Climate Policy under Uncertainty

vorgelegt von Diplom Physiker

Matthias G.W. Schmidt

aus Heidelberg

Technische Universität Berlin

Fakultät VI - Planen Bauen Umwelt

vorgelegt zur Erlangung des akademischen Grades

Doktor der Wirtschaftswissenschaften - Dr. rer. oec.

genehmigte Dissertation

Promotionsausschuss:

Vorsitzende: Prof. Dr. Cordula Loidl-Reisch

Berichter: Prof. Dr. Ottmar Edenhofer

Berichter: Prof. Dr. Hermann Held

Tag der wissenschaftlichen Aussprache: 31.08.2011

Berlin 2011

D83

Contents

Abstract

The challenges posed by climate change are unprecedented in scale and scope. Climate change is global in its origins and impacts. It involves time horizons of hundreds of years and many generations. And, last but not least, it is surrounded by great uncertainty, which is the focus of this thesis. More specifically, this thesis intends to contribute to the identification of climate policies that do justice to the pervasiveness of uncertainty in climate change. In its core it contains four research articles.

The first article shows that the combination of uncertainty about climate damages with the fact that climate damages will be distributed heterogeneously across the population can be an argument for substantially stricter climate policy, i.e. stronger emissions reductions. The article also discusses how insurance and self-insurance can, at least theoretically, mitigate this result and thus permit weaker climate policy.

The second article highlights some major conceptual problems of cost-effectiveness analysis of climate policies for given climate targets. The problems occur once it is taken into account that uncertainty will be reduced in the future, which is an important aspect of climate change. In consequence, we propose an alternative decision criterion that avoids the problems by including a trade-off between the probability of violating the target and aggregate mitigation costs.

The third article investigates the circumstances under which learning about tipping elements in the climate system is an argument for stricter or weaker climate policy. It shows that learning is an argument for stricter policy if it is expected to happen in a narrow "anticipation window" in time, and that it can be neglected otherwise.

The fourth article reviews approaches to uncertainty in integrated assessment models of climate change with corresponding results. The complexity of the matter demands a variety of complementary approaches and a later synthesis of results. This article intends to summarize and structure this process and the respective literature.

The research articles are framed by an introduction to the field and general conclusions.

Acknowledgements

I would like to thank my family and friends as well as my colleagues at the Potsdam Institute for their support. In particular, I want to thank Hermann Held and Elmar Kriegler for creating and leading a wonderful collaboration in the Risk Group. Thanks to Alexander Lorenz for valuable discussions, and thank you to Ottmar Edenhofer and Hermann Held for supervising this thesis.

Chapter 1

Introduction

This chapter lays out the context of the thesis and specifies its objectives. Sections 1.1 and 1.2 give brief overviews of the science and economics of climate change, respectively. Section 1.3 introduces the main questions raised by uncertainty and how they have been approached. Section 1.4 then specifies the thesis objective and outline.

1.1 The Science of Climate Change

The basic cause-effect chain of anthropogenic climate change is straightforward. The burning of fossil fuels, land-use change, livestock production, and many other human activities produce greenhouse gases (GHGs), such as CO_2, CH_4, N_2O, and others. This increases the GHG concentration in the atmosphere. GHGs are essentially transparent for the incoming visible radiation from the sun but absorb and diffusely re-radiate the outgoing infrared radiation from the earth surface. Increased GHG concentrations thus lead to an imbalance between incoming and outgoing radiant energy. In consequence, earth surface temperature and the corresponding outgoing radiation increase until a new energy balance is reached. In 2004, for instance, global anthropogenic emissions of the GHGs included in the Kyoto protocol amounted to 49GtCO_2-eq (IPCC, 2007c) and were growing at roughly 3%/yr, mainly due to growth of emissions in China. The overall concentration of these GHGs had increased from 278ppm CO_2-eq at preindustrial times (around 1850) to 433ppm CO_2-eq (IPCC, 2007a), i.e. by roughly 50%. This had led to an energy imbalance of 1.6 W/m^2 and an increase of global mean temperature of about 0.7°C. Due to the inertia of the climate system and the warming of the oceans, in particular, committed warming was higher. Hare and Meinshausen (2006) predict a 1.2°C expected equilibrium warming if GHG concentrations were kept constant at 2004 levels.

Although the basic cause-effect chain of climate change is well understood, quantifying and predicting its dynamics is notoriously difficult and requires an understanding of all major components and feedbacks of the climate system. Major processes still to be understood include the cloud feedback, the ice sheet response to warming, and the carbon cycle feedback amongst others. Furthermore, substantial uncertainty is associated with the devel-

opment of the global economy and the resulting GHG emissions. As a result, average global warming in 2100, for instance, is highly uncertain. The IPCC AR4 specifies likely ranges based on the SRES scenarios for emissions in the absence of climate policy. For the most benign scenario, called B1, average global warming in 2100 is likely to be between 1.1 and 2.9°C. For the worst scenario, called A1F1, it is likely to be between 2.4 and 6.4°C (IPCC 2007a). The policy implications of this uncertainty are the subject matter of this thesis and will be introduced in greater detail in Section 1.3

An increase of global mean temperature by 2°C can already pose a major threat, because it implies considerably stronger local and seasonal changes in temperature in many places. It is also likely to alter rainfall patterns and to increase the frequency and intensity of extreme weather events, such as storms, droughts, floods, and heat waves. Furthermore, it might trigger irreversible tipping-element-like processes in the climate system. Melting of the Greenland Ice Sheet, for instance, lowers the altitude of the ice sheet surface and hence increases the surface temperature, which in turn reinforces the melting. This could eventually lead to a complete disintegration of the ice sheet with an associated sea-level rise of up to 7 meters. Other potential tipping-elements include the shutoff of the Atlantic thermohaline circulation (see also Chapter 4), the collapse of the Amazon rain forest, and the release of methane from melting permafrost (see Lenton et al., 2008, for a complete list).

Another layer of complexity is added when considering the actual damages inflicted on human societies as a consequence of climate impacts. Damages crucially depend on the vulnerability and adaptive capacity of societies, both of which are hard to estimate. The main damages include losses in food security and ecosystem services, increased water stress, diminishing biodiversity, the spread of infectious diseases, and direct losses of life due to extreme weather events (IPCC, 2007b). These direct impacts might lead to further indirect ones in the form of social conflicts and migration.

It is expected that damages will be distributed unequally across the globe, with poor countries experiencing greater impacts, higher vulnerability, and smaller adaptive capacity than rich countries. African countries, in particular, are expected to suffer severe damages, whereas the USA, for instance, might even experience some benefits from climate change, at least initially (IPCC, 2007b). At the same time, GHG emissions, as the root cause of climate damages, are and have been predominantly produced by rich countries. This implies a first ethical dimension to the climate problem (e.g. Edenhofer et al., 2010). Another ethical dimension originates from the fact that damages will predominantly be experienced by future generations, whereas emissions are produced by current ones.

1.2 The Economics of Climate Change

Science has shown that climate change is global, that it involves big uncertainties and long time horizons. Economics adds at least two qualitative aspects to this picture: Firstly, climate change comprises several market failures, the most important one being the externality of climate damages. A negligible share of the damages resulting from GHG emissions are

borne by the person producing them. Secondly, climate change involves difficult trade-offs between mitigation costs and avoided risks, between the interests of rich countries and poor countries, and between current and future generations. We briefly discuss these two aspects of climate change - externalities and trade-offs - in turn.

The classic approach to externalities is to put a price on the activity creating it, thus internalizing external costs and aligning private and social interests. In our context this amounts to putting a tax on GHG emissions or similarly demanding emitters to purchase tradable emissions permits whose overall supply is limited by the aggregate emission target. The latter is called a cap-and-trade system and essentially equivalent to a commensurate carbon tax unless there is asymmetric information between the regulator and the regulated (Weitzman, 1974, see end of Section 1.3). In the absence of additional market failures, both instruments are efficient, i.e. they achieve given emission reductions at minimum cost. However, climate policy involves other market failures, the arguably most important one being technology spillovers. Progress made by one firm in a technology partly spills over to other firms. This positive externality generally leads to underinvestment from a social point of view and justifies targeted technology subsidies in addition to a carbon price (Jaffe et al., 2005).

Reaching a global agreement that puts a price on GHG emissions has turned out to be difficult, as witnessed by the recent UNFCCC Conferences of the Parties (COP). There are at least three major reasons for this. Firstly, individual countries try to free-ride on the emissions reductions of others. Secondly, future generations suffering the bulk of climate damages are not present in today's negotiations. Thirdly, the main historical emitters including the USA and Europe are not the ones likely to suffer the greatest damages. Hence, the challenge is to agree on a fair burden sharing including damages and the costs of mitigation and adaptation, and at the same time to prevent free-riding.

Meanwhile, individual countries and regions have already set themselves emissions reduction targets. The European Union and California, for instance, aim for at least 80% reductions by 2050. Besides, COP16 could raise US$5 billion for the UN REDD (Reducing Emissions from Deforestation and Forest Degradation) program. However, it is doubtful that these efforts will effectively counteract climate change, and a global agreement including the main emitters, USA and China, seems indispensable.

In the following we neglect market failures and free-riders and turn to the second economic aspect of climate change mentioned at the beginning of this section: the trade-offs involved in first-best climate policy. Broadly speaking, global emissions reductions should be determined by a trade-off between avoided damages and mitigation costs. There are two mainstream approaches to do that: The first approach monetizes and aggregates climate damages globally and then performs a formal cost-benefit analysis (Cline, 1992, and Nordhaus, 1994, have pioneered here). The second one might be called a risk management approach. It clarifies and quantifies climate risks but avoids the monetization of climate damages. A somewhat informal comparison of mitigation costs and risk reduction then determines desirable policies (e.g. Schneider & Mastrandrea, 2005). The policies are often

formulated as GHG concentration or temperature stabilization targets. The popularity of the risk management approach mainly stems from the controversies surrounding the monetization of damages summarized below.

A straightforward way to monetize damages, or more generally risks, is by asking the affected what she would be willing to pay to get rid of the risk. At best this willingness-to-pay can be observed in markets. But how to aggregate these individual values over the entire population? Most global impact studies simply add them up. Yet this implies a risk to the life of a rich person is valued higher than the same risk to a poor person, for instance. Another question is how to aggregate these values over present and future generations. A sound way to aggregate is by assuming a social welfare function, which is discussed in Chapter 2. This is a very strong normative assumption, though, and results will be very sensitive to it. Especially notorious parameters in this context are the pure rate of time preference at which future welfare is discounted, as well as society's degree of inequality aversion (see Dasgupta, 2008, for an overview). A further concern with the cost-benefit approach is that some damages and risks simply cannot be quantified (e.g. Stern, 2008). However, the author is skeptical about this concern. As long as quantification does not trick oneself into a false sense of certainty, it should be superior to qualitative or intuitive reasoning.

The costs of mitigation are somewhat less controversial than climate damages. GHG emissions have several drivers that can be targeted by mitigation. We decompose emissions into

$$\text{Emissions} = \text{Pop} \times \text{GDP}/\text{Pop} \times \text{E}/\text{GDP} \times \text{Emissions}/\text{E},$$

The first factor, population, is controlled by factors outside the realm of climate policy. However, Kelly and Kolstad (2001) show that the severity of climate change is very sensitive to population growth. Directly lowering GDP per capita in the second factor is a highly controversial and supposedly inefficient way to reduce emissions. The third factor, energy efficiency, is widely believed to offer some no-regret options that would be desirable even without climate change (1.4 GtC annually in a recent study by McKinsey, 2007). The last term, emissions intensity of energy, encompasses the bulk of mitigation options such as renewable energy, carbon capture and sequestration, and avoided deforestation.

The primary tool to quantify mitigation costs are integrated assessment models (IAMs) of climate change. This thesis makes extensive use of these models. They normally consist of a dynamic model of the economy coupled to a simple climate model, and they vary in macroeconomic, energy economic, and climate dynamic complexity. The pioneering work in this field is due to Cline (1992) and Nordhaus (1994). The latter consistently argues for a "policy ramp", i.e. a slow increase of mitigation effort, and an eventual stabilization of GHG concentrations at around 650 ppm CO_2-eq. However, this result stems from cost-benefit analysis and is very sensitive to the aforementioned controversial choices of normative and climate damage parameters. A growing number of studies just calculate mitigation costs for various stabilization targets in a cost-effectiveness analysis. The IPPC AR4 reports costs

of at most 3% GDP by 2030 and 5.5% by 2050 for the 2°C target (IPCC, 2007c). The EU project ADAM reports at most 2% discounted overall costs for the same target and various IAMs (Edenhofer et al., 2010). A desirable target could then be determined based on these costs in the above-mentioned risk management approach. Stern (2008), for example, argues for at most 550 ppm CO_2-eq in this vein.

Mitigation in a strict sense only targets the first link in the cause-effect chain of climate change (see Section 1.1): emissions. In recent years, increasing attention has been paid to geoengineering, which applies further down the cause-effect chain. One can distinguish between carbon management and radiation management. Carbon management aims at removing CO_2 from of the atmosphere or the ocean. Technologies and mechanisms to do so are still in early stages of development, though, and it is not clear wether they will eventually become cost-effective. Radiation management aims at directly reducing the radiative forcing instead of reducing GHG concentrations. A prominent option is the injection of reflective sulfur particles in the stratosphere (Crutzen, 2006). Radiation management is likely to be cheap but associated with substantial risks of its own. It is therefore mostly seen as a measure of last resort (Victor et al., 2009).

1.3 Uncertainty and Climate Policy

The pervasiveness of uncertainty is one of the main challenges in climate policy. It has an influence on both optimal policy stringency and optimal policy instruments. We focus on the former but comment on the latter at the end of this section. Uncertainty affects optimal stringency in two major ways:

Firstly, most people simply dislike uncertainty. This is called risk aversion. The question is then whether climate policy decreases or increases uncertainty. On the one hand, it reduces uncertainty, because the closer the climate is to the current one, the better it is known. On the other hand, it introduces uncertainty about mitigation costs. Which uncertainty dominates is an empirical question.

Secondly, uncertainty will be reduced in the future due to scientific progress. We will call this "learning". It is sometimes used as an argument for deferring mitigation action. This can be illustrated by drawing an admittedly odd but nonetheless helpful analogy between mitigation and buying a used car. We assume, the value of the car is uncertain, but the price is low enough to justify its purchase. Now, if a mechanically more savvy friend, who plays the role of science here, could tell us the true value of the car tomorrow, we should certainly consider deferring the purchase until tomorrow and only buy it if it turns out to be in good shape, even if the price is acceptable as of today. In analogy, it should be considered to defer mitigation effort until more is known about its benefits. However, this argument neglects the fact that the car might already be sold tomorrow, and we might therefore buy it right now anyway. In the climate context, deferring mitigation action commits the planet to ever more warming and impacts, mainly because GHG emissions and the tipping-elements mentioned in Section 1.1 are essentially irreversible. Again, whether the irreversibility of

investments in mitigation or the irreversibility of climate processes dominates is an empirical question.

The preceding questions and arguments have to be evaluated by quantifying the uncertainty, applying a suitable decision criterion, and thus obtaining a desirable or at best optimal policy. The most common way to quantify uncertainty is in terms of probabilities. Probabilities are called "objective" if they have a frequentist interpretation: an experiment is repeated many times, and the frequencies of different outcomes define their probabilities. A weaker notion of probability is called "subjective" and based on the level of confidence a person has in different outcomes. The most common probabilistic decision criterion is expected utility maximization, or (marginal) cost-benefit analysis. Decision makers maximize the expected value of their utility. As a descriptive criterion, this has been contested by a growing body of experiments, but for normative purposes, it is still firmly grounded on cogent axioms of rationality (e.g. Gilboa, 2009). The risk management approach mentioned in Section 1.2 often makes use of probabilities as well.

However, the non-repeatability of climate change prevents the derivation of objective probability distributions, and there is no general agreement on subjective ones, either. The lack of unique probability distributions is often called "deep uncertainty", or "Knightian uncertainty" after Frank H. Knight (1921). It demands alternative decision approaches. Several of them have been applied to the climate problem, including Dempster-Shafer theory (Luo & Caselton, 1997), ambiguity (Lange & Treich, 2008), and robust decision making (Lempert et al., 2000). The latter is a non-probabilistic variety of the risk management approach mentioned above. Approaches to deep uncertainty are generally involved and their normative foundation is still controversial to some extent. This explains why most integrated assessments of climate change including this thesis presume a unique probability distribution.

In the following, we will shortly summarize the main conclusion from the literature. A more detailed review including an introduction to different probabilistic approaches is given in Chapter 5. We organize the summary around the two key questions mentioned at the beginning of this section: (i) How does uncertainty change the stringency of optimal climate policy in terms of emissions reductions as compared to a world where all uncertain parameters are fixed at their expected value? (ii) How does future learning about uncertainty change the stringency of optimal near-term policy as compared to a world with uncertainty but without learning? It is often helpful to treat these questions separately, because they demand different simplifications.

(i) A lot of the recent discussion has evolved around the upper tail of the distribution of climate damages, or the fact that currently we cannot exclude truly catastrophic consequences of climate change. This point has been made most forcefully by Weitzman's dismal theorem (Weitzman, 2009), which shows that the said upper tail might actually lead to an unbounded utility and thus render a cost-benefit analysis impossible. It has been argued by Nordhaus (2009) and others, though, that the preconditions of the dismal theorem are too restrictive to hold for the climate problem. The presumed exponential dependence

of climate damages on temperature is particularly controversial. Further research will be needed to settle this discussion. More applied studies using IAMs neglect potential tails in the probability distributions. They generally find that uncertainty argues for a somewhat stricter climate policy. The damage uncertainty dominates the mitigation cost uncertainty. The magnitude of the effect ranges from small (Peck & Teisberg, 1993; Webster et al., 2008) up to 30% stronger emission reductions (Pizer, 1999) depending on the model used and the degree of uncertainty considered. Chapter 2 shows that uncertainty can have a substantial effect on optimal policy, if the heterogeneity of the distribution of climate damages across the global population is taken into account.

(ii) The effect of future learning on optimal near-term climate policy is generally found to be small in cost-benefit analysis (Peck & Teisberg, 1993; Ulph & Ulph, 1997; Webster, 2002; Webster et al., 2008, O'Neill & Sanderson, 2008). The irreversibility of emissions and investments cancel each other out. This result changes, though, if a highly non-linear climate tipping element is included. This was shown first by Keller et al. (2004) and is extended in chapter 4. Future learning generally leads to substantially stricter near-term policy in cost-effectiveness analysis (Webster et al., 2008; Bosetti et al., 2009). However, Schmidt et al. (2011, Chapter 3) argue that this result is due to a controversial interpretation of climate targets under uncertainty as strict targets that have to be met with certainty.

Up to now, we have not been concerned with how to achieve a given policy by actual policy instruments, which is also not the focus of this thesis. Uncertainty influences the choice of an instrument to internalize the climate externality described in Section 1.2 in mainly three ways: First, Weitzman (1974) shows that the equivalence of price and quantity instruments breaks down under asymmetric information about abatement costs. More specifically, he assumes that private emitters know their abatement costs while the regulator doesn't. A price instrument, such as a carbon tax, then implies uncertainty about emissions and damages, whereas a quantity instrument, such as a cap-and-trade system, implies uncertainty about costs. Hence, the former is preferred if damages as a function of emissions are less convex than abatement costs. This is arguably the case in climate change (Pizer, 1999), at least as long as climate tipping elements are neglected. Second, Stavins (1996) points out that additional uncertainty about damages can reverse this result if damages are positively correlated with abatement costs. This correlation is likely to be weak, though, in climate change. Third, Baldursson & von der Fehr (2004) highlight that a price instrument might be superior to a quantity instrument if firms are overly risk-averse. The price volatility associated with a quantity regulation then implies an inefficiently low level of trade in emissions permits and investments in R&D. In summary, uncertainty is likely to favor a price instrument over a quantity instrument for climate change. However, arguments not linked to uncertainty might still favor and explain the current focus on quantity regulation (see Hepburn, 2006, for an overview).

