narrow, spanning approximately one decade. This is due to the high flexibility of the model to change emissions on short notice to moderate costs. The location and width of the window is strongly sensitive to this flexibility. When introducing an additional inflexibility into the model, whether it represents political or social barriers to fast emission reductions, the anticipation window moves towards the present and becomes broader. The maximum value of anticipation decreases, as the inflexible emission trajectories are necessarily suboptimal compared to the flexible ones.

A straight forward extension of the work of Chapter 4 would be the inclusion of multiple thresholds that might depend on different climate variables and not only on global mean temperature. The estimates of damages arising from the crossing of the thresholds should be revised in the light of new research on climate change impacts. However, the main result that future learning only influences near-term decisions if strong nonlinearities, e.g. from irreversibility, come into play, will remain. Another extension would be to replace the simple one-time perfect learning by a more advanced representation of the learning process, whether it is learning at multiple points in time or even endogenous learning. The latter in particular might lead to stronger dependencies between future learning possibilities and short-term decisions as the expected marginal welfare after learning is then potentially influenced by both, potential path dependencies (stock effects, thresholds, irreversibilities) from the first period, and changes in the post-learning probability distribution due to first period actions. However, the formal requirements in terms of computational needs (like efficiently implementing dynamic programming for large models) and informational needs (formalizing endogenous learning functions for different technologies and sectors) to represent the uncertainties and learning processes are huge; at this point of time they seem prohibitively so.

One very interesting extension of the work would be its application to the so-called Dismal Theorem by Weitzman (2009). He proposed that under certain conditions, that include fat-tailed probability distributions for uncertain climate responses, exponential damage from climate change and sufficiently high risk aversion of the DM, the marginal utility of mitigation can be unbounded. This means that any single ton of mitigated carbon would be worth nearly the whole economic output, or putted the other way round, the DM would spend nearly the whole economic output to avoid a very unlikely but catastrophic tail event. Several critics of the applicability of Weitzman's argument have been put forward, e.g. by Nordhaus (2009). Another crucial assumption that has received less attention, is the fact that Weitzman's DM acts in an a-temporal framework, i.e. she cannot change her decisions in the mid-term. Thus an argument against the Dismal Theorem put forward (not formally) by Myles Allen and David Frame (communication at the 2011 Tanner lecture, on 2011-05-21 in Oxford, UK) refers to the possibility of mid-term corrections once the DM can foresee that the unlikely tail events are becoming reality. From the perspective of Chapter 4 this argument refers to the flexibility of the decision maker to harvest the (really large) expected value of future information from learning about the (potentially catastrophic) climate damage nearly completely by post-learning decisions. From what we have learned from Chapter 4, however, this is only true when either the climate damage are reversible or the learning leading to mid-term corrections takes place early enough to be outside the anticipation window. If the DM finds herself inside the anticipation window, the learning about the potentially catastrophic damages is only valuable if she reacts in the short-term, thus the fat tails again have a paramount influence on current decision making. As the processes that may lead to catastrophic climate damage are arguably all highly nonlinear and include significant hysteresis behavior, only a situation outside the anticipation window could be used as an argument against Weitzman. Thus the open question again is an empirical one, for the position of the thresholds in the system and for the possibility of our future ability to cope with irreversibility by means of geo-engineering or climate control options.

Another interesting extension of this work, especially in the context of potentially catastrophic climate change is the combination of the framework from Chapter 4 with the preference structure by Traeger (2009), mentioned above. In his framework discounting due to pure impatience is no longer allowed, instead the DM discounts the future due to increasing uncertainty. The interesting part of this extension is that now different future scenarios are discounted differently according to the confidence in this specific projection. Thus the application of Traeger's framework to catastrophic climate change would probably deliver a solution for the problem of dominating catastrophic fat tails of climate damages. The framework established in Chapter 4 could be used to check the validity of such a proposition, especially when including mid-term corrections and endogenous learning as above.

5.4 General Outlook and further Research Questions

The results of this thesis rest on the general assumptions of the welfare economic framework employed. Thus a general direction for further research is to release or extend some of the underlying assumptions.

One possible expansion is the recognition of heterogeneity within different parts of the model setup. Schmidt et al. (2011a) investigate the influence of heterogeneous climate damage on the certainty premium a decision maker would be willing to pay for ruling out uncertainty. They find that if the climate damage is only imposed on a minor fraction of the population the risk premium can increase substantially. This can partly be countered by implementing efficient insurance markets or allowing for measures of self-insurance. In expansion one could be interested in additionally considering an initially heterogeneously wealthy population. This would surely again increase the effect of damage uncertainty.