1.4 Thesis Outline

The preceding discussion has highlighted the pervasiveness of uncertainty in climate change and the challenges this poses to economics and climate policy. This thesis intends to contribute to the identification of policies that rise to this challenge. It's core comprises four articles contained in chapters 2 to 5. The articles have already been put in the context of the literature in the previous section and will be outlined in the following. The author of this thesis will also state his contributions to the individual articles.

Chapter 2: This article shows that substantially stricter climate policy can be desirable if both the uncertainty of climate damages and the fact that climate damages are distributed heterogeneously across the population are taken into account. More specifically, the joint effect of uncertainty and heterogeneity on climate policy can be sizable even if the separate effects are negligible. The reason is that the same risk borne by fewer people implies a larger risk premium. We discuss how insurance markets could eliminate this effect by spreading the risk over the entire population. Furthermore, we show how self-insurance can still mitigate the effect if insurance markets are not available. Individuals that are strongly affected by climate damages increase their savings today to partly compensate the damages in the future. The different effects are first presented in a simple analytical model that allows closed-form solutions. Numerical results are then provided for the IAM DICE.

This article has been submitted to Environmental & Resource Economics as "Schmidt, M.G.W., H. Held, E. Kriegler, A. Lorenz. *Stabilization Targets under Uncertain and Heterogeneous Climate Damages.*" M.G.W. Schmidt conceived the idea for this research, performed the analysis and wrote the article. A. Lorenz provided model input. All four authors contributed by extensive discussions.

Chapter 3: Climate Targets are becoming ever more influential as witnessed by the recent adoption of the 2°C target by COP15. 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 once future learning about uncertainty is taken into account, and learning is an essential aspect 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, per-

formed the analysis and wrote the article. The co-authors, and A. Lorenz in particular, contributed by extensive discussions helping to structure the argument.

Chapter 4: As mentioned in Section 1.1, one of the crucial uncertainties of climate change are tipping-elements in the climate system. This article sheds some light on the implications of uncertainty and future learning about these tipping-elements for optimal near-term climate policy. The main finding, obtained from the IAM called MIND, is that optimal near-term policy should be substantially stricter if learning about the severity of the tipping-point is expected to happen in a specific, narrow time window. Stricter policy then serves to keep the option open to avoid the tipping-point in case it is learned to be severe. Future learning has no relevant effect on near-term policy other-wise. However, learning about the severity of the tipping-element and adjusting post-learning decisions accordingly is valuable in either case. The article furthermore provides some novel concepts for the analysis and interpretation of results under anticipated future learning.

This article has been accepted for publication in Environmental Modeling 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 numerical analysis was performed by A. Lorenz, who also wrote the larger fraction of the article. M.G.W. Schmidt made substantial contributions in conceptualizing the results and writing the article.

Chapter 5: Investigating uncertainty in IAMs, the primary tool of climate policy analysis, is both conceptually and numerically demanding. Various complementary approaches and simplifications are required to grasp the full implications of uncertainty. This article provides both an overview of approaches and a review and synthesis of results in the literature. This blend allows us to structure the literature and to identify the main drivers of the results, such as the quality of representation of uncertainty and learning. Included is an accessible introduction to a novel approach based on real options analysis, which has recently been proposed by one of the co-authors of this article. We also identify a number of future research needs.

An earlier version of this article has been published as "Golub, A., D. Narita, M.G.W. Schmidt, 2011. *Uncertainty in Integrated Assessment Models of Climate Change: Alternative Analytical Approaches*. FEEM Working Paper No. 2.2011". The three authors contributed equally to its conception and writing.

Finally, Chapter 6 draws some general conclusions from this thesis and indicates future research needs in the context of climate policy and uncertainty.

1.5 References

Baldursson, F. M., & N.-H. M. von der Fehr 2004. Price volatility and risk exposure: on market-based environmental policy instruments. Journal of Environmental Economics and Management, 48(1):682–704.

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

Cline, W. R. 1992. The Economics of Global Warming. Institute for International Economics,U.S., first edition, first printing edition.

Crutzen, P. J. 2006. Albedo Enhancement by Stratospheric Sulfur Injections: A Contribution to Resolve a Policy Dilemma? Climatic Change, 77(3-4):211–220.

Dasgupta, P. 2008. Discounting climate change. Journal of Risk and Uncertainty, 37(2-3):141–169.

Edenhofer, O., J. Wallacher, M. Reder, & H. Lotze-Campen (eds) 2010. Global aber Gerecht. Klimawandel bekaempfen, Entwicklung ermoeglichen. C.H. Beck, Muenchen.

Enkvist, P.-E., T. Naucler, & J. Rosander 2007. A Cost Curve for Greenhouse Gas Reduction. The McKinsey Quarterly, 1.

Gilboa, I. 2009. Theory of Decision under Uncertainty. Cambridge University Press, 1 edition.

Hare, B., & M. Meinshausen 2006. How Much Warming are We Committed to and How Much can be Avoided? Climatic Change, 75(1-2):111–149.

Hepburn, C. 2006. Regulation by Prices, Quantities, or Both: A Review of Instrument Choice. Oxford Review of Economic Policy, 22(2):226 –247.

IPCC 2007a. Fourth Assessment Report: Working Group I Report "The Physical Science Basis". Cambridge University Press, Cambridge, United Kingdom and New York, USA.

IPCC 2007b. Fourth Assessment Report: Working Group II Report "Impacts, Adaptation and Vulnerability". Cambridge University Press.

IPCC 2007c. Fourth Assessment Report: Working Group III Report "Mitigation of Climate Change". Cambridge University Press, Cambridge, United Kingdom and New York, USA.

Jaffe, A. B., R. G. Newell, & R. N. Stavins 2005. A tale of two market failures: Technology and environmental policy. Ecological Economics, 54(2-3):164–174.

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

Kelly, David L., & Charles D. Kolstad 2001. Malthus and Climate Change: Betting on a Stable Population. Journal of Environmental Economics and Management, 41(2):135–161.

Knight, F. H. 1921. Risk, Uncertainty, and Profit. Hart, Schaffner, and Marx Prize Essays. Houghton Mifflin, Boston and New York.

Lange, A., & N. Treich 2008. Uncertainty, learning and ambiguity in economic models on climate policy: some classical results and new directions. Climatic Change, 89(1-2):7–21.

Lempert, R.J., M.E. Schlesinger, S.C. Bankes, & N.G. Andronova 2000. The impacts of climate variability on near-term policy choices and the value of information. Climatic Change, 45(1):129–161.

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

Luo, W.B., & B. Caselton 1997. Using Dempster-Shafer theory to represent climate change uncertainties. Journal of Environmental Management, 49(1):73–93.

Nordhaus, W.D. 1994. Managing the Global Commons: The Economics of Climate Change. The MIT Press, 1st edition.

Nordhaus, W. D. 2009. An Analysis of the Dismal Theorem. Cowles Foundation Discussion Paper, 1686.

O'Neill, B. C., & W. Sanderson 2008. Population, uncertainty, and learning in climate change decision analysis. Climatic Change, 89(1-2):87–123.

Peck, S. C., & T. J. Teisberg 1993. Global Warming, Uncertainties and the Value of Information - An Analysis Using CETA. Resource and Energy Economics, 15(1):71–97.

Pizer, W. A. 1999. The optimal choice of climate change policy in the presence of uncertainty. Resource and Energy Economics, 21(3-4):255–287.

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.

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

Stavins, R. N. 1996. Correlated Uncertainty and Policy Instrument Choice. Journal of Environmental Economics and Management, 30(2):218–232.

Stern, N. 2008. The Economics of Climate Change. American Economic Review, 98(2):1–37.

Ulph, A., & D. Ulph 1997. Global Warming, Irreversibility and Learning. The Economic Journal, 107(442):636–650.

Victor, D.G., M.G. Morgan, F. Apt, J. Steinbruner, & K. Ricke 2009. The Geoengineering Option: A Last Resort Against Global Warming? 88(2):64.

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

Weitzman, M.L. 2009. On Modeling and Interpreting the Economics of Catastrophic Climate Change. The Review of Economics and Statistics, 91(1):1–19.

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

Chapter 2

Stabilization Targets under Uncertain and Heterogeneous Climate Damages[1]

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

[1]This chapter has been submitted to Environmental & Resource Economics as "Schmidt, M.G.W., H. Held, E. Kriegler, A. Lorenz. *Stabilization Targets under Uncertain and Heterogeneous Climate Damages.*

Stabilization Targets under Uncertain and Heterogeneous Climate Damages

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

Abstract

We highlight that uncertainty about climate damages and the fact that damages will be distributed heterogeneously across the global population can jointly be an argument for substantially stricter climate stabilization targets even if uncertainty and heterogeneity in isolation are not. The reason is that a given climate risk borne by fewer people implies greater welfare losses. However, these losses turn out to be significant only if society is both risk and inequality averse. We discuss how insurance and self-insurance of climate risk could theoretically mitigate this joint effect of uncertainty and heterogeneity and thus admit weaker stabilization targets. Insurance provides more efficient risk sharing and self-insurance allows strongly impacted individuals to compensate damages by increasing savings. We first use a simple analytical model to introduce the different concepts and then provide more realistic results from the integrated assessment model DICE.

Keywords: climate change, stabilization target, uncertainty, heterogeneity, damages, insurance

1 Introduction

Climate change is surrounded by great uncertainty. However, a number of integrated assessment studies have found that uncertainty has only a minor effect on first-best climate policy and emissions (Peck & Teisberg, 1993; Nordhaus, 1994; Nordhaus & Popp 1997; Ulph & Ulph, 1997; Webster, 2002). These studies are based on the assumption of a representative agent. The real society with heterogeneous preferences, income levels, and climate damages is replaced by a fictitious homogeneous one that is supposed to lead to the same equilibrium prices and savings. A representative agent is known to exist for complete market economies (Constantinides, 1982), which in this context would include insurance markets for climate damages. However, markets are far from complete, and risks are not shared efficiently. As an example, currently only about 20% of catastrophic damages are insured (Mills, 2005).

Introducing explicit damage heterogeneity, or any heterogeneity for that matter, immediately raises questions of equity that where conveniently omitted in representative agent models. How should impacts imposed on different people be valued and aggregated? Global impact studies are still rare and mostly just add up the willingness-to-pay for avoiding damages of all individuals (Cline, 1992; Nordhaus 1994; Fankhauser 1994; Nordhaus 2006; Hope 2006), thus valuing impacts on poor people lower than the same impacts on rich people (see Fankhauser et al., 1997). An

*M.G.W. Schmidt (corresponding author), E. Kriegler, A. Lorenz
Potsdam Institute for Climate Impact Research, 14412 Potsdam, Germany
E-mail: schmidt@pik-potsdam.de. Tel.: +49-331-2882566
†H. Held
University of Hamburg - Klima Campus, Bundesstr. 55, 20146 Hamburg, Germany
and Potsdam Institute for Climate Impact Research, 14412 Potsdam, Germany

exception is Tol (2002). We will, in contrast, reconstruct the heterogeneous damages from the aggregate estimates and then, following Fankhauser et al. (1997), explicitly use a social welfare function to aggregate to the overall population. In the welfare function we separate risk aversion from inequality aversion in order to clarify the interaction between the two.

There are several papers that analyze regional damage heterogeneity without uncertainty (Nordhaus & Yang, 1996; Azar 1999; Fankhauser & Tol, 2005; Anthoff et al., 2009; Anthoff & Tol, 2010). This paper is closest to Tol (2003) and Anthoff & Tol (2009), who take uncertainty into account. Using the integrated assessment model FUND they show that damage- and income heterogeneity in combination with uncertainty can have a big effect on the benefits of emission reductions and even lead to a break-down of cost-benefit analysis if the uncertainty is fat-tailed in some regions. Using a simple analytical model, our paper intends to clarify and separate the effects of heterogeneity and uncertainty. It also shows how insurance markets and self-insurance can mitigate them. Numerical results for the benefits of various concentration targets are then obtained with the integrated assessment model DICE.

More specifically, the five main points of this article are the following: (i) Uncertainty and damage heterogeneity can jointly have a strong effect on optimal climate policy even if their separate effects are negligible. The reason is that the same risk borne by fewer people implies greater welfare losses. The fact that uncertainty has only a small effect in other studies is hence at least partly due to their assumption of a representative agent and, more specifically, the implicit assumption of efficient risk sharing. (ii) Under constant relative risk aversion and inequality aversion, income inequality favors stricter climate policy only if people with low income either suffer higher relative damages or bear lower relative abatement costs than people with high income. (iii) The introduction of complete insurance markets essentially lowers the aggregate risk premium associated with heterogeneous damages to the one for homogeneous damages of the same amount. Complete insurance would therefore allow a significant relaxation of climate policy. (iv) Even in the absence of insurance markets, individuals can still mitigate the effect of damage heterogeneity substantially by self-insuring, i.e. increasing savings. This is particularly effective under lax climate policy, because it allows to shift consumption from the short term, where abatement costs are low, to the long term, where damages are high. (v) For all results we shortly discuss their dependence on the available information about aggregate climate damages and the distribution of damages across the population. As known in the literature, better information decreases the effectiveness of insurance markets in the absence of market failures but increases the effectiveness of self-insurance.

The article is structured as follows. In Section 2 we introduce the different concepts in an analytical model, where we can derive closed-form solutions. After a short introduction of the model assumptions, we discuss three settings: In Subsection 2.1 neither insurance with others nor self-insurance is possible. In Subsection 2.2, a perfect insurance market is available. In Subsection 2.3 self-insurance is possible whereas insurance with others is not. Section 3 then shows numerical results from the integrated assessment model DICE. Parallel to Section 2, Subsections 3.1, 3.2, and 3.3 discuss the three different settings. Finally, Section 4 concludes.

2 Analytical Model

In this section, we use a simple analytical model and convenient functional forms to define and discuss the effects of uncertainty and damage heterogeneity on welfare.

We make the following assumptions: (i) All agents have a constant absolute risk aversion utility

function, $u(c) = -e^{-A\,c}/A$, with the same degree of absolute risk aversion A. (ii) Aggregate, additive climate damages are normally distributed: $D \sim \mathcal{N}(\mu, \sigma)$. (iii) The heterogeneity of damages can be described by only two cohorts: one cohort is affected by climate damages, the other one is not. The affected cohort constitutes a share k of the population. Thus, if average per capita damages equal D for the overall population, per capita damages are $D^{(1)} = D/k$ for the affected and $D^{(2)} = 0$ for the unaffected, where the superscripts indicate the cohort. The homogeneous case is obtained for $k = 1$.

All three assumptions will be replaced by more realistic ones in the numerical model in Section 3. Furthermore, we assume that the climate risk is the only risk in the economy, i.e. there is no systemic macroeconomic or idiosyncratic income risk. For most of this section we neglect inequality in gross income before damages in order to isolate the effect of damage heterogeneity, but we consider income inequality at the end of Subsection 2.1.

2.1. No Insurance

In this subsection, we assume the cohort that is exposed to the climate risk cannot insure with the rest of the population. This is not a completely unrealistic assumption. As mentioned in the introduction, currently only about 20% of catastrophic risks are insured (Mills, 2005), and big part of climate impacts will be in the form of catastrophes such as floods, heat waves, storms and so on. Gollier (2005) gives an overview of possible reasons for the difficulties of insuring catastrophic risks.

The certainty equivalent (CE) consumption of an affected individual, $\bar{c}^{(1)}$, where the over-bar refers to the CE and the superscript indicates the cohort, is implicitly defined over

$$E\left[u\left(c - D/k\right)\right] \quad = \quad u\left(\bar{c}^{(1)}\right), \tag{1}$$

where c is gross consumption. For simplicity we do not explicitly specify the dependence of $\bar{c}^{(1)}$ on u, c, D, and k. Under the functional assumption at the beginning of this section, we get

$$\bar{c}^{(1)} \quad = \quad c - \frac{\mu}{k} - \frac{A}{2}\frac{\sigma^2}{k^2}. \tag{2}$$

The risk premium is then given by $\pi^{(1)} = E[c^{(1)} = c - D/k] - \bar{c}^{(1)} = (A/2)\,\sigma^2/k^2$.

For the aggregation to the overall population, we use a social welfare function. In order to separate the effect of risk / risk aversion from the effect of inequality / inequality aversion on welfare, we assume it to be of the following form, which will be explained below,

$$\begin{aligned} W(c^{(1)}, c^{(2)}, k) \quad &= \quad k\,v\left(u^{-1}\left(E\left[u(c^{(1)})\right]\right)\right) + (1-k)\,v\left(u^{-1}\left(E\left[u(c^{(2)})\right]\right)\right) \\ &= \quad k\,v(\bar{c}^{(1)}) + (1-k)\,v(c), \end{aligned} \tag{3}$$

Here, $u^{-1}(\cdot)$ is the inverse of the utility function, and $v(\cdot)$ is an increasing function expressing inequality aversion. In the second step we used the fact that the second cohort does not suffer damages. Thus, we take the function v of the certainty equivalent consumption levels of the individuals and then sum over the individuals. Society is risk averse, if it is worse off with uncertain consumption than with consumption fixed at its expected values, $W(c^{(1)}, c^{(2)}, k) < W(E[c^{(1)}], E[c^{(2)}], k)$. Since v is an increasing function, this translates to $\bar{c}^{(i)} < E[c^{(i)}]$, which in turn implies strict concavity of u, $E[u(c^{(i)})] < u(E[c^{(i)}])$. Society is risk-neutral, if the inequalities are replaced by

equalities so that u has to be linear. Thus, risk aversion is determined by the curvature of u. Society is inequality averse, if it is worse off with a heterogeneous distribution of certain consumption over individuals than with a homogeneous distribution, where all individuals enjoy average consumption, $W(\bar{c}^{(1)}, \bar{c}^{(2)}, k) < W(k\bar{c}^{(1)} + (1-k)\bar{c}^{(2)}, k\bar{c}^{(1)} + (1-k)\bar{c}^{(2)}, k)$. This implies strict concavity of v, $kv(\bar{c}^{(1)}) + (1-k)v(\bar{c}^{(2)}) < v(k\bar{c}^{(1)} + (1-k)\bar{c}^{(2)})$. Society is inequality-neutral, if the inequalities are replaced by equalities so that v has to be linear. Thus inequality aversion is determined by the curvature of v. This way of separating inequality aversion from risk aversion is analogous to the way Kreps & Porteus (1978) separate the elasticity of inter-temporal substitution from risk aversion.

As customary, we define the certainty and equity equivalent (C&EQE) consumption level as the certain and homogeneous (across the population) consumption level that gives the same welfare as an uncertain heterogeneous one (e.g. Anthoff & Tol, 2009). We denote it by $\hat{\bar{c}}$, where the bar still refers to the CE and the hat refers to the EQE. More formally, we define

$$W(c^{(1)}, c^{(2)}, k) \quad = \quad v\left(\hat{\bar{c}}(u, v)\right) \tag{4}$$

where we omit the dependence of $\hat{\bar{c}}$ on $c^{(1)}$, $c^{(2)}$, and k. We consider four special cases:

(i) Society is both risk- and inequality averse. More specifically, we assume $v \equiv u$ and get the utilitarian welfare function $W(c^{(1)}, c^{(2)}, k) = k\,E\left[u(c^{(1)})\right] + (1-k)E\left[u(c^{(2)})\right]$. Somewhat sloppily we will denote $\hat{\bar{c}}(u, u)$ shortly by $\hat{\bar{c}}$. Under the functional assumptions of this section we get

$$\hat{\bar{c}} \quad = \quad c - \ln\left(1 - k\left(1 - e^{A\left(\mu/k + (A/2)\sigma^2/k^2.\right)}\right)\right)/A. \tag{5}$$

(ii) Society is only risk averse. For a linear $v(c) = c$, we simply add the certainty equivalents of all individuals, $W(c^{(1)}, c^{(2)}, k) = k\,\bar{c}^{(1)} + (1-k)\bar{c}^{(2)}$. We call the resulting consumption the CE consumption of the population and denote it by $\bar{c} = \hat{\bar{c}}\,(u, v(c) = c)$. Under the functional assumptions of this section we get

$$\bar{c} \quad = \quad k\bar{c}^{(1)} + (1-k)c = c - \mu - \frac{A}{2}\frac{\sigma^2}{k}. \tag{6}$$

Holding average damages D fixed: The smaller k, the greater is the risk of the affected individuals, namely D/k. This leads to an increase proportional to $1/k^2$ of the risk premium of the affected (Eq. 2) and hence an increase proportional to $1/k$ of the risk premium of the overall population (Eq. 6). Hence, the risk premium increases five times, for instance, if only 20% of the population are affected by climate damages. It is straightforward to verify that $\bar{c} \geq \hat{\bar{c}}$, i.e. inequality aversion decreases C&EQE consumption.

(iii) Society is only inequality averse. For a linear $u(c) = c$ we get $W(c^{(1)}, c^{(2)}, k) = k\,v(E[c^{(1)}]) + (1-k)\,v(E[c^{(2)}])$. We call the resulting C&EQE consumption the EQE consumption of the population and denote it by $\hat{c} = \hat{\bar{c}}\,(u(c) = c, v)$. Under the functional assumptions of this section and assuming $v(c) = -e^{-Ac}/A$ we get

$$\hat{c} \quad = \quad c - \ln\left(1 - k\left(1 - e^{A\mu/k}\right)\right)/A. \tag{7}$$

(iv) Society is neither risk nor inequality averse. For both linear $u(c) = c$ and $v(c) = c$ we get $W(c^{(1)}, c^{(2)}, k) = k\,E[c^{(1)}] + (1-k)\,E[c^{(2)}]$ and welfare is given by expected average consumption. For the example of this section, we have $W(c^{(1)}, c^{(2)}, k) = c - \mu$.

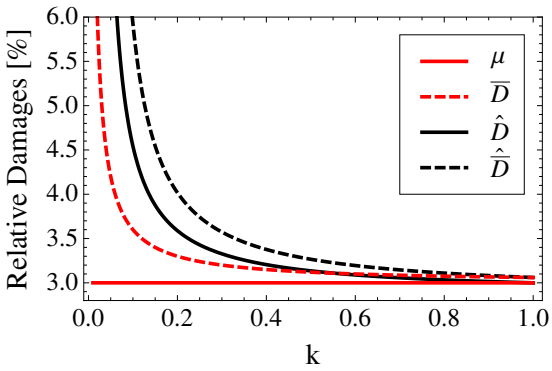

Figure 1: Expected (μ), CE ($\bar{D}$), EQE ($\hat{D}$), and C&EQE ($\hat{\bar{D}}$) damages in relative terms of per capita consumption of $c = \$7000/yr$ and as a function of k. The parameter values are $A = 3/7$, $\mu/c = 3\%$, $\sigma/c = 2\%$. Color code: Red lines refer to results without inequality aversion, black lines to results for a utilitarian. Solid lines refer to results without risk aversion, dashed lines to results with risk aversion.

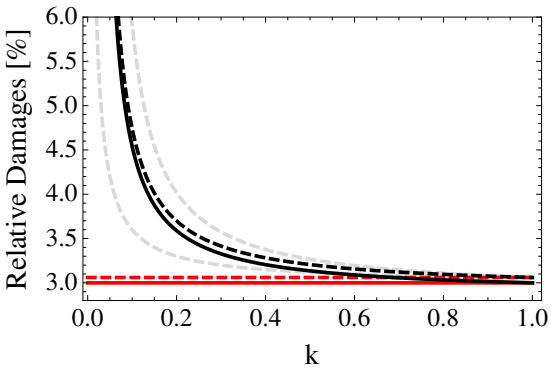

Figure 2: The same as in Fig. 1, in the same color code, but for the market solution. Damages without insurance are shown in light gray.

Fig. 1 shows exemplary results for expected, CE, EQE, and C&EQE damages, which are defined as the difference in the corresponding values for consumption with and without damages, i.e. $\hat{\bar{D}} = \hat{\bar{c}}_{D=0} - \hat{\bar{c}}$, for instance. It shows that uncertainty has a substantial effect on damages both with and without inequality aversion if k is small.

What happens if gross consumption, i.e. consumption before damages, differs between the cohorts? There are two effects (i) If absolute risk aversion A depends on the consumption level and more specifically decreases in consumption, then the risk premium will increase if the affected cohort is poorer than average. Under the assumption of constant absolute risk aversion in this section, though, this effect is absent. It will be present in the numerical model in Section 3. (ii) Gross consumption inequality has an effect on net consumption inequality and hence EQE consumption. Net inequality is decreased by gross inequality compared to the case of equal gross consumption, if the affected are richer than the non-affected by an amount smaller than twice $\bar{D}^{(1)}$. The initial wealth then partly compensates for damages. Inequality is increased otherwise. Smaller net consumption inequality leads to an increase in EQE consumption, or equivalently a decrease in EQE damages.

2.2. Perfect Insurance Market

In the last section, the affected individuals didn't have the possibility to insure with the rest of the population. Heterogeneity then leads to a substantial increase in C&EQE damages. Now we investigate to what extent a complete contingent claims, or insurance, market can mitigate this result. Since we assume no other risks in the economy, the benefits from such a market are due to risk sharing not diversification.

For each state of the world, characterized by average damages D, we introduce a tradable contingent claim that pays off average damages in the corresponding state of the world. We denote the prices of these claims by p_D, and the amounts of claims purchased by the affected and unaffected

by $x_D^{(1)}$ and $x_D^{(2)}$, respectively. The equilibrium conditions for the affected and unaffected are

$$\begin{aligned}
&\max_{x_{1,D}} \left\{ E\left[u\left(c - D/k + x_D^{(1)} D - \int_{-\infty}^{\infty} x_{D'}^{(1)} p_{D'} dD' \right) \right] \right\}, \\
&\max_{x_{2,D}} \left\{ E\left[u\left(c + x_D^{(2)} D - \int_{-\infty}^{\infty} x_{D'}^{(2)} p_{D'} dD' \right) \right] \right\}, \\
&\text{s.t. } \forall D: \; k\, x_D^{(1)} + (1-k) x_D^{(2)} = 0.
\end{aligned} \tag{8}$$

The integrals $\int_{-\infty}^{\infty} x_{D'}^{(i)} p_{D'} dD'$ equal the overall amount spent on the contingent claims portfolio, and the $x_D^{(i)} D$ equal the random payoffs of the portfolio. The last equality in Eq. (8) is the market clearing condition, which has to hold in every state of the world.

In the following we verify that under the assumptions of this section, individuals purchasing the same amount of per capita damages in all states of the world, i.e. $x_D^{(i)} = x^{(i)}$, $i = 1, 2$, is an equilibrium. Hence, in our simple setting, it is not necessary to have separate contingent claims for all states of the world to obtain the complete market equilibrium, but it is sufficient to have a single claim that pays off average per capita damages however high they turn out to be. This is a consequence of the linearity of individual damages in average damages in our simple formulation of heterogeneity. It won't hold for the formulation of heterogeneity used in the numerical model in Section 3.2. Substituting $x_D^{(i)} = x^{(i)}$, $i = 1, 2$ into Eq. (8) and denoting $p = \int_{-\infty}^{\infty} p_{D'} dD'$, we get

$$\begin{aligned}
&\max_{x_1} \left\{ E\left[u\left(c - D/k + x^{(1)}(D - p) \right) \right] \right\}, \\
&\max_{x_2} \left\{ E\left[u\left(c + x^{(2)}(D - p) \right) \right] \right\}, \\
&\text{s.t. } k\, x^{(1)} + (1-k) x^{(2)} = 0.
\end{aligned} \tag{9}$$

These conditions are solved under the functional assumptions of this section by

$$p = \mu + A\,\sigma^2, \tag{10}$$

$$x^{(2)} = -1 = -\frac{k}{1-k} x^{(1)}. \tag{11}$$

In the equilibrium described by Eqs. (10) and (11), every individual suffers per capita damages, the risk is equally distributed between all individuals. This result is due to the assumption of constant absolute risk aversion. For decreasing absolute risk aversion, the affected would carry a smaller risk in equilibrium because the insurance premium they have to pay makes them poorer and hence more risk averse. This will be the case, albeit weakly, in Subsection 3.2. The price of per capita damages in Eq. (10) equals the marginal certainty equivalent damages if the individual already suffers per capita damages, $p = d/dx\,(x\mu + A/2\,x^2\sigma^2)|_{x=1}$. Like the allocation of per capita damages, it does not depend on k.

For the CE and C&EQE consumption, we get

$$\bar{c} = c - \mu - \frac{A}{2}\sigma^2, \tag{12}$$

$$\hat{\bar{c}} = c + \frac{A}{2}\sigma^2 - \ln\left(1 - k\left(1 - e^{A\left(\mu + A\sigma^2\right)/k}\right)\right)/A \tag{13}$$

The corresponding damages are shown in Fig. 2. Due to the efficient risk sharing, the risk premium in the market allocation is reduced to the premium for the homogeneous case, i.e. Eq. (6) for $k = 1$. The EQE is not affected by an insurance market. If individuals are risk-neutral there is no reason for buying insurance.

The market equilibrium crucially depends on the information structure. The main dimensions are: (i) whether it is known how many individuals are affected and who they are, (ii) the probability distribution on aggregate damages, and (iv) whether all this information is public or private.

(i) If nobody knows whether she is affected and all individuals have the same probability of being affected, then individuals are homogeneous ex ante. A perfect insurance market then leads to a homogeneous distribution of consumption net of damages ex post, as well. This homogeneity is obtained via contracts that transfer consumption ex post, i.e. once damages have been realized, from the unaffected to the affected. This leads to an increase of the EQE compared to the case where it is known who is affected. Thus, less information about who is affected increases EQE consumption. This is an instance of the well-known Hirshleifer effect (Hirshleifer, 1971), which might be summarized as "realized risks cannot be insured and shared".

(ii) The same effect applies to information about the value of aggregate climate damages. Once damages are known, there is no way to share damage risk. Since it can be expected that uncertainty will be resolved over time, insurance contracts will either have to be made soon, which, of course, brings problems of its own, or insurance will loose some of its effectiveness.

(iii) If the information is asymmetric, i.e. if, for instance, the affected know that they are affected but others don't, or if hidden actions influence damages, the classical problems of adverse selection and moral hazard would also hamper insurance markets and bring the resulting allocation closer to the one in Subsection 2.1.

2.3. Self-Insurance

Even if insurance contracts are not available, affected individuals can use savings, or self-insurance, to mitigate utility losses. We assume there are two periods, where the first period covers t_1 years. Damages occur only in the second period. Self-insurance is done by increasing savings in the first period and thereby shifting consumption to the second period. We can decompose damages into expected damages and a zero-mean risk, $D = \mu + D_0$ and then distinguish a deterministic and a stochastic reason for increasing savings, namely (i) consumption smoothing and (ii) precautionary savings (see e.g. Gollier, 2004). (i) Expected damages decrease the second period consumption level. This increases marginal utility and hence the propensity to save in the first period. (ii) If the decision maker is prudent, i.e. if she has convex marginal utility, then the zero mean risk increases marginal utility and hence savings (Jensen's inequality).

More formally, we denote the interest rate by r and the pure rate of time preference by β. The endowments in the two periods for both cohorts are denoted by c_1 and c_2, respectively, where the subscripts denote the time period not the cohort. We assume individuals maximize the sum of discounted utility over time. Thus the affected and unaffected cohorts solve the independent maximization problems

$$\max_{s^{(1)}} \left\{ u(c_1 - s^{(1)}) + e^{-\beta t_1} E\left[u\left(c_2 + s^{(1)} e^{rt_1} - D/k \right) \right] \right\}$$
$$\max_{s^{(2)}} \left\{ u(c_1 - s^{(2)}) + e^{-\beta t_1} u\left(c_2 + s^{(2)} e^{rt_1} \right) \right\}.$$

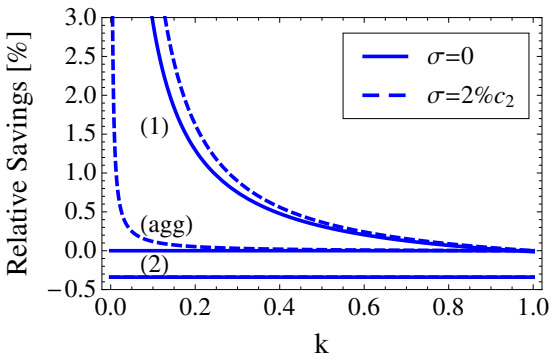

Figure 3: Additional savings relative to consumption as a function of the infliction rate k. Savings for affected, unaffected, and aggregate savings are denoted by (1), (2), and (agg), respectively. The first period covers $t_1 = 50\,yrs$, it is $\beta = 0.1\%$, $c_1 = \$7000$ and $c_2 = \$21000/Cap/yr$. The other parameter values are as in Fig. 1.

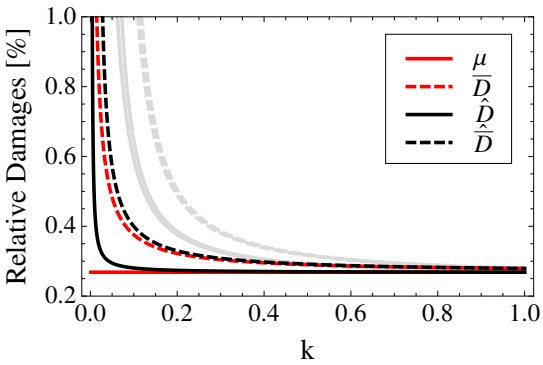

Figure 4: The same as in Fig. 1 in the same color code but for the two-period model with self-insurance. Parameter values are as in Fig. 3. Damages without self-insurance are shown in light gray.

For the functional forms assumed in this section we get

$$
\begin{aligned}
s^{(1)*} &= \left(1+e^{rt_1}\right)^{-1}\left(c_1 - c_2 + \frac{t_1(r-\beta)}{A} + \frac{\mu}{k} + \frac{A}{2}\frac{\sigma^2}{k^2}\right), \\
s^{(2)*} &= \left(1+e^{rt_1}\right)^{-1}\left(c_1 - c_2 + \frac{t_1(r-\beta)}{A}\right), \qquad (14)
\end{aligned}
$$

The last two terms in the second factor of $s^{(1)*}$ equal the certainty equivalent damages of the affected and describe additional savings due to damages. The former of them is due to consumption smoothing, the latter is due to prudence. Savings are increasings in the interest rate (and the first period length) if it is low, but decreasing if it is high. The reason for the latter is that a higher interest rate provides higher consumption in the second period, which decreases the incentive to save.

In order to isolate the additional savings due to heterogeneity, we can choose the interest rate r such that individuals have no incentive to save in the homogeneous case ($k = 1$), i.e. $u'(c_1) = e^{(r-\beta)t_1}E[u'(c_2 - D)] \rightarrow s^{(1)*}|_{k=1} = 0$. Under the functional forms of this section, we get $r = \beta + (A/t)(c_2 - \mu - (A/2)\sigma^2 - c_1)$. Substituting this into Eq. (14) leads to rather lengthy and little intuitive expressions. A numerical example is therefore shown in Fig. 3. The affected save substantially more, whereas the unaffected save less than in the homogeneous case. The latter is because the unaffected enjoy greater consumption in the second period than in the homogeneous case and hence shift consumption to the first period. The additional savings of the affected are mainly due to consumption smoothing (solid lines in Fig. 3). The aggregate additional savings are small down to about $k = 0.05$, which would justify the assumption of a fixed interest rate even in the presence of non-constant returns.

In order to measure the impact of self-insurance on welfare, we have to accommodate the temporal dimension. Therefore we generalize the certainty equivalent to a certainty and zero-growth equivalent (C&ZGE) consumption $\tilde{\bar{c}}^{(1)}$, where the tilde refers the the ZGE and the bar to

the CE. For the affected without self-insurance, for instance, it is implicitly defined over

$$u(c_1) + e^{-\beta t_1} E\left[u\left(c_2 - D/k\right)\right] \quad = \quad u\left(\tilde{\bar{c}}^{(1)}\right)\left(1 + e^{-\beta t_1}\right), \tag{15}$$

An arbitrary consumption vector $(c_1\, c_2)$ is replaced by a constant one $\left(\tilde{\bar{c}}^{(1)}\, \tilde{\bar{c}}^{(1)}\right)$ that yields the same utility. For the more general concept of balanced growth equivalents, where consumption grows at a constant rate instead of being constant, see Mirrlees & Stern (1972) and Anthoff & Tol (2009). Parallel to Eq. (3), the welfare function is defined as the following sum over the two cohorts:

$$W(k) \quad = \quad k\, g\left(\tilde{\bar{c}}^{(1)}\right) + (1-k)\, g\left(\tilde{\bar{c}}^{(2)}\right),$$

and the C&ZG&EQE $\hat{\bar{c}}$ is then defined analogous to Eq. (4). Its explicit form under the functional assumptions of this section is lengthy, so that we show a numerical example in Fig. 4 instead. It shows that self-insurance substantially reduces C&ZG&EQE damages for low k. The fact that damages without risk and inequality aversion are not independent of k if self-insurance is not available (lower solid gray line) is due to the finite elasticity of intertemporal substitution. Comparing Fig. 4 to Fig. 2 shows that self-insurance mainly lowers welfare losses due to inequality, whereas insurance mainly lowers the risk premium.

In the last Section we highlighted that the equilibrium in an insurance market crucially depends on the information available about who is affected and about the value of aggregate damages, and that utilitarian social welfare is larger the greater the uncertainty. In contrast, self-insurance is hampered by uncertainty. Only having a certain probability of being affected, for instance, lower savings to a level, which is ex post inefficient if the individual is actually affected. As it should be expected in a situation where individuals are independent of each other, information enhances the welfare gains from self-insurance.

3 Numerical Model

We now use the integrated assessment model DICE (Nordhaus, 2008) to obtain more realistic results. DICE is a Ramsey-type growth model coupled to a simple climate box model that translates greenhouse gas emissions resulting from economic production to concentration, radiative forcing, atmospheric and oceanic warming and finally economic impacts. In addition to the investment into the aggregate capital stock, there is a second decision variable called the emissions control rate, which reduces emissions at given abatement costs.