The same is true when considering other sources of inequity between different actors and parts of the total population, like inequity between regions (e.g. see Anthoff et al. 2009).

An even stronger impact of uncertainty in terms of changes in welfare gains from different policies would be expected when adopting a framework put forward by Sterner and Persson (2008), who argued that climate impacts would be more severe when not hitting one aggregated global output providing sector but would hinder the production of very specific single sectors, like an envisaged "environmental good production". If the goods from those separate sectors where only limited in their possibility to being substituted for each other, damage to one sector could not only lead to limited growth or even recession of this single sector but due to the low elasticity of substitution, the overall welfare impact from uncertainty about damage would be highly increased.

The assumption of infinitely ongoing exponential aggregated economic growth itself is a very strong one. Besides the normative question of whether such a growth regime is a "good" thing, in that it optimally provides happiness, it is also questionable whether such an exponential trajectory is an adequate description of realistic assumptions about future economic growth. As the theoretical foundation of long-term economic growth is still not satisfying, a first step for further research about systematic uncertainty would be a sensitivity study with respect to the assumption of exponential economic growth.

Gerst et al. (2010) go in this direction, by modelling economic growth as an exponential process with a stochastic yearly growth rate. However, one could argue, as done e.g. by Ayres & Warr (2009), that the processes driving economic growth are not stochastic in nature but result from several processes that are partly continuous and partly discontinuous in nature. These are population change, expansion and interconnection of markets, and technological, scientific, and societal innovation. To include those processes in a satisfying way into a general framework of long-term economic growth proposes an enormous challenge to the economic profession. Once this is achieved, the climate problem, or better, the ecological constraints of our planet, can be understood as (temporary) boundaries that induce technological change. So it might even be the other way round, that the climate challenge and the planetary constraints work as drivers for the economy and overcoming them, whether by behavioral change, by decarbonizing the economy, introduction of climate management or even the expansion of markets beyond the planetary boundaries through access to cheap space travel, will be what ensures ongoing exponential growth in the future. To investigate the role of scientific uncertainty about the system boundaries of the earth system in such an integrated growth framework is an exciting challenge for future research.

Summarizing, this thesis has contributed to evaluating the importance of including uncertainty and learning into the integrated assessment of climate change mitigation policy. This importance turns out to be low under standard assumptions about the cause-effect chain of climate change. However, several plausible structural changes in the representation of the climate cause-effect chain have been positively tested with respect to their potential to significantly increase the importance of uncertainty and learning. This thesis has also provided the means to conduct more of these tests with regard to changes proposed in this final chapter. Without waiting for the results of those tests, a, in this sense somewhat hasty, personal conclusion can be drawn from this thesis: that I personally tend to lean towards our initial intuition that uncertainty and learning are dominant aspects of climate change analysis and that our formal framework is far from being able to represent this important feature properly. Some important first steps have already been made to improve this situation. But another bold effort has to be undertaken, maybe even more so than the original idea to put together economic and physical models, to overcome the shortcomings of model representation of uncertainty and its impacts. Only than can the integrated assessment of climate change lead to policy implications that real world decision makers could base their real decisions upon.

References

(2010). *Report of the Conference of the Parties on its fifteenth session, held in Copenhagen from 7 to 19 December 2009*. UNFCCC.

Alvarez, L. W., Alvarez, W., Asaro, F., and Michel, H. V. (1980). Extraterrestrial cause for the Cretaceous-Tertiary extinction. *Science*, 208(4448):1095 –1108.

Anthoff, D., Tol, R. S. J., and Yohe, G. W. (2009). Risk aversion, time preference, and the social cost of carbon. *Environmental Research Letters*, 4(2):024002.

Arrhenius, S. (1896). On the influence of carbonic acid in the air upon the temperature on the ground. *Philosophical Magazine and Journal of Science*, 41:237–276.

Ayres, R. U. and Warr, B. (2009). *The Economic Growth Engine: How Energy and Work Drive Material Prosperity*. Edward Elgar Publishing.

Baker, E. (2006). Increasing risk and increasing informativeness: Equivalence theorems. *Operations Resaerch*, 54(1):26–36.