We replace the assumptions of the analytical model in the last section by the following more realistic ones: (i) All agents are described by a constant relative risk aversion utility functions, $u(c) = (c^{1-\gamma} - 1)/(1-\gamma)$, with the same relative risk aversion $\gamma = 3$. In contrast to Nordhaus, who uses a pure rate of time preferences that declines from $\beta = 1.5\%/yr$ to zero over time, we choose a constant $\beta = 0.1\%/yr$ (see Dasgupta, 2008, for a justification). (ii) We use $\mathcal{L}og - \mathcal{N}(ln(2.6), 0.33)$ (Wigley & Raper, 2001) as probability density function (PDF) on climate sensitivity. In the DICE function for aggregate relative damages as a function of global mean temperature $d(T) = aT^b/(1+aT^b)$ we use the joint PDF on a and b derived from an expert elicitation by Roughgarden & Schneider (1999). Fig. 5 shows exemplary distributions. In the following we use an equiprobable descriptive sampling with 10×10 sample points to represent the uncertainty. (iii) Estimates of the

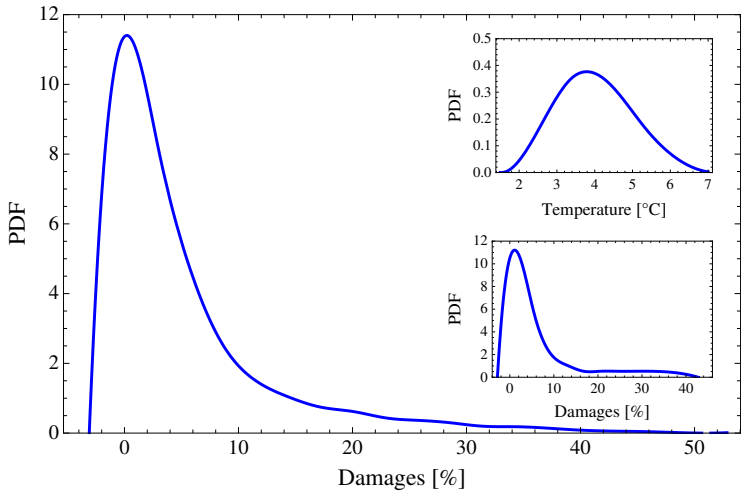

Figure 5: Climate damage PDF in 2100 for a 1000 ppmv concentration target. The inset upper graph shows the warming for the same year and scenario, and the lower one shows the damage distribution for the fixed average warming in this year of $T = 4.04°C$.

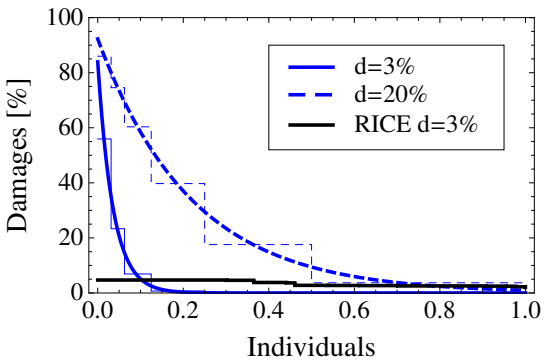

Figure 6: Distribution of relative damages over individuals for three different values of aggregate damages indicated in the plot legend. The value of the heterogeneity parameter is $\eta = 0.05$.

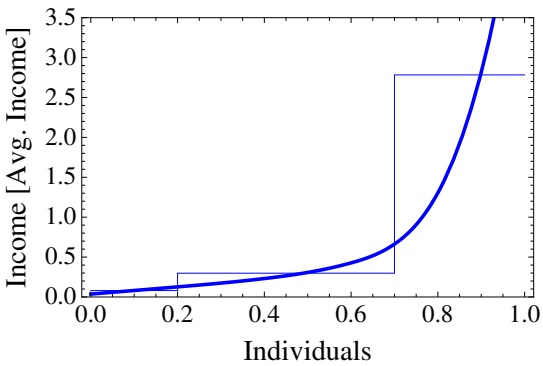

Figure 7: Global income in units of average income sorted in descending order. The thin straight lines show the 3-point discretization. Income inequality will be discussed at the end of Subsection 3.1.

geographic distribution of climate damages are only available on a world regional or at best on a country scale. The damage heterogeneity in the model RICE (Nordhaus & Yang, 1996) is included in Fig. 6. These estimates neglect the intra-regional heterogeneity. It can be expected that this heterogeneity is substantial for damages e.g. from extreme weather events. Therefore we use a conceptual parametrization of heterogeneity and perform sensitivity analysis. In contrast to the simple parametrization in the analytical model in Section 2, increasing aggregate damages lead to both higher damages for the affected and a greater share of affected individuals. To take this into account, we use the following parametrization: We index individuals by $i \in [0, 1]$ and assume individual damages are described by $\delta(i) = d^{\eta} e^{-b\,i}$, where d are average damages, η is a free parameter for the degree of heterogeneity, which now replaces the k of Section 2, and b is chosen such that the average damages actually equal d, $\int_0^1 \delta(i) di = d$. The homogeneous distribution is obtained for $\eta = 1$. The distribution of damages over individuals for a fixed value of $\eta = 0.05$ is shown in Fig. 6. In the following we use a discretization of the parametrization with six cohorts at $i = 1$, $1/2$, $1/4$, $1/8$, $1/16$, $1/32$, which also shown in Fig. 6.

Solving DICE with uncertainty about climate sensitivity and heterogeneous damages is numerically very intensive. Therefore, we structure the decision space into 13 concentration targets from 400 to 1000 ppmv CO_2eq in steps of 50 ppmv. For each target, we maximize utility for homoge-

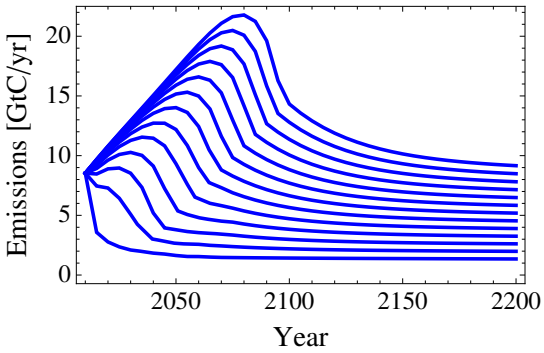

Figure 8: CO_2 emissions over time for the 13 different concentration targets. The lowest target of 400 ppmv, of course, implies the lowest emissions.

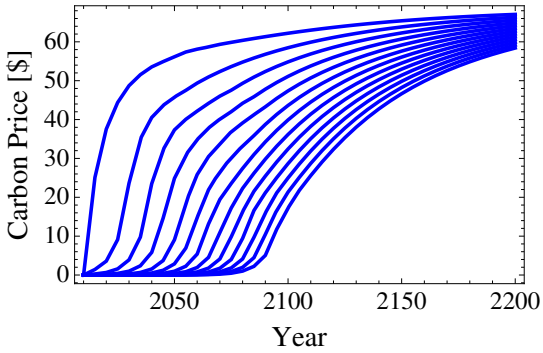

Figure 9: (Current value) carbon price over time for the 13 concentration targets. The lowest target of 400 ppmv, of course, implies the highest carbon price.

neous damages and under the respective concentration constraint in a deterministic optimization, i.e not taking uncertainty nor heterogeneity into account. This gives us the time paths of the decision variables that we will associate with the targets. The emissions and carbon price paths are shown in Figs. 8 and 9.We then evaluate these 13 targets and associated decisions taking uncertainty, damage heterogeneity, insurance markets and self-insurance into account. Thereby, we assume that abatement costs are distributed homogeneously among the population. For most of this section we also neglect income inequality in order to isolate the effect of damage heterogeneity. The interaction between income and damage inequality is discussed at the end of Subsection 3.1.

As we will see, structuring the decision space into 13 targets and evaluating them in different settings is not only more convenient, but also brings some added value. It allow us to analyze the differential effect of heterogeneity, insurance and so on across the policy space. We can also easily calculate opportunity costs of choosing suboptimal policies, which is necessary to assess whether changes in optimal decisions are accompanied by significant welfare improvements.

Parallel to Section 2, we will discuss the results without insurance, with insurance and with self-insurance in Subsections 3.1, 3.2, and 3.3, respectively.

3.1. No Insurance

For each target and associated control path, we calculate the discrete probability distributions on average consumption and damages. We then calculate heterogeneous damages and net consumption for each cohort in each state of the world. Subsequently we calculate expected utility for each cohort and aggregate to overall welfare.

Fig. 10 shows the ZGE consumption levels with and without damage heterogeneity and for the different targets. Uncertainty, despite its considerable dispersion shown in Fig. 5, has a very small effect on welfare and consequently almost no differential effect on the different targets in the homogeneous case. The optimal target is lowered from 650 ppm to 600 ppm but the resulting welfare improvement is negligible.

In contrast, uncertainty has a strong effect on welfare, if we introduce a pronounced heterogeneity described by $\eta = 0.05$ and shown in the right panel of Fig. 10. This effect is roughly doubled if utilitarian inequality aversion is assumed. The heterogeneity also has a strong differential effect: It makes high concentration targets less attractive by penalizing the bigger uncertainty they imply. The effect of inequality aversion without uncertainty, however, is small. Hence the separate effects

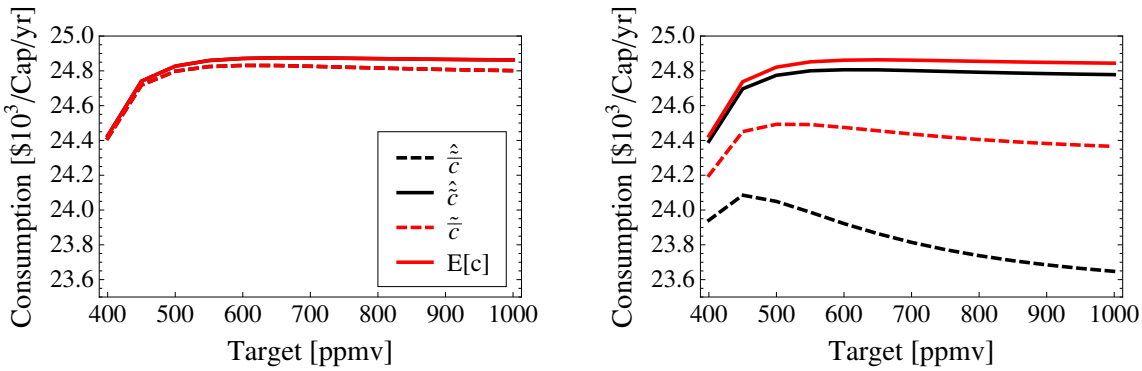

Figure 10: ZGE consumption with (black lines) and without (red lines) inequality aversion, and with (dashed lines) and without (solid lines) risk aversion for different concentration targets. The legend in the left graph applies to the right-hand graph as well. The left-hand graph shows results for a perfectly homogeneous distribution of damages ($\eta = 1$), where inequality aversion doesn't have an effect (black and red lines coincide). The right-hand graph shows results for heterogeneous damages with $\eta = 0.05$. .

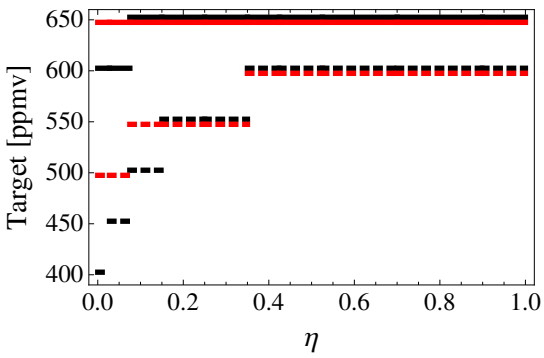

Figure 11: The optimal concentration target as a function of the heterogeneity parameter η: with inequality aversion (black) and without (red); with risk aversion (dashed lines) and without (solid line) risk aversion, i.e. the same color code as in Fig. 10.

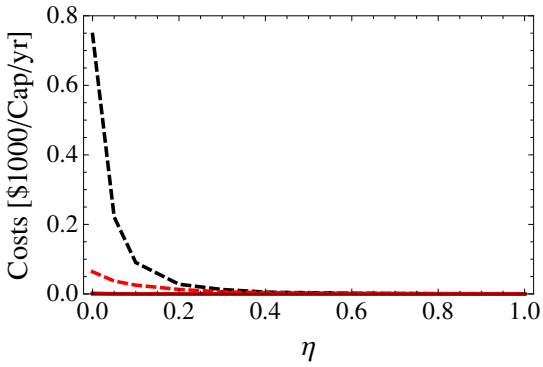

Figure 12: Consumption losses resulting from applying the optimal target under homogeneity and without uncertainty, which is 650 ppmv, rather than the optimal target with heterogeneity and uncertainty. The color code is the same as in Fig. 11.

of uncertainty and damage heterogeneity are negligible, whereas the joint effect is substantial.

The optimal target as a function of the heterogeneity parameter η and with and without inequality aversion is shown in Fig. 11. The optimum decreases down to 400 ppmv for very heterogeneous damages with, and to only 500 ppmv without inequality aversion. Without uncertainty, inequality aversion has only a minor effect on the optimal target. It changes the optimal target from 650 to 600 ppmv for small $\eta < 0.1$. The welfare losses measured in ZGE that are incurred if the optimal target without heterogeneity and without uncertainty, which is 650 ppmv, is applied under heterogeneity and uncertainty, again as a function of η, are shown in Fig. 12. The losses from pursuing the 650 ppmv target are only severe if inequality aversion is assumed and heterogeneity is pronounced. For $\eta = 0.05$ these losses amount to roughly \$200 C&EQ&ZGE consumption/Cap/yr. This is about a third of the benefit of taking any action against climate change at all, i.e. not following business as usual.

We now perform a sensitivity analysis with respect to the parameter γ in the utility function. This parameter not only determines risk aversion but also the elasticity of inter-temporal substitution. Therefore the efficient policies to achieve the 13 concentration targets without uncertainty,

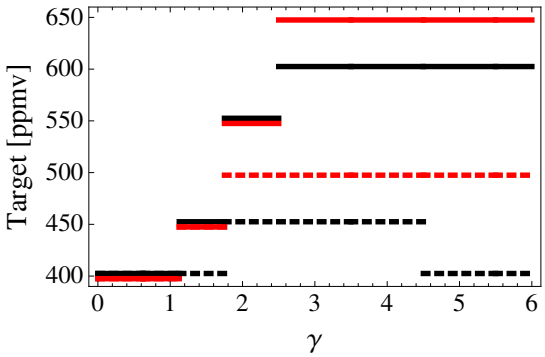

Figure 13: The optimal concentration target as a function of the parameter γ. It is $\eta = 0.05$. The color code is the same as in Fig. 10.

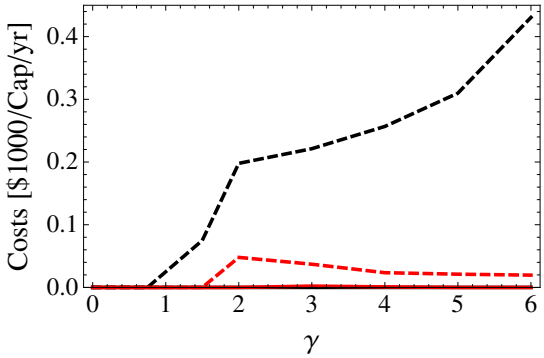

Figure 14: ZGE consumption losses resulting from applying the optimal target under homogeneity and without uncertainty, which depends on γ, rather than the optimal target with heterogeneity and uncertainty. It is $\eta = 0.05$. .

which we use to structure the decision space, depend on γ. We take this into account and change the policies with the value of γ. We do not change the pure rate of time preference, though, so that the consumption discount rate changes with γ.

The dependence of the optimal target on γ and for a heterogeneity parameter $\eta = 0.05$ is shown in Fig. 13. We note that the optimal target strongly decreases for decreasing γ even without uncertainty and inequality aversion. This is due to the fact that for decreasing γ, marginal utility decreases less rapidly, future consumption, which is higher than present consumption, becomes more valuable and hence future damages more painful, thus favoring strict targets (see also Nordhaus, 2007). With uncertainty, the target decreases at high values of γ. This is due to the fact that a large γ also implies large risk aversion and thus favors strict targets, which lead to less risk. The flatness of the dashed curves between $\gamma = 1$ and $\gamma = 5$ is then explained by the two opposite effects canceling out: an increasing γ puts less emphasis on future consumption but at the same time puts more emphasis on risk. Fig. 14 shows the losses resulting from choosing the optimal target without homogeneity and uncertainty rather than the truly optimal target. Again, without inequality aversion these losses can be neglected. Since the optimal target with inequality aversion does not change between $\gamma = 1$ and $\gamma = 5$, the increase of losses is due to a different valuation of same consumption losses and particularly the associated risk.

Up to now we have neglected income inequality in this section, in order to isolate the effect of damage heterogeneity. We now use a three-point discretization of the unequal income distribution shown in Fig. 7. It displays the income of individuals as a multiple of average income and is based on data from the World Development Report (World Bank, 2004). The index of individuals is not the same as the one for damage heterogeneity in Fig. 6, i.e. we do not assume perfect (anti-)correlation between relative damages and income. We rather consider two cases: (i) Relative damages are the same for all income classes. (ii) Relative damages are higher for low-income individuals. In both cases we assume that growth is distribution neutral, i.e. the income of all income classes grows at the same rate.

(i) Fig. 15 shows the ZGE consumption for the different concentration targets. Inequality aversion now makes a huge difference and has a far bigger effect than risk aversion even for a strong damage heterogeneity described by $\eta = 0.05$. For a utilitarian, income inequality is obviously the primary concern. (ii) It can be expected that relative damages are higher for poor countries and

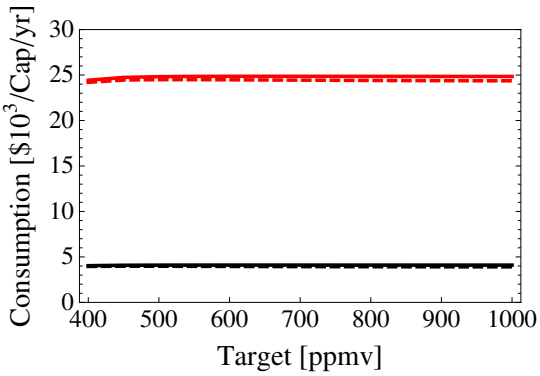

Figure 15: ZGE consumption under income inequality. It is $\eta = 0.05$. The color code is the same as in Fig. 11. The red lines are the same as in the right panel of Fig. 10.

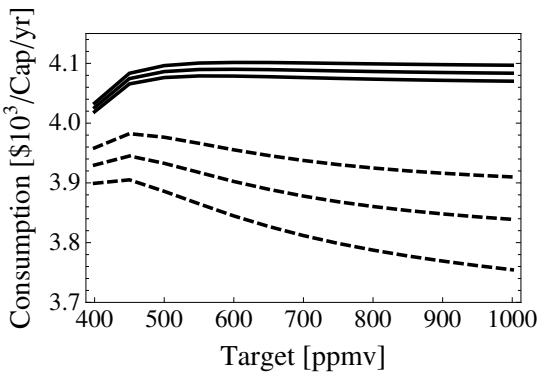

Figure 16: ZGE consumption under income inequality and biased relative damages. Relative damages of the low income cohort are increased by 0, 20, and 40% for the three solid and three dashed lines. Lower lines correspond to higher percentages.

low income classes (Yohe and Schlesinger, 2002). Fig. 16 shows the effect of such a negative relation between income and relative damages. More specifically, we increase relative damages of the poorest income class by 0, 20, and 40% and compensate this by decreasing relative damages of the richest income class thus keeping aggregate damages constant. The resulting effect on ZGE consumption is negligible without inequality aversion, the risk premium is the same as in Fig. 10 and therefore not shown. With inequality aversion, 40% bigger relative damages on the low income class instead of homogeneous relative damages decrease the C&EQ&ZGE consumption by about \$150 or roughly 3.5%. Hence, bigger relative risk for poor individuals notably increases the risk premium only under inequality aversion. Again, a substantial effect is only generated by compounding risk aversion and inequality aversion.

For the preceding results, we assumed that abatement costs are shared in proportion to income. Obviously, a progressive cost-sharing scheme, where percentage costs are higher for rich individuals than for poor ones, would have a welfare-enhancing redistributional effect under inequality aversion and hence favor stricter stabilization targets.

3.2. Perfect Insurance Market

We now introduce an insurance market parallel to Subsection 2.2. More specifically, there is a contingent claims market for each time period and contingent claims are paid for and pay off in the same period. Endowments are determined by the 13 concentration targets. In contrast to Subsection 2.2, it is not sufficient to introduce one claim that pays off aggregate per capita damages in all states of the world characterized by aggregate per capita damages. Individual damages are now non-linear in aggregate damages, and individuals therefore want to buy different multiples of per capita damages in different states of the world.

Technically, for each of the 10×10 uncertain states of the world we introduce a contingent claim. We derive the first order conditions for each cohort and contingent claim analytically and furthermore impose market clearing conditions in each state of the world. The resulting system of equations is solved numerically.

Fig. 17 shows the ZGE consumption with insurance. Comparing the results with the ones without insurance shows that insurance substantially reduces the risk premium both with and

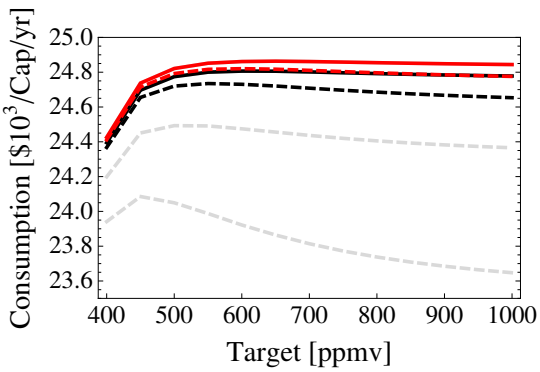

Figure 17: The same as in Fig. 10 for the perfect insurance market solution. The solution without insurance is shown in light gray for comparison.

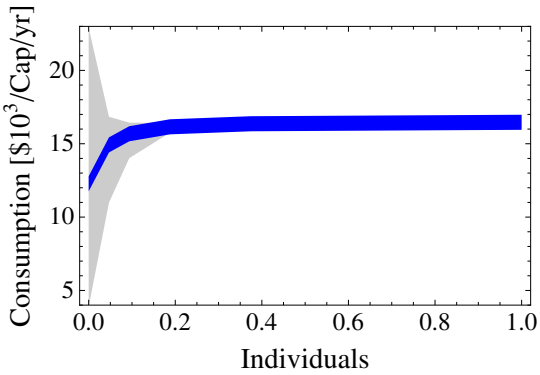

Figure 18: 10% and 90% quantiles of the PDF of consumption for the 600 ppmv target and for different individuals. The blue area is with insurance market and the gray area is without. The upper edge of the gray area is due to negative damages, i.e. benefits, from climate change.

without inequality aversion. The reason, as discussed in detail in the analytical model in Subsection 2.2, is that the market efficiently distributes the risk over the entire population thereby reducing the individual and aggregate risk premium. This is also highlighted in Fig. 18, which shows how strongly the risk born by the low-i individuals is reduced. Actually, the most affected individuals carry a (slightly) lower than average risk due to their lower expected consumption and resulting higher absolute risk aversion. The reduction of individual risk premia also leads to a reduction of inequality, which explains the diminished effect of inequality aversion in Fig. 17.

In the presence of a perfect insurance market, the optimal targets under heterogeneity differ only very little from the ones under homogeneity, and the losses from not taking this into account are negligible. Hence, under the strong assumption that the distribution of damages over individuals is known and that a perfect insurance market can be installed for all periods, even strong heterogeneity would not have a significant impact on optimal climate policy.

3.3. Self-Insurance

Parallel to Subsection 2.3, we now assume that heterogeneous individuals can make additional savings or dis-savings at a fixed interest rate. In other words, we assume approximately constant returns to scale justified by the presumed smallness of aggregate additional savings. For each target, the time-varying interest rate is determined in the homogeneous and deterministic case, for which the policies were optimized,

$$P_t u'(c_t - E[D_t]) \quad = \quad e^{(r(\Delta t)-\beta)\Delta t} P_{t+\Delta t} u'(c_{t+\Delta t} - E[D_{t+\Delta t}]), \tag{16}$$

where P_t is the population at time t. At this interest rate, no additional savings are optimal under certainty and homogeneity. Zero additional savings are generally not optimal, though, if uncertainty is taken into account, even if damages remain homogeneous. Due to the smallness of the risk premia for homogeneous damages, though, optimal additional savings due to uncertainty are less than 1% of overall savings for all targets and periods.

For heterogeneous damages, however, individual savings change considerably. Fig. 20 shows additional savings in 2010. The most affected individuals save about 30% more under the 400

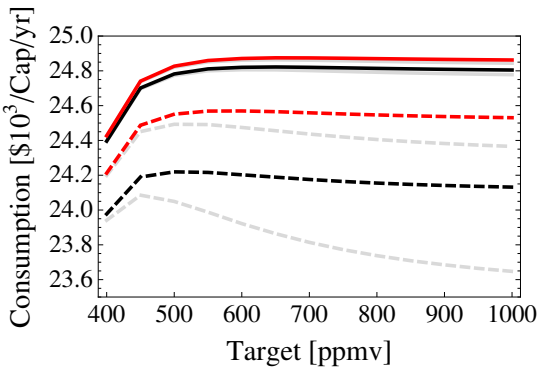

Figure 19: The same as in Fig. 10 but including self-insurance. The results without self-insurance are shown in light gray for comparison.

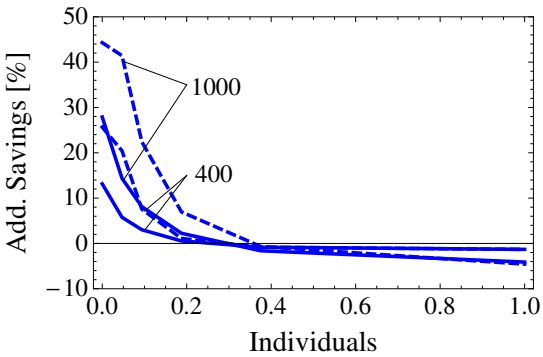

Figure 20: Additional savings due to heterogeneity in 2010 for 2 targets, 400 and 1000 ppmv, where the latter is the steeper curve. The heterogeneity parameter is $\eta = 0.05$.

ppmv target and about 45% more under the 1000 ppmv target, of which about 13% and 30%, respectively, are due to deterministic consumption smoothing and the rest is due to precautionary saving. Aggregate savings increase by 2.9% and 1.5% respectively. These results are obtained by numerically solving independent consumption-savings problems with exogenous interest rate for the different cohorts.

The welfare effect of self-insurance is depicted in Fig. 19. Self-insurance, of course, improves welfare for all targets but particularly for high concentration targets. The improvement for 1000 ppmv, for instance, is about $500 ZGE. Self-insurance is particularly effective for lax targets because mitigation costs for these targets are low and thus consumption in early periods is high. Savings can shift this consumption to later periods with high damages.

An important caveat for these results is the assumption that increased savings do not lead to increased damages. The results would change dramatically if the savings of the affected were diminished by the same damage factor as their gross consumption. However, this is not quite realistic either. In well functioning capital markets it should be possible to choose investments that are impacted at least only by the average damage factor across the population. Under this assumption, impacts on savings turn out not to have a significant effect on the results shown in Figs. 19 and 20. The truth presumably lies somewhere in between.

4 Conclusions

We have first demonstrated how climate damage heterogeneity and uncertainty can jointly have a big effect on certainty- and equity equivalent damages, particularly with but also without inequality aversion. Numerical results from the DICE model later showed that this can lead to a substantially stricter optimal stabilization target even if the separate effects of uncertainty and heterogeneity are negligible. This latter result hinges on the presence of inequality aversion and thus emphasizes again the importance of equity considerations in climate change. Taking heterogeneity into account becomes more important the higher the relative risk aversion of the individuals.

Income inequality is presumably a far greater concern to a utilitarian than climate change. However, we showed to what extent it favors strict targets if there is a pronounced negative correlation between income and relative damages.

We then studied two “instruments” that can mitigate the effect of damage heterogeneity and

uncertainty: insurance markets and self-insurance. A perfect insurance market leads leads to an efficient distribution of climate damages and the associated risk across the entire population. This reduces the risk premium essentially to the one for homogeneous damages. Some heterogeneity persists, though, because affected individuals have to pay insurance premia. The resulting effect on the optimal target under inequality aversion, however, turns out to be small in DICE. The presence of insurance markets thus would allow a weakening of the stabilization target and lead to substantial welfare gains. This indicates a large theoretical potential of insurance of climate damage uncertainty. However, the large time horizon and multiple market failures involved will certainly impede these markets.

Self-insurance, i.e. the increase in savings of the above-average impacted individuals, is not as effective as insurance markets in mitigating damage heterogeneity but still improves the attractiveness especially of less stringent concentration targets. The reason is that these targets imply low costs in the short run but high damages in the long run, which can partly be offset by increased savings. As a result, welfare differences between concentration targets of 500-1000 ppmv CO_2eq vanished in DICE even for pronounced damage heterogeneity.

Improved information about who is affected by climate change and about the aggregate amount of damages decreases the effectiveness of insurance and increases the effectiveness of self-insurance resulting in an ambiguous overall effect of this information on welfare.

The following main caveat applies to the analysis. Its results are conceptual to the same extent as our parametrization of climate damage heterogeneity, which was a result of the lack of geographically explicit global estimates of climate damages and the associated uncertainty. However, the results emphasize the need for such estimates and for subsequent analyses that explicitly take heterogeneity and uncertainty into account.

References

D. Anthoff and R. S.J. Tol. On international equity weights and national decision making on climate change. *Journal of Environmental Economics and Management*, 60(1):14–20, July 2010.

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.

D. Anthoff, C. Hepburn, and R. S. J. Tol. Equity weighting and the marginal damage costs of climate change. *Ecological Economics*, 68(3):836–849, 2009.

C. Azar. Weight factors in Cost-Benefit analysis of climate change. *Environmental & Resource Economics*, 13(3):249–268, 1999.

W. R. Cline. *The Economics of Global Warming*. Institute for International Economics,U.S., first edition, first printing edition, September 1992.

G.M. Constantinides. Intertemporal asset pricing with heterogeneous consumers and without demand aggregation. *The Journal of Business*, 55(2):253–267, April 1982.

P. Dasgupta. Discounting climate change. *Journal of Risk and Uncertainty*, 37(2-3):141–169, 2008.

S. Fankhauser. The economic costs of global warming damage: A survey. *Global Environmental Change*, 4(4):301–309, December 1994.

S. Fankhauser and R.S.J. Tol. On climate change and economic growth. *Resource and Energy Economics*, 27(1):1–17, 2005.

S. Fankhauser, R. Tol, and D. Pearce. The aggregation of climate change damages: a welfare theoretic approach. *Environmental & Resource Economics*, 10(3):249–266, 1997.

C. Gollier. *The Economics of Risk and Time*. The MIT Press, new edition edition, September 2004.

C. Gollier. Some aspects of the economics of catastrophe risk insurance. Technical Report 1409, CESifo Group Munich, 2005.

J. Hirshleifer. The private and social value of information and the reward to inventive activity. *American Economic Review*, 61(4):561–74, 1971.

C. W. Hope. The marginal impacts of CO2, CH4 and SF6 emissions. 6(5):537–544, 2006.

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

E. Mills. Insurance in a climate of change. *Science*, 309(5737):1040 –1044, 2005.

J. A. Mirrlees and N. H. Stern. Fairly good plans. *Journal of Economic Theory*, 4(2):268–288, 1972.

W. Nordhaus. Critical assumptions in the stern review on climate change. *Science*, 317(5835):201 –202, July 2007.

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

W. D. Nordhaus. Geography and macroeconomics: New data and new findings. *Proceedings of the National Academy of Sciences of the United States of America*, 103(10):3510 –3517, March 2006.

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

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 and Z.L. Yang. A regional dynamic general-equilibrium model of alternative climate-change strategies. *American Economic Review*, 86(4):741–765, September 1996.

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

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

R. Tol. Estimates of the damage costs of climate change. part 1: Benchmark estimates. *Environmental & Resource Economics*, 21(1):47–73, 2002.

R.S.J. Tol. Is the uncertainty about climate change too large for expected cost-benefit analysis. *Climatic Change*, 56:265—289, 2003.

A. Ulph and D. Ulph. Global warming, irreversibility and learning. *The Economic Journal*, 107 (442):636–650, May 1997.

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

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

WorldBank. World development report, 2004.

G. Yohe and M. Schlesinger. The economic geography of the impacts of climate change. *Journal of Economic Geography*, 2(3):311 –341, July 2002.

Chapter 3

Climate Targets under Uncertainty: Challenges and Remedies[1]

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

[1]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. $Z = 2°\text{C}$, 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

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

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 $E_1(t)$, $t \le t_l$. 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

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

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. $\sum_{m\in M} q_m R_m(E_1, E_m) \le Q$. 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, 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 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

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

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 $\max_{E_1}\left\{\sum_{m\in M} q_m \max_{E_m}\left\{\int f(\theta|m) A\left(E_1, E_m, \theta\right)\right\}\right\}$ 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.

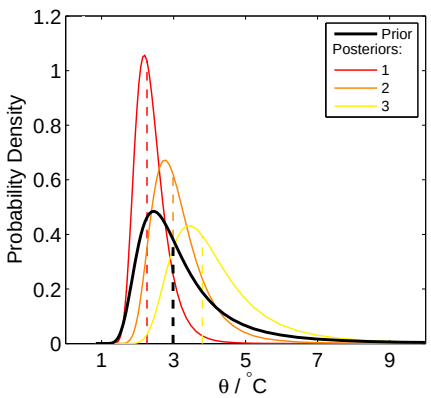

Fig. 1 Information structure with prior distribution and three posterior distributions. The vertical dashed lines indicate the medians.

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 4

Anticipating Climate Threshold Damages[1]

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

[1]This chapter has been accepted for publication in Environmental Modeling and Assessment as "Lorenz, A., M.G.W. Schmidt, E. Kriegler, H. Held. *Anticipating Climate Threshold Damages.*"

Anticipating Climate Threshold Damages

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

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 MIND 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".

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. Lorenz (corresponding author), M.G.W. Schmidt, E. Kriegler
Potsdam Institute for Climate Impact Research, 14412 Potsdam, Germany
E-mail: lorenz@pik-potsdam.de. Tel.: +49-331-2882562
†H. Held
University of Hamburg - Klima Campus, Bundesstr. 55, 20146 Hamburg, Germany
and Potsdam Institute for Climate Impact Research, 14412 Potsdam, Germany

[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. (2008) and Kriegler (2009).

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 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 (1980), who showed that it depends on the properties of the 2nd-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, 1951) unambiguously, i.e. independent of the specific form of the information structure (in the sense of Marschak & Miyasawa, 1968), 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, 1996; Gollier et al., 2000). 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 (2006), and Salanie & Treich (2009) 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. (2006), Bosetti et al. (2008), Johansson et al. (2008), and parts of Webster et al. (2008). 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. (2011) 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 & Teisberg (1993), Yohe & Wallace (1996), Kelly & Kolstad (1999), Leach (2007), and parts of Webster et al. (2008). See Lange & Treich (2008) 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 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. (2004). 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 & 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) $x_1 = I(t)$, $t_0 < t \leq t_{lp}$ and second period decisions $x_2(y) = I(y)(t)$, $t_{lp} < t \leq T$, 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 inter-temporally 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 & Miyasawa (1968) as presented in Jones & Ostroy (1984). 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

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

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.

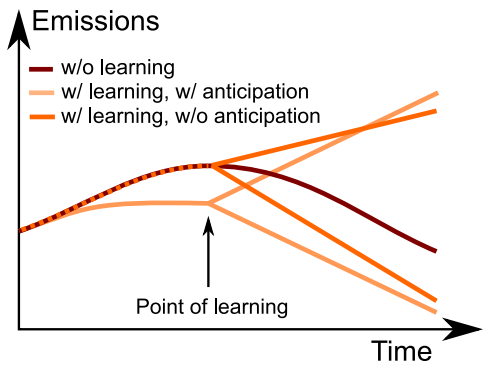

Figure 1: Schematic plot of optimal emissions over time under different information scenarios and for two learning paths.

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

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

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.

The anticipatory adjustment of first period decisions to future learning can be measured by the *Expected Value of Anticipation* (EVOA):

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

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

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

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 & Stern (1972). 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 & Tol (2009), 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:

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

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 & Tol, 2009). Intuitively, a 1% reduction in CBGE, for instance, can be interpreted as a permanent loss of consumption of 1%.

2.3 The Integrated Assessment Model MIND

We use the *M*odel of *In*vestment and Technological *D*evelopment (MIND) (Edenhofer et al., 2005)[2]. We use the version from Held et al. (2009) and add anticipated learning about uncertainty (see 2.1), but we leave out carbon capturing and sequestration (CCS) for tractability. Edenhofer et al. (2005) and Held et al. (2009) perform cost-effectiveness analysis for a given climate target. We have shown elsewhere (Schmidt et al., 2011) 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, 1993). 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 (**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 social planner 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). 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., 1999; Kriegler et al., 2007). SO_2 emissions are coupled to CO_2 emissions with an exogenously declining ratio of sulfur per unit CO_2 representing desulfurization. Radiative forcing from other greenhouse gases and aerosols is included as exogenous scenario (see Held et al., 2009).

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:

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

where $I = (I_{\mathrm{K}}, I_{\mathrm{R\&D}}, I_{\mathrm{Fossil}}, I_{\mathrm{Renewables}})$ is the vector of investment flows in the different

[2]Modified model versions feature an endogenous carbon capturing and sequestration (CCS) module (Bauer, 2005), a more elaborate carbon cycle and atmospheric chemistry module (Edenhofer et al., 2006), and parametric uncertainty (Held et al., 2009).

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

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:

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

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

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

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 upwards and downwards. 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 emissions 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, 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 (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 ($t_{lp} = 2030$), intermediate ($t_{lp} = 2050$), and late learning ($t_{lp} = 2070$). The uncertainties are described by probability distribution functions, that are given explicitly in Appendixes A and B. 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. Thereby we do not chose an equiprobable spacing, but chose 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. (2004) have found significant changes in emissions due to anticipation of learning if a highly non-linear irreversible threshold is included in the analysis. More specifically they considered a possible shut-down of the North-Atlantic thermohaline (THC) circulation (Broecker, 1997). 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 (2007), we consider explicit tipping point-like threshold damages. Similar to Keller et al. (2004), 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 DF_{thresh} are uncertain. The damages are added to Nordhaus's damage factor DF leading to output net of damages, $Y_{\text{net}} = Y_{\text{gross}} \cdot DF_{\text{thresh}}$. 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

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

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

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

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 $ED_{\text{thresh}} = 1.5\%$ of net GDP is in accordance with empirical estimates for the expected impact of a THC shut-down by Tol (1998). 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 DF_{thresh} changes), and learning-points t_{lp}.

	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

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

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 Tab. 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 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 Tab. 1 is the proportion of the EVOI that is obtained by anticipatory changes in pre-learning decision, i.e. the ratio EVOA/EVOI (see Subsection 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 & Ulph, 1997; Nordhaus & Popp, 1997; Webster, 2002; O'Neill & Melnikov, 2008; Webster, 2008). Why could we have expected an effect in the model MIND? As shortly discussed in the introduction, 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, 1980), $\alpha\partial/\partial x_1 V_2(x_1, \pi_s^i) + (1-\alpha)\partial/\partial x_1 V_2(x_1, \pi_s^j) \neq \partial/\partial x_1 V_2(x_1, \alpha\pi_s^i + (1-\alpha)\pi_s^j)$. 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. Fig. 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 Subsection 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 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 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.