Bosello, F., Carraro, C., and Cian, E. D. (2010). Climate policy and the optimal balance between mitigation, adaptation and unavoided damage. *Climate Change Economics*, 1(2):71–92.

Bosetti, V., Carraro, C., Sgobbi, A., and Tavoni, M. (2008). Delayed action and uncertain targets. how much will climate policy cost? Technical report, CESifo Group Munich.

Boyd, P. W., Jickells, T., Law, C. S., Blain, S., Boyle, E. A., Buesseler, K. O., Coale, K. H., Cullen, J. J., de Baar, H. J. W., Follows, M., Harvey, M., Lancelot, C., Levasseur, M., Owens, N. P. J., Pollard, R., Rivkin, R. B., Sarmiento, J., Schoemann, V., Smetacek, V., Takeda, S., Tsuda, A., Turner, S., and Watson, A. J. (2007). Mesoscale iron enrichment experiments 1993-2005: Synthesis and future directions. *Science*, 315(5812):612 –617.

Broccoli, A. J., Dahl, K. A., and Stouffer, R. J. (2006). Response of the itcz to northern hemisphere cooling. *Geophys. Res. Lett.*, 33:L01702.

Conway, T. and Tans, P. (2011). Trends in atmospheric carbon dioxide. *www.esrl.noaa.gov/gmd/ccgg/trends/*, page accessed on 2011/06/11.

COP (2010). The cancun agreements: Outcome of the work of the ad hoc working group on long-term cooperative action under the convention. In *United Nations Framework Convention on Climate Change*.

Dakos, V., Scheffer, M., van Nes, E. H., Brovkin, V., Petoukhov, V., and Held, H. (2008). Slowing down as an early warning signal for abrupt climate change. *Proceedings of the National Academy of Sciences*, 105(38):14308 –14312.

Dasgupta, P. (2008). Discounting climate change. *Journal of Risk and Uncertainty*, 37(2):141–169.

Edenhofer, O., Wallacher, J., Reder, M., and Lotze-Campen, H. (2010). C.H. Beck.

Forster, P., Ramaswamy, V., Artaxo, P., Berntsen, T., Betts, R., Fahey, D., Haywood, J., Lean, J., Lowe, D., Myhre, G., Nganga, J., Prinn, R., Raga, G., Schulz, M., and Dorland, R. V. (2007). *Climate Change 2007: The Physical Science Basis. Contribution of Working Group I to the Fourth Assessment Report of the Intergovernmental Panel on Climate Change*, chapter Changes in Atmospheric Constituents and in Radiative Forcing. Cambridge University Press, Cambridge, United Kingdom and New York, NY, USA.

Fung, I. Y., Doney, S. C., Lindsay, K., and John, J. (2005). Evolution of carbon sinks in a changing climate. *Proceedings of the National Academy of Sciences*, 102(32):11201–11206.

Gerst, M. D., Howarth, R. B., and Borsuk, M. E. (2010). Accounting for the risk of extreme outcomes in an integrated assessment of climate change. *Energy Policy*, 38(8):4540–4548.

Gollier, C. (2004). *The Economics of Risk and Time*. The MIT press.

Hardin, G. (1968). The tragedy of the commons. *Science*, 162(3859):1243 –1248.

Held, H. and Kleinen, T. (2004). Detection of climate system bifurcations by degenerate fingerprinting. *Geophys. Res. Lett.*, 31:L23207.

Held, H., Kriegler, E., Lessmann, K., and Edenhofer, O. (2009). Efficient climate policies under technology and climate uncertainty. *Energy Economics*, 31:S50–61.

Hoffman, P. F., Kaufman, A. J., Halverson, G. P., and Schrag, D. P. (1998). A neoproterozoic snowball earth. *Science*, 281(5381):1342 –1346.

IEA (2011). Co2 emissions from fossil fuel combustion - highlights -. Technical report, International Energy Agency.

Ingham, A., Ma, J., and Ulph, A. (2007). Climate change, mitigation and adaptation with uncertainty and learning. *Energy Policy*, 35(11):5354–5369.

IPCC (2007). *Contribution of Working Groups I, II and III to the Fourth Assessment Report of the Intergovernmental Panel on Climate Change*, chapter Climate Change 2007: Synthesis Report, page 104. IPCC, Geneva, Switzerland.