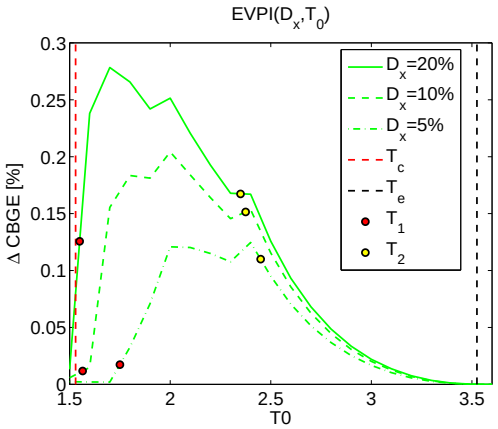

Figure 2: The EVPI for different values of D_x and T_0. The EVPI is measured in % $CBGE$ of the no-learning case. T_c denotes the temperature the decision maker is already committed to cross. For $T_0 > T_1(D_x)$ avoiding the threshold is optimal for perfect information that $D_{thresh} = D_x$. For $T_0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

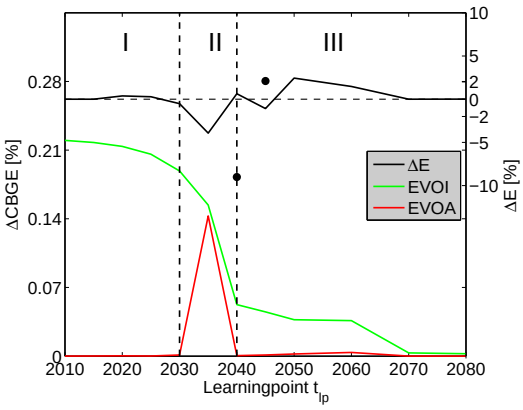

Figure 3: Expected Value of Information (EVOI), Expected Value of Anticipation (EVOA) and relative changes in cumulative pre-learning emissions in anticipation of learning (ΔE) is shown depending on the time of learning t_{lp}. The dashed lines mark three distinct regimes of anticipation (I-III). The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

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}. Fig. 3 shows the EVOA and EVOI for learning-points between the year $t_{lp} = 2010$ and $t_{lp} = 2080$ 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 $t_{lp} = 2200$. The latter is essentially the no-learning case. The EVOA has to be zero for $t_{lp} = 2010$ because there are simply no pre-learning decisions to be made. It is also zero for $t_{lp} = 2200$ 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 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 $D_{thresh} = D_x$ is learned leads to a substantial EVOI without the need for downward anticipation. Not having to anticipate downwards benefits the case where $D_{thresh} = 0$ is learned. There is even some upward anticipation to come closer towards the solution that would be optimal for perfect information about $D_{thresh} = 0$.

(II) For increasing t_{lp} there there is less time between learning and crossing the treshold (without adjustments). Since mitigation costs are convex, this increases the costs of avoiding the threshold in the "bad case" ($D_{thresh} = D_x$) by post-learning adjustments alone. Therefore, in regime II the DM lowers pre-learning emissions compared to the no-learning

[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.

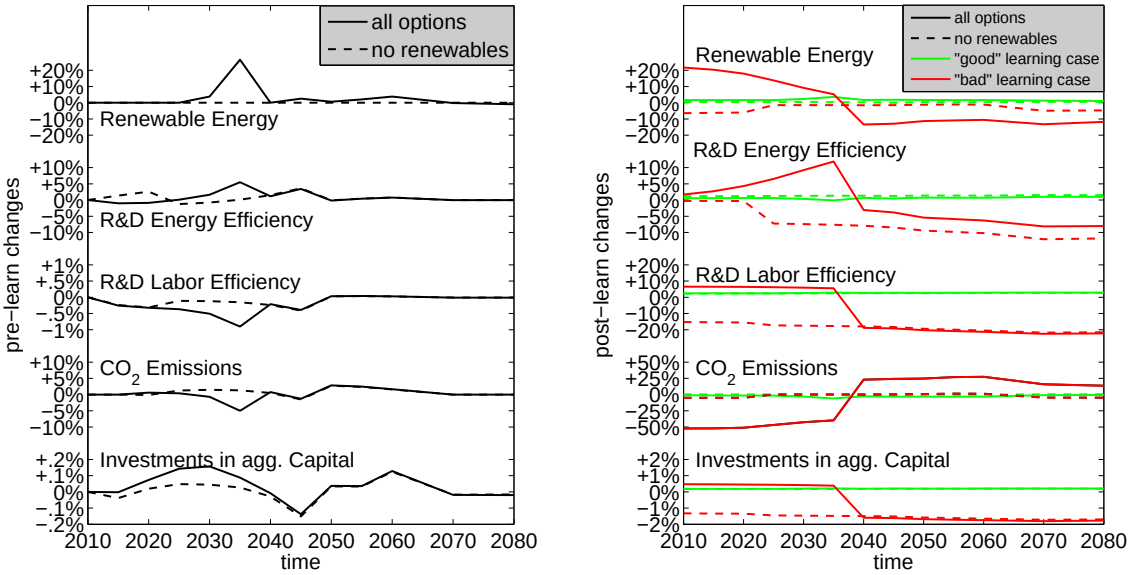

Figure 4: The anticipation effect (left) and post-learning decisions (right) both in cumulative decision variables and with and without the availability of renewable energy.

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 regime 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 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 $t_{lp} > 2040$ 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 $t_1 < t_{lp} < t_2$. This "anticipation window" is narrow, 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 Subsection 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 macroeconomic 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 $D_{\text{thresh}} = 0$ and $D_{\text{thresh}} = D_x$. The resulting EVOI and EVOA are shown in Fig. 5.

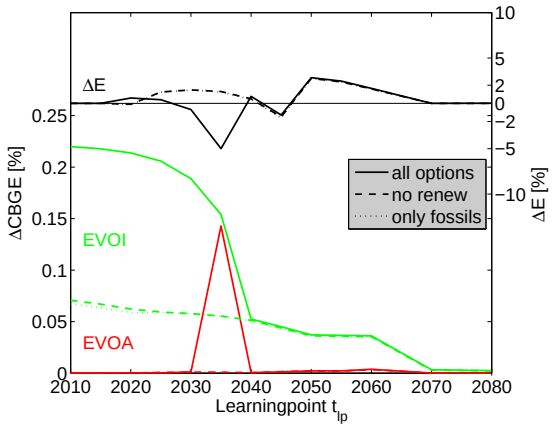

Figure 5: Expected Value of Information (EVOI), Expected Value of Anticipation (EVOA) and relative changes in cumulative pre-learning emissions in anticipation of learning (ΔE) is shown depending on the time of learning t_{lp}. Shown are three scenarios differing in the availability of mitigation options. 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 22nd 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.

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 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.

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 panels a-c of Fig. 6.

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 towards 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

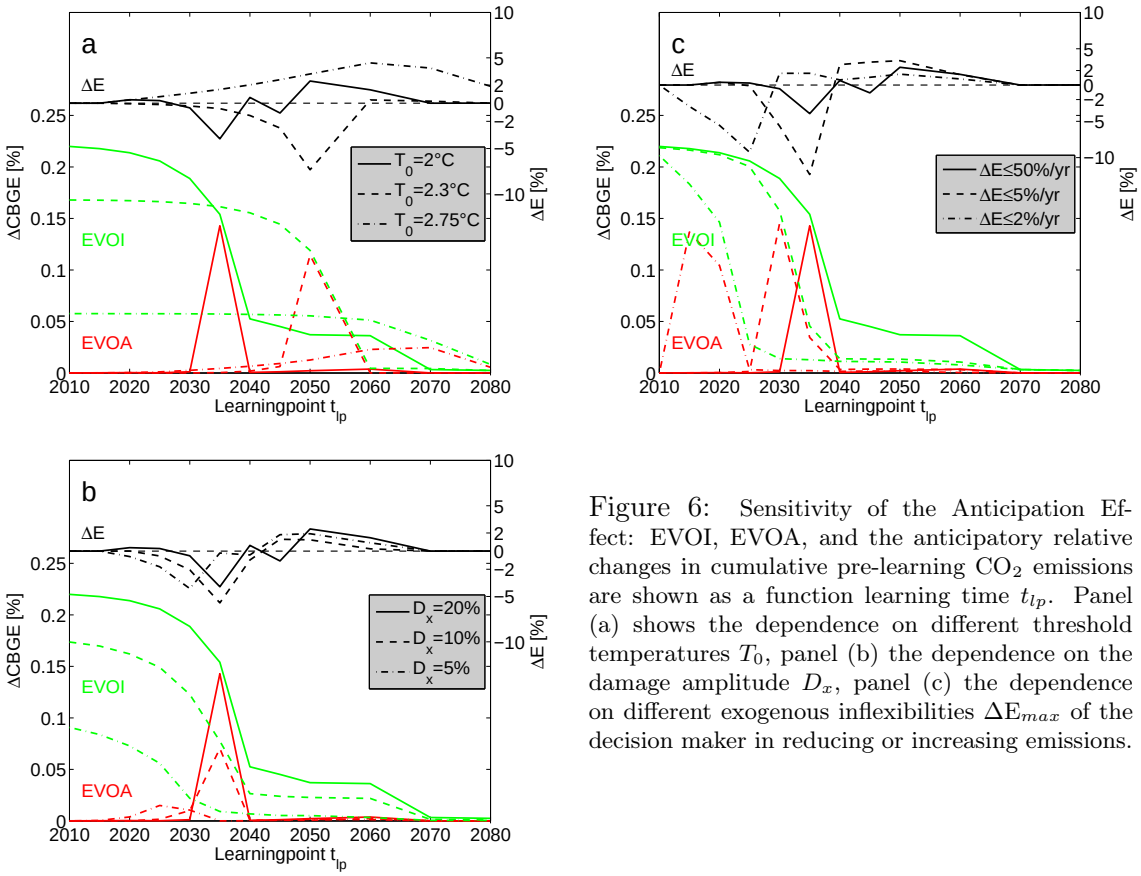

Figure 6: Sensitivity of the Anticipation Effect: EVOI, EVOA, and the anticipatory relative changes in cumulative pre-learning CO_2 emissions are shown as a function learning time t_{lp}. Panel (a) shows the dependence on different threshold temperatures T_0, panel (b) the dependence on the damage amplitude D_x, panel (c) the dependence on different exogenous inflexibilities ΔE_{max} of the decision maker in reducing or increasing emissions.

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 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 towards 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 towards 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 towards 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 ($\Delta E < 1\%/yr$) (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 lead 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 occured 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 socio-economic 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 towards a dynamic programming framework. Studying multiple, and partly reversible thresholds occuring 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: Allthough 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.

Acknowledgments

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

A 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:

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

where C is current and C_{pi} pre-industrial atmospheric CO_2 concentration, T denotes global mean temperature anomaly, 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:

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

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., 2002; Frame et al., 2005). 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 & Raper (2001): $\bar{\pi}(\text{CS}) = \mathcal{LN}(0.973, 0.4748)$.

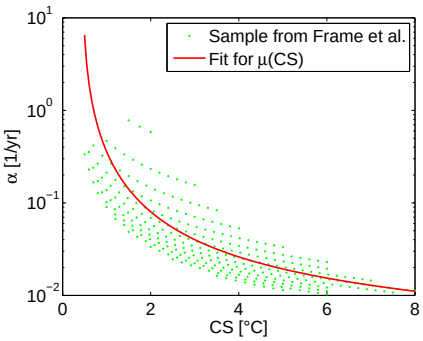

Figure 7: Correlation of α and CS from the temperature record of the last two centuries (see Frame et al., 2005) [green dots]. By assuming a strict relationship between $\frac{1}{\mu}$ and CS as $\frac{1}{\mu} = \frac{1}{\mu} - 10 \cdot \exp(-0.5\,\mathrm{CS})$ the correlation narrows to the [red curve].

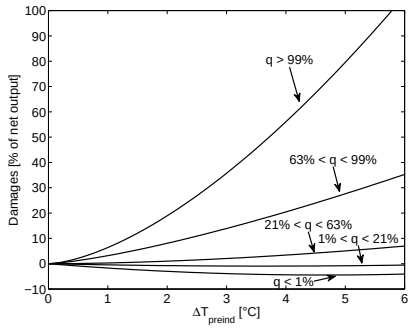

Figure 8: Samples taken according to a descriptive sampling scheme from a joint probability distribution of the damage function parameters a and b from Roughgarden & Schneider (1999). Shown are the damage functions representative for the quantiles q with probability weights $\omega_i = [1, 20, 60, 18, 1]\%$ that have been used within the experiments.

B Climate Damages

The uncertain parameters a and b in the exponential damage function $DF(T) = \frac{1}{1+a\,T^b}$ are determined from an expert-based assessment done by Roughgarden & Schneider (1999). 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. Fig. 8 shows the damage functions that represent the quantiles chosen for our experimental setup: $\omega_i = [1, 20, 60, 18, 1]\%$.

References

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

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

Bauer, N. (2005). *Carbon Capture and Sequestration: An Option to Buy Time?* PhD thesis, Potsdam University, Germany.

Blackwell, D. (1951). Comparison of experiments. In Neyman, J., editor, *2nd Berkeley Symposium on Mathematical Statistics and Probability*. University of California Press.

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.

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.

Edenhofer, O., Bauer, N., and Kriegler, E. (2005). The impact of technological change on climate protection and welfare: Insights from the model MIND. *Ecological Economics*, 54:277–292.

Edenhofer, O., Lessmann, K., and Bauer, N. (2006). Mitigation strategies and costs of climate protection: The effects of ETC in the hybrid model MIND. *Energy Journal*, pages 207–222.

Epstein, L. G. (1980). Decision making and the temporal resolution of uncertainty. *International Economic Review*, 21(2):269–83.

Forest, C. E., Stone, P. H., Sokolov, A. P., Allen, M. R., and Webster, M. D. (2002). Quantifying uncertainties in climate system properties with the use of recent climate observations. *Science*, 295(5552):113–117.

Frame, D. J., Booth, B. B. B., Kettleborough, J. A., Stainforth, D. A., Gregory, J. M., Collins, M., and Allen, M. R. (2005). Constraining climate forecasts: The role of prior assumptions. *Geophysical Research Letters*, 32(9).

Gollier, C., Jullien, B., and Treich, N. (2000). Scientific progress and irreversibility: an economic interpretation of the 'Precautionary principle'. *Journal of Public Economics*, 75(2):229–253.

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

Johansson, D. J. A., Persson, U. M., and 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.

Jones, R. A. and Ostroy, J. M. (1984). Flexibility and uncertainty. *Review of Economic Studies*, 51(1):13–32.

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.

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

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.

Kriegler, E. (2009). Updating under unknown unknowns: An extension of bayes' rule. *International Journal of Approximate Reasoning*, 50:583–596.

Kriegler, E., Held, H., and Bruckner, T. (2007). Climate protection strategies under ambiguity about catastrophic consequences. In *Advanced Methods for Decision Making*, pages 3–42. Nova Science Publishers, j. kropp, j. scheffran edition.

Lange, A. and 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.

Leach, A. J. (2007). The climate change learning curve. *Journal of Economic Dynamics & Control*, 31(5):1728–1752.

Marschak, J. and Miyasawa, K. (1968). Economic comparability of information systems. *International Economic Review*, 9(2):137–174.

Mirrlees, J. A. and Stern, N. H. (1972). Fairly good plans. *Journal of Economic Theory*, 4(2):268–288.

Nordhaus, W. (2007). The challenge of global warming: Economic models and environmental policy. Yale University.

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

Nordhaus, W. D. 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.

O'Neill, B., Ermoliev, Y., and Ermolieva, T. (2006). Endogenous risks and learning in climate change decision analysis. In *Coping with Uncertainty, Modeling and Policy Issues*, pages 283–300. Springer Verlag, Berlin, Germany.

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.

Oppenheimer, M., O'Neill, B., and Webster, M. (2008). Negative learning. *Climatic Change*, 89(1):155–172.

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.

Petschel-Held, G., Schellnhuber, H. J., Bruckner, T., Toth, F. L., and Hasselmann, K. (1999). The tolerable windows approach: Theoretical and methodological foundations. *Climatic Change*, 41(3-4):303–331.

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

Salanie, F. and Treich, N. (2009). Option value and flexibility: A general theorem with applications. Technical report, LERNA, University of Toulouse.

Saliby, E. (1997). *Descriptive Sampling: An Improvement Over Latin Hypercube Sampling.*

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

Tol, R. (1998). Potential slowdown of the thermohaline circulation and climate policy. *Institute for Environmental Studies Vrije Universiteit Amsterdam*, Discussion Paper DS98/06.

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

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.

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

Yohe, G. and 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

Uncertainty in Integrated Assessment Models of Climate Change[1]

Alexander Golub
Daiju Narita
Matthias G.W. Schmidt

[1]An earlier version of this chapter has been published as Golub, A., D. Narita, M.G.W. Schmidt 2011. *Uncertainty in Integrated Assessment Models of Climate Change: Alternative Analytical Approaches.* FEEM Working Paper No. 2.2011

Uncertainty in Integrated Assessment Models of Climate Change

Alternative Analytical Approaches

Alexander Golub[1], Daiju Narita[2], Matthias G.W. Schmidt[3]

Abstract

Uncertainty plays a key role in the economics of climate change, and research on this topic has led to a large body of literature. However, the discussion on the policy implications of uncertainty is still far from settled. Due to the complexity of the problem, an increasing number of analytical approaches have been used to examine the policy implications of uncertainty in integrated assessment models of climate change. We review these approaches, the corresponding literature and respective policy implications.

Keywords: Uncertainty, learning, economics of climate change, integrated assessment models, real options, dynamic programming

D 81, Q54, C61

1 Introduction

Although a large body of scientific evidence confirms the existence of the problem, the detailed mechanism and impacts of climate change are still uncertain. Climate change mitigation measures incur significant costs for society. These costs are mostly sunk and cannot be recouped if climate change turns out to be less severe than expected. Therefore, uncertainty about the benefits of mitigation is sometimes stated as an argument for deferring mitigation effort until more is known. In fact, this argument appears to have some intuitive appeal to the public. In a recent American poll[4], for instance, 30% of respondents agreed with the statement that "we don't know enough about global climate change, and more research is necessary before we take any actions."

However, holding actions against climate change renders the planet ever warmer, and weak current actions may be regretted later if the impacts of climate change turn out to be severe. This is an argument for precautionary mitigation actions and counteracts the argument above. Both arguments are based on the assumption that society will learn about the severity of

[1] Corresponding author: Environmental Defense Fund; 1875 Connecticut Ave., NW, Washington, DC 20009, USA Email: agolub@edf.org

[2] Kiel Institute for the World Economy; Hindenburgufer 66, 24105 Kiel, Germany Email: daiju.narita@ifw-kiel.de

[3] Potsdam Institute for Climate Impact Research; Telegraphenberg A31, 14473 Potsdam, Germany. Email: schmidt@pik-potsdam.de

[4] NBC News/WSJ, December, 2009

climate change in the future.

Another argument associated with uncertainty is based on risk aversion, or the fact that people are inclined to avoid uncertainty. On the one hand, it can be assumed that mitigation reduces uncertainty since we know better about climatic responses to low atmospheric concentrations of greenhouse gases than about those to elevated concentrations. On the other hand, mitigation costs are uncertain as well. Whether risk aversion is an argument for more or less mitigation depends on which type of uncertainty dominates.

Examination of these arguments has yielded a body of literature in the field of economics of climate change. It revolves around the basic question of how uncertainty and the reduction of uncertainty in the future affect optimal climate policy. This question can be broken down into the following sub-questions: How do individual types of uncertainty influence the optimal timing and stringency of climate policy? How does future learning influence optimal policy? What is the value of information about different uncertainties?

Theoretical studies have shed some light on these questions (e.g. Epstein, 1980; Baker, 2009) but the estimation of quantitative benchmarks usable in the policy debate necessitates the application of detailed models featuring the key mitigation technologies and key uncertainties of climate change. Therefore, an increasing number of studies utilize numerical integrated assessment models (IAMs), the primary tool for the investigation of complex climate-economy interactions.

Climate change involves various types and sources of uncertainty. There are parametric uncertainty and stochastic uncertainty in every link of the cause-effect relationship of climate change from emissions to global warming and impacts. Parametric uncertainty describes the incomplete knowledge of model parameters such as climate sensitivity. Stochasticity describes the persistent randomness of the system due to unresolved processes, e.g. in the development of global mean temperature or global economic output.

Another layer of complexity is the fact that knowledge is continuously updated due to scientific progress and measurements. The various uncertainties and their updating can be called the information dynamic complexity of the problem. Meanwhile, economic growth, technological progress, and climate change are undoubtedly complex processes. As a result, estimation of realistic probability distributions for mitigation costs and for avoided benefits under different policies is difficult. This can be called the system dynamic complexity of the problem.

The information and system dynamic complexity of the problem suggests that multiple complementary analytical approaches with different tradeoffs between the two complexities are required to grasp the full implications of uncertainty for climate policy. This article reviews the approaches that have been applied in the literature, and summarizes their respective policy implications. It follows the previous review articles by Kann & Weyant (2000), Heal &

Kriström (2002), and Peterson (2006). Apart from a discussion of recent contributions published after those reviews, a distinctive feature of our review is a detailed look at the complementarity of varied analytical approaches highlighting different facets of climate change uncertainty and at what induces their respective results for optimal policy.

We discuss analytical approaches in association with the issues for which each of them has a methodological advantage. More precisely, our focus lies in the following: (i) Non-recursive stochastic programming (NSP) is the most common approach for finding optimal decision under uncertainty in IAMs and particularly useful for investigating the implications of parametric uncertainty. We discuss NSP studies both with and without learning about uncertainty. By and large, NSP in IAMs shows that uncertainty without learning favors stronger mitigation. The climate damage uncertainty dominates the mitigation cost uncertainty. Future learning about uncertainty has only a small effect on the optimal level of mitigation unless a highly non-linear climate threshold is included in the analysis. Hence, learning serves as an argument for neither deferring nor advancing mitigation action. (ii) This last conclusion changes in a more novel approach presented in Section 3, which applies real options analysis (ROA) to IAMs. ROA highlights the value of flexibility in future actions in the face of uncertain climate change. The analysis shows that substantially stricter interim targets become economical if the value of the option to switch to a laxer target later on is taken into account. This result stems from the skewness and long upper tail in the probability distribution of avoided damages. (iii) Stochastic dynamic programming (SDP) in discrete and continuous time is the most comprehensive approach to uncertainty and discussed in Section 4. The use of SDP is practically a necessity for investigating implications of stochasticity. However, due to its intricacy, it has rarely been applied in the climate change context. Two studies show that learning about key uncertainties in the climate problem might take a long time. A first continuous time study shows that the stochasticity of climate damages has only a small effect on optimal climate policy.

This review limits its scope to the body of IAM studies on uncertainty and does not cover the entire range of literature about uncertainty and the economics of climate change. The reader is referred to the following studies for different sub-topics: Baker (2009) and citations therein discuss the theoretical literature on climate policy and uncertainty. There is a set of studies examining deep uncertainty, in which a probability distribution is not available (see e.g. Lange & Treich, 2008, for ambiguity; Lempert, 2002, for exploratory modeling; and Luo & Caselton, 1997, for Dempster-Shafer theory). As an alternative to the welfare theoretic approaches utilized by the majority of IAM analyses, some studies adopt risk management approaches (e.g. Scott et al., 1997). Also, in this paper, we only review the studies on first-best climate policy, but there is an extensive literature on the choice of policy instruments under uncertainty (see e.g. Hepburn, 2006). Finally, quantitative results obtained from intertemporal welfare analysis, including those

of all the studies reviewed in this article, are notoriously sensitive to the chosen normative parameters. We will not address this issue explicitly and refer the reader to Dasgupta (2008).

As the basis for the discussion of different approaches below, we here give a generic formulation of an IAM. We denote the utility function of the representative agent by u. This could also be a multi-regional welfare function and would then depend on the average consumption in the different regions. We denote the pure rate of time preference by δ, time-step length by Δt, the vector of state and decision variables at time t by X_t and I_t, respectively, consumption by c_t, the vector of uncertain parameters by θ, stochastic shocks by η_t and the measurement error by γ_t. The state variables include the production capital and the atmospheric carbon stock, for instance, while investments are formulated as decision variables. We omit a time-varying population in this article, which would simply lead to a time-dependent factor to the utility function. Finally, the vector of messages containing information about the uncertain parameters up to time t is denoted by $m^t = (m_1,...,m_t)$. A generic stochastic first-best IAM can then be written as

$$\begin{aligned} \max_{\{I_t(m^t)\}} \quad & \left\{ E_0 \sum_{t=0}^{\infty} e^{-\delta t \Delta t} u\big(c_t(X_t, I_t(m^t))\big) \right\} \\ \text{s.t.} \quad & X_{t+1} = f\big(X_t, I_t(m^t), \theta\big) + g\big(X_t\big)\eta_t \\ & m_t = X_t + h(X_t)\gamma_t \end{aligned} \tag{1}$$

The expectation (E_0) of utility is taken conditional on the information available at time t = 0. The first constraint in Eq. (1) specifies the system dynamics, which contains both uncertain parameters θ and a stochastic term η_t. The measurement error in the second constraint has not been considered in IAMs yet and will also be neglected in the following. What makes Eq. (1) difficult to solve is the fact that decisions I_t generally depend on the history of messages that have been received.

2 Implications of Parameter Uncertainty: Non-Recursive Stochastic Programming

Numerical modeling of climate change uncertainty is mostly based on stochastic programming, which denotes an optimization including random parameters, whether it be uncertain model parameters or stochastic shocks. We denote methods that do not use dynamic programming, which is discussed separately in Section 4, by non-recursive stochastic programming (NSP).

NSP is the simplest approach to actual optimization under uncertainty and especially useful for the investigation of parametric uncertainty. In even simpler approaches, such as sensitivity or scenario analysis, the performance of a policy in the multiplicity of possible states

of the world and the agents' risk aversion is not taken into account. We do not discuss those approaches (see Kann & Weyant, 2000).

2.1 Effect of Uncertainty on Optimal Policy

A few studies conduct NSP without taking learning about uncertainty into account. The main question of these studies is how uncertainty affects optimal policy in terms of timing and stringency. Exclusion of learning from the modeling substantially reduces the information dynamic complexity and thus allows including multiple uncertainties and detailed system dynamics.

NSP can be formulated as follows. First, a sample is drawn from the joint probability distribution of all uncertain parameters θ and shocks η_t. The sample points $(\theta_s, \eta_{t,s})$ can be called states of the world. We denote the probability of state *s* by p_s. Without learning and a with a finite time horizon *T*, problem (1) then reads as

$$\begin{aligned} \max_{\{I_t\}} \quad & E_0\left\{\sum_{t=0}^{T} e^{-\delta t \Delta t} u\big(c_t(X_t, I_t)\big)\right\} \\ \text{s.t.} \quad & X_{t+1} = f(X_t, I_t, \theta_s) + g(X_t)\eta_{t,s} \end{aligned} \tag{2}$$

Pizer (1999) shows a way to apply NSP by approximating the optimal consumption paths under climate change policy by analytical functions of the state variables. He finds the optimal policy path under uncertainty by evaluating the intetemporal welfare from the optimal consumption paths with differing policies and inclusive of uncertainty. Taking account of various forms of uncertainty, he finds that uncertainty justifies roughly 30% stricter emissions reductions.

Problem (2) can be further simplified by not actually performing a continuous optimization but only finding the most desirable policy in a given set of policies I^p. This is called policy evaluation. However, whether the resulting policy approximates the optimal choice well is not clear, particularly in models with a high-dimensional decision space. Gjerde et al. (1999), for example, use this simplification to show that a potential climate catastrophe justifies substantially stronger mitigation action. They do not separate the effect of uncertainty about this catastrophe, but it can be conjectured that it strengthens the argument, because mitigation costs are not uncertain in their study.

NSP can also be used to estimate the value of the immediate and complete resolution of uncertainty. This is done by comparing expected utility resulting from problem (2) with the expectation of utility over separate deterministic optimizations in each state of the world. Peck & Teisberg (1993) report for the CETA model that climate sensitivity and climate damages are the most useful uncertainties to learn about with values of about US$150 and 100 billion, respectively. Gjerde et al. (1999) report an even higher value of learning about potential climate

catastrophes of almost US$600 billion.

2.2 Effect of Learning on Policy Stringency

By taking learning into account, NSP can address the questions of how future learning changes optimal near-term climate policy and of how valuable future information about different uncertainties is.

Due to increasing computing power, NSP with learning has become more widely applicable to IAMs in recent years. It is still limited to a single or at most a few learning steps, though, and information arrives continuously in reality. However, information pooling and climate policy formation are slow processes. The IPCC publishes its reports every 7 years, it took five years to negotiate the Kyoto Protocol, 15 years to build consensus on the 2°C threshold as a long-term environmental target, and it may still take several years to get the major developing countries to commit to absolute emission targets. In this light, it might not be unrealistic to assume that an initial near-term climate policy up to 2030 or 2050, for instance, is revised only once or a couple of times.

For one learning step and only parametric uncertainty, we can rewrite (1) as

$$\begin{aligned} \max_{\{I_t^l\}} \quad & \sum_i q_i \sum_j p_j^i \sum_{t=0}^{T} e^{-\delta t \Delta t} u\left(c_t(X_t(\theta_j), I_t^i)\right), \\ s.t. \quad & X_{t+1}(\theta_j) = f\left(X_t, I_t^i, \theta_j\right), \\ & \forall t < t_l : I_t^l = I_t^1, \end{aligned} \tag{3}$$

where q_i and p_j^i are the probability of message *i* and the probability of state of the world *j* after receipt of message *i*, respectively. The latter is characterized by the vector of parameter values θ_j.

The last constraint in (3) ensures that decisions can only be tailored to the individual messages after receiving them at time t_l. This "trick" is sometimes called "discrete stochastic programming" and was first proposed by Cocks (1968). It allows solving recourse problems such as (3) by efficient optimization solvers in modeling systems such as AMPL and GAMS. This in turn allows using IAMs with comparably high system dynamic complexity and often without changing the modeling system. However, the number of decision variables and constraints increases exponentially for more than one learning step quickly rendering the problem unsolvable. For several learning steps, solution methods based on a recursive formulation are superior (see Section 5).

The formulation in (3) is particularly suited to consider parameter uncertainty but could in principle also be used to incorporate stochasticity. Since a stochastic process introduces a random shock for each time-step, however, a sufficient sample would render the sums in Eq. (3) unmanageable. Therefore, recursive methods are the preferred choice if stochasticity is

Figure 1:

Scheme of optimal emissions in different scenarios. The cases with learning are depicted for only two messages.

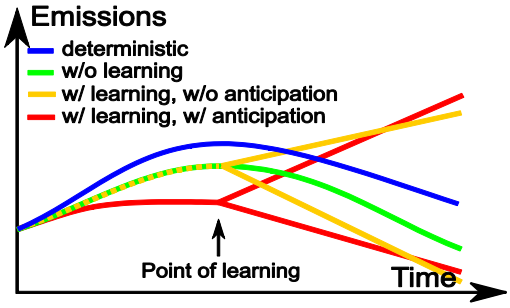

considered (see Section 5).

Considering only one learning point as in (3) simplifies not only the solution but also the interpretation of results. It makes the distinction between the four cases shown in Fig. 1 particularly intuitive: (a) The deterministic case, in which the uncertain parameters are fixed at their expected value. This is the blue line in Fig. 1. (b) The case without learning, in which the parameters are uncertain and uncertainty is not resolved. This is the green line in Fig. 1. (c) The case of non-anticipated learning, in which the uncertainty is at least partly resolved but this is not anticipated. Decisions before learning coincide with decisions without learning. These are the orange lines. (d) The case of anticipated learning, in which learning is additionally anticipated, potentially leading to different optimal pre-learning decisions. These are the red lines in Fig. 1. The key results are differences in optimal policies in these scenarios and the associated welfare differences.

More specifically, we can distinguish two effects. Firstly, static uncertainty has an effect on optimal emissions and associated welfare as compared to the deterministic scenario. This is the difference between the blue and the green line in Fig. 1. This effect stems from the non-linearity of the objective function in the uncertain parameters and is also investigated in uncertainty propagation. It is generally found to be small in studies using NSP (Webster, 2000; Webster et al., 2008; O'Neill & Sanderson, 2008; Lorenz et al., 2011). Uncertainty propagation has shown that uncertainty can have a substantial effect on optimal emissions. This indicates that the smallness of the effect of uncertainty in DSP studies is likely to be at least partly due to a crude representation of uncertainty.

Secondly, learning has an effect on optimal emissions as compared to the no-learning case. This is the difference between the green and the red lines. The associated welfare increase is called the expected value of information (EVOI). The EVOI is generally found to be significant. Particularly learning about climate damages and climate sensitivity are found to be very valuable compared to current research budgets. This was shown in different IAMs by Nordhaus & Popp (1997), and Lorenz et al. (2011) amongst others.

The effect of learning can be decomposed into two parts. Firstly, optimal policy after learning will depend on what is learned. This is the difference between the orange lines and the

green line. The associated welfare difference can be called an option premium, which is further discussed in Section 4.

Secondly, anticipation of future learning changes optimal near-term climate policy before learning. This is the difference between the orange and the red lines. The associated welfare increase can be called the expected value of anticipation (EVOA). Anticipation of learning is valuable if decisions are irreversible and anticipation generates flexibility. There are two main irreversibilities involved in the climate problem that counteract each other. Investments in mitigation, which are at least partly sunk, and emissions stay in the atmosphere for decades to centuries. As a result, most studies performing cost-benefit analysis find that anticipation of learning has only a small effect on optimal emissions (Ulph & Ulph, 1997; Webster, 2000; Webster et al., 2008; O'Neill & Sanderson, 2008; and Lorenz et al., 2011).

However, a substantial effect of anticipation on optimal near-term policy was shown in the presence of an irreversible climate threshold with uncertain corresponding damages (Keller et al., 2004; Dumas & Ha-Duong, 2004; and Lorenz et al., 2011). A stricter policy can then be justified because it keeps the option open to avoid the threshold if it is learned to be severe.

A strong effect of anticipation is also found in studies performing cost-effectiveness analysis in general find lower optimal emissions with uncertainty and learning (parts of Webster et al., 2008; Bosetti et al., 2009). But Schmidt et al. (2011) argue that these latter results should be taken with caution because they stem from a disputable interpretation of climate targets under uncertainty as strict targets that have to be met with certainty.

Since learning is exogenous in NSP, it is not suitable for studying the risks of geo-engineering or learning about technologies, for instance. What will be learned on geo-engineering strongly depends on whether and to what extent it is applied. Similarly uncertainty about the floor costs and learning rates of technologies is mostly reduced by applying them. An endogenous dependence of what is learned at a fixed point in time could, in principle, be included in DSP. But it is unclear whether the resulting most likely non-convex model can still be solved. Besides, the authors don't see how the time of learning could be endogenized. Another consequence of the exogeneity of learning is that NSP is not suited to answer the question of what can be expected to be learned. It uses the answer as an input.

3 Value of Flexibility: Real Options Analysis

Recently, Anda et al. (2009) have presented a way to apply real options analysis (ROA) in integrated assessment models. ROA uses methods from financial option pricing, in particular contingent claims analysis, to value the managerial flexibility inherent in real investment decisions. It sometimes also uses stochastic dynamic programming, which we discuss separately in Section 4. ROA also provides a language and intuition for talking about this flexibility.

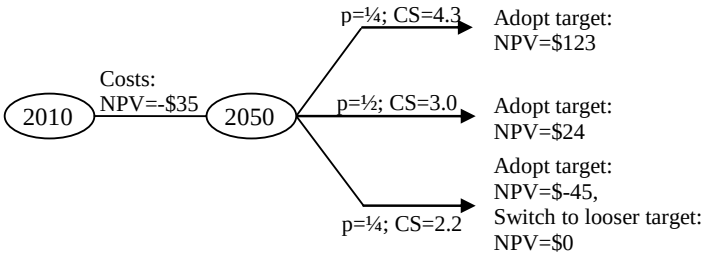

Figure 2 Simple demonstration of the option value of an interim target. CS = 4.3 with probability ¼, CS = 3.0 with prob-ability ½ , and CS = 2.2 with probability ¼ .

The simple example depicted in Fig. 2 demonstrates the option value concept. There is uncertainty about climate sensitivity (CS), which can only take one of three possible values. The true value is learned in 2050. The mitigation costs of an interim target up to 2050 are $35 trillion. Benefits in the form of avoided climate damages accrue only after 2050 and depend on the true value of climate sensitivity. Adoption of the long-term target whatever the value of CS has a negative NPV of -$35 + ¼ $123 + ½ $24+ ¼ (-$45) = -$3.5 trillion. This assumes risk-neutrality for simplicity and would have to be risk-adjusted under risk aversion. If we add the option to abandon the target for a looser one with zero NPV if the post-learning NPV of the target is negative, we get an expanded NPV of -$35 + ¼ $123 + ½ $24+ ¼ (-$0) = $7.75 trillion. Thus, the option premium is $11.25 trillion. Due to this premium, it is economical to adopt the target as an interim target but not as a one-shot, long-term target.

The interim target can be seen as a European option on the long-term target. It is the option but not the obligation to adopt the long-term target at a given future time. The time of expiration is 2050. We say the option is executed if the target is abandoned (put option). The strike price is then the NPV of the best alternative, which is assumed to be zero for simplicity.

The value of an option depends on the volatility and price of the underlying asset. For financial options and in standard ROA these characteristics can be observed in markets, or at least the characteristics of a highly correlated twin security. There are, of course, no markets for long-term climate targets, though, but one can derive the characteristics from an IAM.

This is done by a Monte-Carlo simulation of the policy under consideration denoted by I^p. The NPV of the target right after learning the true state of the world θ_s at t_l is obtained by discounting net benefits over business-as-usual ($I^{BA'U}$) at the ex post risk-free rate $r_{s,t}$,

$$NPV^{ep}(t_l) = \sum_{t=t_l}^{T} e^{-r_{s,t}t}\left(c_t(\theta_s, I^p) - c_t(\theta_s, I^{BAU})\right) \tag{4}$$

This is the random price of the underlying asset right after learning. The NPV of the target right before learning is obtained by discounting certainty equivalent benefits at the ex ante risk-free rate r_t

$$NPV^{ea}(t_l) = \sum_{t=t_l}^{T} e^{-r_t t}\left(\bar{c}_t(I^p) - \bar{c}_t(I^{BAU})\right) \tag{5}$$

This is the spot price of the target right before learning. The NPV of the costs of the interim target is given by

$$NPV^{costs}(t=0) = -\sum_{t=0}^{t_l} e^{-r_t t}\left(\bar{c}_t(I^p) - \bar{c}_t(I^{BAU})\right) \tag{6}$$

If the probability distribution of benefits (Eq. 4) can be approximated by a convenient distribution function, one can use analytical option pricing formulas. Most option pricing formulas, including Black-Scholes's, imply non-negative prices. The price of a long-term target, however, can be negative. Therefore, Anda et al. (2009) apply Bachelier's model, in which prices are normally distributed and possibly negative. If we denote the strike price by K, and the volatility of the asset, i.e. the standard deviation of NPV^{ep}, by σ, Bachelier's formula for the value of the option right before learning reads

$$V^B = (NPV^{ea} - K)\Phi\left(\frac{NPV^{ea} - K}{NPV^{ea}\sigma}\right) + NPV^{ea}\phi\left(\frac{NPV^{ea} - K}{NPV^{ea}\sigma}\right), \tag{7}$$

where Φ and ϕ are the cumulative distribution function and the probability density function of the standard normal distribution, respectively, and where we have omitted the time argument of the *NPV*. Since we consider an instantaneous resolution of uncertainty, neither the time to expiration nor the interest rate occur in the pricing formula. The interim target should be adopted if

$$NPV^{costs} + V^B > 0. \tag{8}$$

Due to a long upper tail in the probability distribution of climate damages, a normal distribution for the benefits of climate policy is not realistic and Bachelier's model underestimates the option value of an interim target. Anda et al. (2009) show that more sophisticated pricing models taking higher moments of the distribution into account lead to substantially higher option values.