Keith, D. (2001). Geoengineering. *Nature*, 409:420.

Keller, K., Bolker, B. M., and Bradford, D. F. (2004). Uncertain climate thresholds and optimal economic growth. *Journal of Environmental Economics and Management*, 48(1):723–741.

Keller, K., Yohe, G. W., and Schlesinger, M. (2007). Managing the risks of climate thresholds: uncertainties and information needs. *Climatic Change*.

Klibanoff, P., Marinacci, M., and Mukerji, S. (2009). Recursive smooth ambiguity preferences. *Journal of Economic Theory*, 144(3):930–976.

Kreps, D. M. and Porteus, E. L. (1978). Temporal resolution of uncertainty and dynamic choice theory. *Econometrica*, 46(1):185–200.

Lackner, K. S., Grimes, P., and Ziock, H.-J. (2008). Capturing carbon dioxide from air. In *Proceedings of the 2nd U.S.-China Symposium on CO2 Emissions Control Science & Technology*.

Leimbach, M. and Baumstark, L. (2010). The impact of capital trade and technological spillovers on climate policies. *Ecological Economics*, 69:2341–2355.

Lenton, T. M. (2011). Early warning of climate tipping points. *Nature Climate Change*, 1:201–209.

Lenton, T. M., Held, H., Kriegler, E., Hall, J. W., Lucht, W., Rahmstorf, S., and Schellnhuber, H. J. (2008). Tipping elements in the earth's climate system. *Proceedings of the National Academy of Sciences*, 105(6):1786 –1793.

MACHINA, M. J. (1987). Decision-Making in the presence of risk. *Science*, 236(4801):537 –543.

Meehl, G., Stocker, T., Collins, W., Friedlingstein, P., Gaye, A., Gregory, J., Kitoh, A., Knutti, R., Murphy, J., Noda, A., Raper, S., Watterson, I., Weaver, A., and Zhao, Z.-C. (2007). *Climate Change 2007: The Physical Science Basis. Contribution of Working Group I to the Fourth Assessment Report of the Intergovernmental Panel on Climate Change*, chapter Global Climate Projections. Cambridge University Press, Cambridge, United Kingdom and New York, NY, USA.

National Research Council (2011). *Climate Stabilization Targets: Emissions, Concentrations, and Impacts over Decades to Millennia*. The National Academies Press, Washington, D.C.

Nordhaus, W. (2008a). *A Question of Balance: Weighing the Options on Global Warming Policies*. Yale University Press, illustrated edition edition.

Nordhaus, W. and Popp, D. (1997). What is the value of scientific knowledge? an application to global warming using the PRICE model. *Energy Journal*, 18(1):1–45.

Nordhaus, W. D. (1994). *Managing the global commons*. MIT Press Cambridge, MA.

Nordhaus, W. D. (2006). The "Stern review" on the economics of climate change. Technical report, National Bureau of Economic Research, Inc.

Nordhaus, W. D. (2008b). *A question of balance*. Yale University Press.

Nordhaus, W. D. (2009). An analysis of the dismal theorem. Technical report, David K. Levine.

O'Neill, B. C. and Melnikov, N. B. (2008). Learning about parameter and structural uncertainty in carbon cycle models. *Climatic Change*, 89(1-2):23–44.

Peck, S. C. and Teisberg, T. J. (1993). Global warming, uncertainties and the value of information - an analysis using CETA. *Resource and Energy Economics*, 15(1):71–97.

Pizer, W. A. (1997). Optimal choice of policy instrument and stringency under uncertainty: The case of climate change. *Working Paper: Resources for the Future*, 97:17.

Ramanathan, V. and Feng, Y. (2008). On avoiding dangerous anthropogenic interference with the climate system: Formidable challenges ahead. *Proceedings of the National Academy of Sciences*.

Roe, G. H. and Baker, M. B. (2007). Why is climate sensitivity so unpredictable? *Science*, 318(5850):629 –632.

Saphores, J. M. (2004). Environmental uncertainty and the timing of environmental policy. *Natural Resource Modeling*, 17(2):163–190.

Schmidt, M. G. (2011). *Climate Policy under Uncertainty*. PhD thesis, TU Berlin.

Schmidt, M. G. W., Held, H., Kriegler, E., and Lorenz, A. (2011a). Stabilization targets under uncertain and heterogeneous climate damages. *submitted to Environmental & Resource Economics*.

Schmidt, M. G. W., Lorenz, A., Held, H., and Kriegler, E. (2011b). Climate targets under uncertainty: Challenges and remedies. *Climatic Change: Letters*, 104(3-4):783–791.

Schneider, S. H. and Mastrandrea, M. D. (2005). Probabilistic assessment of dangerous climate change and emissions pathways. *Proceedings of the National Academy of Sciences of the United States of America*, 102(44):15728–15735.

Solomon, S., Qin, D., Manning, M., Alley, R., Berntsen, T., Bindoff, N., Chen, Z., Chidthaisong, A., Gregory, J., Hegerl, G., Heimann, M., Hewitson, B., Hoskins, B., Joos, F., Jouzel, J., Kattsov, V., Lohmann, U., Matsuno, T., Molina, M., Nicholls, N., Overpeck, J., Raga, G., Ramaswamy, V., Ren, J., Rusticucci, M., Somerville, R., Stocker, T., Whetton, P., Wood, R., and Wratt, D. (2007). *Climate Change 2007: The Physical Science Basis. Contribution of Working Group I to the Fourth Assessment Report of the Intergovernmental Panel on Climate Change*, chapter Technical Summary. Cambridge University Press, Cambridge, United Kingdom and New York, NY, USA.

Stern, N. (2007). *The Economics of Climate Change: The Stern Review.* Cambridge University Press, UK.

Sterner, T. and Persson, U. (2008). An even sterner review: Introducing relative prices into the discounting debate. *Review of Environmental Economics and Policy*, 2(1):61–76.

Stott, P. A., Jones, G. S., Lowe, J. A., Thorne, P., Durman, C., Johns, T. C., and Thelen, J. (2006). Transient climate simulations with the HadGEM1 climate model: Causes of past warming and future climate change. *Journal of Climate*, 19(12):2763–2782.

Stroeve, J., Serreze, M., Holland, M., Kay, J., Maslanik, J., and Barrett, A. P. (2011). The arctic's rapidly shrinking sea ice cover: a research synthesis. *Climatic Change*, in press.

Traeger, C. P. (2009). Recent developments in the intertemporal modeling of uncertainty. *Annual Review of Resource Economics*, 1(1):261–286.

Trenberth, K., Jones, P., Ambenje, P., Bojariu, R., Easterling, D., Tank, A. K., Parker, D., Rahimzadeh, F., Renwick, J., Rusticucci, M., Soden, B., and Zhai, P. (2007). *Climate Change 2007: The Physical Science Basis. Contribution of Working Group I to the Fourth Assessment Report of the Intergovernmental Panel on Climate Change*, chapter Observations: Surface and Atmospheric Climate Change. Cambridge University Press, Cambridge, United Kingdom and New York, NY, USA.

Ulph, A. and Ulph, D. (1997). Global warming, irreversibility and learning. *ECONOMIC JOURNAL*, 107(442):636–650.

Vermeer, M. and Rahmstorf, S. (2009). Global sea level linked to global temperature. *Proceedings of the National Academy of Sciences*, 106(51):21527 –21532.

Victor, D. G., Morgan, M. G., Apt, J., Steinbruner, J. D., and Ricke, K. (2009). The geoengineering option: A last resort against global warming? *Foreign Affairs*.

Waggoner, P. and Ausubel, J. (2002). A framework for sustainability science: A renovated ipat identity. *Proceedings of the National Academy of Sciences*, 99(12):7860 – 7865.

Webster, M. (2002). The curious role of learning in climate policy: Should we wait for more data? *Energy Journal*, 23(2):97–119.

Webster, M., Jakobovits, L., and Norton, J. (2008). Learning about climate change and implications for near-term policy. *Climatic Change*, 89(1-2):67–85.

Weitzman, M. (2011). An economist grappling with catastrophic climate change. In *Tanner Lecture, Said Business School, Oxford*.

Weitzman, M. L. (1974). Prices vs. quantities. *Review of Economic Studies*, 41(4):477–91.

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

Weitzman, M. L. (2010). GHG targets as insurance against catastrophic climate damages. *National Bureau of Economic Research Working Paper Series*, No. 16136.

Yin, J., Schlesinger, M. E., and Stouffer, R. J. (2009). Model projections of rapid sea-level rise on the northeast coast of the united states. *Nature Geosci*, 2(4):262–266.