ROA has two major advantages over DSP. Firstly, it does not demand a stochastic optimization but only a Monte Carlo simulation. Secondly, it allows consideration of continuous distribution functions with tails in analytical option pricing formulas. Besides, it provides a clear intuition and quantification of flexibility as an option value.

However, in the method summarized above, the target and the decisions to reach were only evaluated and not derived in an optimization. Thus there is no way to tell whether the decisions are efficient and the target optimal. See Anda et al. (2009) for conditions under which the method can be extended to an optimization. A further limitation of the method above is that it can only consider a perfect one-time learning. Partial learning and multiple learning steps

cannot be considered. Besides, learning is exogenous, so that it shares the corresponding limitations of DSP.

Up to now, ROA has only been applied in the IAM DICE by Anda et al. (2009). Their main finding is that an upper tail in the avoided damage distribution leads to a large option value and thus justifies an aggressive interim target even without risk aversion. There is a closely related theoretical literature on the quasi-option value in environmental economics (Arrow & Fisher, 1974; Henry, 1974, see Aslaksen & Synnestvedt, 2004, for the relation to ROA).

4 Implications of Stochasticity: Stochastic Dynamic Programming

The approaches we have discussed up to now are not suitable for taking stochasticity and the repeated and endogenous updating of probability distributions into account. Stochastic dynamic programming (SDP) is preferred for the examination of such high information dynamic complexity. While most current debates on uncertainty in climate change deal with parametric uncertainty, many aggregate processes in the climate system and the economy are also stochastic. Modeling of stochasticity can answer some interesting research questions. To what extent does stochasticity of the climate system hinder the resolution of parametric uncertainty? How does it change optimal decision? We will first briefly discuss SDP in discrete time and subsequently in continuous time.

4.1 Discrete Time Modeling

The value function is defined as the maximum utility that can be obtained given the current state of the system including the probability distributions on the uncertain parameters. In discrete time this reads as

$$\begin{aligned} J(X_0) = \max_{\{I_t(m^t)\}} \quad & E_0 \sum_{t=0}^{\infty} e^{-\delta\, t\, \Delta t} u\big(c_t(X_t, I_t(m^t))\big), \\ \text{s.t.} \quad & X_{t+1} = f(X_t, I_t(m^t), \theta) + g(X_t)\eta_t, \end{aligned} \tag{9}$$

where have already presumed that the time horizon is infinite and the value function does not depend on time explicitly. See Kelly & Kolstad (1999) for how to achieve this, if some parameters vary exogenously over time. Using the value function and the principle of optimality, we can rewrite the first line of problem (9) recursively as

$$J(X_t) = \max_I \left\{ u\big(c_t(X_t, I\big) + e^{-\delta \Delta t} E_t J(X_{t+1}) \right\}, \tag{10}$$

where we omitted the system and information dynamics. The mathematical conditions for which the principle of optimality holds can be found e.g. in Stokey & Lucas (1989).

A system of functional equations for I and J can be obtained if the first order conditions

of the right-hand side maximization are sufficient for optimality,

$$\begin{aligned} \frac{d}{dI}\left(u\left(c_t(X_t,I)+e^{-\delta\Delta t}E_t J(X_{t+1})\right)\right) &= 0 \\ J(X_t) &= u\left(c_t(X_t,I)+e^{-\delta\Delta t}J(X_{t+1})\right) \\ X_{t+1} &= f(X_t,I_t,\theta)+g(X_t)\eta_t \end{aligned} \tag{11}$$

For simple models these equations can be solved or exploited analytically. Karp & Zhang (2004) use a linear-quadratic multi-period IAM. They find that anticipation of learning about climate damages decreases optimal abatement by about 10-20%.

In more complex models, the value function has to be analyzed numerically. The two main algorithms are value-iteration and policy-iteration, of which only the former has been applied in IAMs. They are both based on Eq. (10). Value-iteration exploits the fact that the right hand side of Eq. (10) is a contraction mapping on J due to the discounting (see Bertsekas, 2005, for details). One starts with a guess of the value function, then maximizes the right-hand side of Eq. (10) and obtains a new guess for the value function until the algorithm converges. Thereby, the value function is parameterized. Kelly & Kolstad (1999) and Leach (2007) use neural networks. This is only manageable for a low dimensionality of the state space. Since the state space includes the probability distributions of the uncertain parameters, these have to be describable by few parameters. Kelly & Kolstad (1999) and Leach (2007) assume normality.

The main advantage of SDP is that it allows taking endogenous and repeated updating of uncertainty into account. Kelly & Kolstad (1999) are mainly interested in how learning about climate sensitivity depends on emissions. They explicitly model the stochasticity of the temperature process and the Bayesian updating on climate sensitivity in DICE. They find that learning the true value of climate sensitivity takes at least 90 years. They also show a trade-off between emissions control and the speed of learning. Leach (2007) extends the analysis to two uncertain parameters in the temperature process and shows that this can delay learning by hundreds or even thousands of years.

4.2 Continuous Time Modeling

A precise foundation of stochastic models and SDP methods in continuous time demands a higher degree of mathematical sophistication than in discrete time. We refer to Chang (2004) for an accessible introduction. At the same time, continuous time SDP provides convenient tools.

These tools are mostly contingent on the assumption that uncertainty can be described by an Ito stochastic process. Accordingly, they are not suited to take parameter uncertainty into account. Most of the discussion on uncertainty in climate change has focused on parametric uncertainty so far. However, many aggregate processes in the climate system and the economy are also stochastic. This fact leads to the question of what such stochasticity implies for optimal climate policy. Continuous time SDP is a complete method for the investigation of this question.

We now briefly outline continuous time SDP. The information about the system is summarized by its current state X_t. The value function is then defined analogous to Eq. (9) as

$$\begin{aligned} J(X_0) = \max_{\{I(X)\}} \quad & E_0 \int_0^\infty e^{-\delta t} u\big(c(X_t, I(X_t))\big) dt, \\ \text{s.t.} \quad & dX_t(I) = f\big(X_t, I(X_t)\big) dt + g\big(X_t, I(X_t)\big) dB, \end{aligned} \tag{12}$$

where dB is a vector of increments of independent standard Wiener process. As in Section 5, it is presumed for simplicity that the value function does not depend on time explicitly, and that the time horizon is infinite. Note that the Ito process in the second line might be multi-dimensional, in which case g is the co-variance matrix. Some state variables might also be deterministic.

The principle of optimality can be written analogous to Eq. (10) as

$$J(X_t) = \max_I \left\{ u\big(c(X_t, I)\big) dt + e^{-\delta dt} E\, J(X_t + dX_t(I) \right\}, \tag{13}$$

Substituting the system dynamics and using Ito's calculus one obtains the (autonomous) infinite horizon Hamilton-Jacobi-Bellman (HJB) equation. For simplicity, we only specify it for a one-dimensional state space,

$$\begin{aligned} \max_{\{I(X_t)\}} \Bigg\{ & u\big(c(X_t, I(X_t))\big) - \delta J(X_t) + f\big(X_t, I(X_t)\big) \frac{d}{dX} J(X_t) + \\ & + \frac{1}{2} g\big(X_t, I(X_t)\big) \frac{d^2}{dX^2} J(X_t) \Bigg\} = 0 \end{aligned} \tag{14}$$

For very simple and specific models, the HJB equation can be solved analytically. This is the main advantage of continuous time SDP and exploited in the climate context e.g. in Pindyck (2000). The solutions normally take the form of domains in a state space where specific discrete decisions are optimal (stopping rules; see also Dixit and Pindyck, 1994). However, due to the simplicity of the models, these results mainly serve to build an intuition and have limited policy relevance.

Recently, Lontzek & Narita (2009) have applied continuous time SDP to a somewhat more complex IAM for the first time. Similar to Pindyck (2000), they investigate the effect of climate damage stochasticity on optimal climate policy. They first derive the control variables as a function of the state variables analytically from the first order conditions and subsequently use Chebychev collocation as proposed by Judd (1998) to solve the HJB equation. They show that stochasticity has only a small and ambiguous effect on optimal emissions reductions as compared to the deterministic case (without shocks). The sign of the effect is found to depend mainly on the level of accumulated production capital.

5 Summary and Conclusions

We have reviewed probabilistic approaches to uncertainty in integrated assessment models and their respective implications for climate policy.

Non-recursive stochastic programming (NSP) is the simplest way to take uncertainty and learning into account in IAMs. Uncertainty generally, and not surprisingly, justifies stronger emissions reductions. Estimates of the extent, however, vary from very little up to 30% depending on how many uncertain parameters and sample points are considered. Future learning is generally found not to be a significant factor to promote more or less mitigation unless potential climate thresholds are taken into account. However, learning can have an impact on the efficient mitigation portfolio, and the optimal level of R&D in particular.

We have then discussed a way to apply real options analysis (ROA) to IAMs. It is characterized by the use of financial option pricing methods to value the option of adjusting policy to future learning. It allows a more comprehensive consideration of uncertainties than discrete stochastic programming, and the representation of tails in particular. It shows that future learning can then be an argument for substantially stronger short-term emissions reductions. Up to now, it has only been used to evaluate given policies. Its application with a direct optimization might be a promising extension.

Whenever stochasticity is taken into account, possibly in conjunction with the endogenous resolution of uncertainty, dynamic programming is the preferred, or rather mandatory, choice. We have briefly described the most common discrete and continuous time methods. Stochastic dynamic programming (SDP) has been used in an IAM to show that learning about climate uncertainty may take a very long time up to thousands of years.

The most important general policy implication from the literature is that despite a wide variety of analytical approaches addressing different types of climate change uncertainty, none of those studies supports the argument that no action against climate change should be taken until uncertainty is resolved. On the contrary, uncertainty despite its resolution in the future is often found to favor a stricter policy.

There are a number of future research needs concerning first-best climate policy under uncertainty. (i) A better representation of tails in the probability distributions of uncertain parameters in IAMs will be necessary to settle the discussion that has emanated from Weitzman (2009). (ii) A better representation of irreversibilities in the climate system including tipping points and the inertia in the economy will be necessary to settle the discussion on the optimal stringency of near-term policy in the face of future learning. (iii) Learning about some uncertainties is endogenous. Risks of geoengineering options will be fully known only after they are applied. The maximum efficiency of various renewable energy technologies will be learned only if the technologies are applied on a large scale (see also Baker & Shittu, 2008).

Modeling endogenous learning demands stochastic dynamic programming, in which the inclusion of sufficient climatological or technological detail poses a great challenge. In addition, endogenous technical change generates non-convexities in the optimization problem, which demand global optimization solvers. (iv) What are the implication of uncertainty and learning for first-best climate policy in developing countries? Significant short-term policy of emission control might steer developing countries into low-carbon economic growth and prevent a lock-in to carbon-intensive production capital. The associated benefits could be estimated by discrete stochastic programming or real options analysis. (v) The question of how alternative preferences, such as habit formation, direct utility from an environmental good, distinction between risk aversion and intertemporal elasticity of substitution and others change optimal policy under uncertainty has not yet received sufficient attention but should be explored. (vi) The analysis of the persistent stochasticity both of the climate system and the economy is still in an initial stage, and investigations of its implications for climate policy in more complex IAMs are needed. (vii) Finally, there is a strong need for reliable probability estimates for the key parameters of IAMs, especially the climate change damage parameters.

References

Anda, J., A. Golub, and E. Strukova, 2009. *Economics of Climate Change under Uncertainty: Benefits of Flexibility*, Energy Policy 37: 1345-1355.

Arrow, K.J., and A.C. Fisher, 1974. *Environmental Preservation, Uncertainty, and Irreversibility*, Quarterly Journal of Economics 88(2): 312-319.

Aslaksen, I., and T. Synnestvedt, 2004. *Are the Dixit-Pindyck and the Arrow-Fisher-Henry-Hanemann Option Values Equivalent?* Statistics Norway, Research Department Discussion Papers No. 390

Baker, E. 2009. *Optimal Policy under Uncertainty and Learning about Climate Change: A Stocahstic Dominance Approach.* Journal of Public Economic Theory 11 (5): 721-747.

Baker, Erin, and Ekundayo Shittu. 2008. *Uncertainty and endogenous technical change in climate policy models.* Energy Economics 30 (6): 2817-2828.

Bertsekas, Dimitri P. 2005. *Dynamic Programming & Optimal Control*, 3rd ed. Athena Scientific.

Bosetti, V., Carraro, C., Sgobbi, A., and Tavoni, M., 2009. *Delayed action and uncertain stabilization targets: How much will climate policy cost?* Climatic Change 96:299–312.

Chang, F.-R., 2004. *Stochastic Optimization in Continuous Time*. Cambridge: Cambridge University Press.

Cocks, K. D. 1968. *Discrete Stochastic Programming*. Management Science 15 (1): 72-79.

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

Dixit, A.K., R.S. Pindyck, 1994. *Investment under Uncertainty*. Princeton University Press.

Dumas, P., M. Ha-Duong, 2004. *An abrupt stochastic damage function to analyse climate policy benefits*. In Alain Haurie and Laurent Viguier (eds.), The coupling of climate and economic dynamics, Essays on Integrated Assessment, Kluwer.

Epstein, Larry G. 1980. *Decision Making and the Temporal Resolution of Uncertainty.* International Economic Review 21 (2): 269-83.

Gjerde, J., S. Grepperud and S. Kverndokk, 1999. *Optimal climate policy under the possibility of a catastrophe*, Resource and Energy Economics 21 (3-4): 289-317.

Heal, G., and B. Kristrom, 2002. *Uncertainty and climate change*. Environment and Resource Economics 22: 3–39.

Henry, C., 1974. *Investment Decisions Under Uncertainty: The "Irreversibility Effect,"* American Economic Review 64(6): 1006-1012.

Hepburn, Cameron. 2006. *Regulation by Prices, Quantities, or Both: A Review of Instrument Choice.* Oxford Review of Economic Policy 22 (2): 226 -247.

Judd, Kenneth L. 1998. *Numerical Methods in Economics.* The MIT Press.

Kann, A., and J.P. Weyant, 2000. *Approaches for performing uncertainty analysis in large-scale energy/economic policy models*, Environmental Modeling and Assessment 5: 171-197.

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

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

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

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

Leach, A. 2007. *The climate change learning curve*. Journal of Economic Dynamics and Control, 31(5):1728–1752.

Lempert, R. J. 2002. *A new decision sciences for complex systems*. Proceedings of the National Academy of Sciences 99 (5): 7309-7313.

Lontzek, T.S. and D. Narita, 2009. *The Effect of Uncertainty on Decision Making about Climate Change Mitigation. A Numerical Approach of Stochastic Control*, Kiel Working Paper, 1539.

Lorenz, A., M.G.W. Schmidt, E. Kriegler, H. Held, 2011. *Anticipating Climate Threshold Damages*, accepted for publication in Environmental Modeling and Assessment.

Luo, WB, and B Caselton. 1997. *Using Dempster-Shafer theory to represent climate change uncertainties.* Journal of Environmental Management 49 (1): 73-93.

Nordhaus, W.D., D. Popp, 1997. *What is the Value of Scientific Knowledge? An Application to Global Warming Using the PRICE Model*. The Energy Journal, 18(1): 1-45

O'Neill, B.C., W. Sanderson, 2008. *Population, uncertainty, and learning in climate change decision analysis*. Climatic Change, 89(1-2): 87-123.

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

Peterson, S., 2006. *Uncertainty and Economic Analysis of Climate Change: A Survey of Approaches and Findings.* Environmental Modeling and Assessment, 11(1): 1-17.

Pindyck, R.S., 2000. *Irreversibilities and the Timing of Environmental Policy*, Resource and Energy Economics 22: 233-259.

Pizer, W.A., 1999. *The optimal choice of climate change policy in the presence of uncertainty*. Resource and Energy Economics 21: 255–287.

Schmidt, M.G.W., A. Lorenz, H. Held, and E. Kriegler. 2011. Climate Targets under Uncertainty: Challenges and Remedies. Climatic Change: Letters 104 (3): 783-791.

Scott, M.J., R.D. Sands, J. Edmonds, A.M. Liebetrau, and D.W. Engel, 1999. *Uncertainty in integrated assessment models: modeling with MiniCAM 1.0*, Energy Policy 27: 855-879.

Stokey, N.L., R.E. Lucas Jr, and E.C. Prescott. 1989. Recursive Methods in Economic Dynamics. Harvard University Press.

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

Webster, M.D., 2000. The *Curious Role of "Learning" in Climate Policy: Should We Wait for More Data?* MIT Joint Program on the Science and Policy of Global Change, Report No. 67.

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

Chapter 6

Conclusions

This thesis set out to contribute to the identification of climate policies that do justice to the pervasiveness of uncertainty in climate change. This chapter summarizes its main contributions, draws some general conclusions and points out remaining research needs.

In Section 1.3 of the introduction, we have discussed the main policy questions regarding uncertainty in climate change. Thereby, we have focused on socially optimal climate policy, and, for the most part, did not take market failures into account. The main questions were: What are the implications of uncertainty for optimal policy? What are the implications of learning about uncertainty for optimal policy? A crucial underlying question is: What is a suitable decision criterion for climate change that accommodates uncertainty, learning, and equity? The contributions of this thesis to each of these three questions as well as respective future research needs are summarized in Sections 6.1 to 6.3. A few final remarks are provided in Section 6.4.

6.1 Uncertainty and Climate Policy

In Section 1.3 of the Introduction we have highlighted that it is helpful to separate the effect of uncertainty and the effect of learning on optimal climate policy. Learning substantially increases the complexity of the analysis and should therefore only be taken into account if necessary. We draw conclusions for uncertainty and for learning in this and the next section, respectively.

Many integrated assessments of climate change use the concept of a representative agent. In Chapter 2 we have pointed out that this implicitly presumes efficient risk sharing. Since currently only a fraction of catastrophic damages is insured and large part of climate impacts will be of the form of catastrophes, we have relaxed the assumption of a representative agent and considered heterogeneous climate damages.

In a first step, we then asked for an optimal climate policy if neither insurance, i.e. the transfer of risk between individuals, nor self-insurance, i.e. the adjustment of savings to heterogeneous climate damages, are available. We found that uncertain and strongly heterogeneous climate damages are an argument for substantially stricter climate policy.

This result, however, depended on the form of the social welfare function. We separated risk aversion, i.e. the aversion of individuals against uncertainty in consumption, from inequality aversion, i.e. the aversion of society against differences in consumption between different individuals. The argument for stricter climate policy is only valid if both types of aversion are present. The intuition for this result is as follows: Firstly, heterogeneity of climate damages concentrates the aggregate risk associated with damages on fewer people. This increases the risk premium of these individuals more than proportionately. The same risk borne by fewer people leads to a larger risk premium. Secondly, the heterogeneity of damages and the associated risk premium lead to inequality between individuals, which lowers welfare under inequality aversion. Only the compound effect justifies substantially stricter climate policy.

In a second step, we introduced insurance markets. Since we did not model the market failures that hinder catastrophe insurance, this mainly served to determine the theoretical potential of insurance and insurance markets. This potential was indeed found to be significant. Perfect insurance spreads the risk over the entire population thus lowering the risk premium and resulting inequality. It essentially eliminates the effect of damage heterogeneity and allows to relax climate policy. This is commensurate with the results from representative agent models, which show that uncertainty has only a minor effect on optimal policy unless the tails of the probability distributions are taken into account.

In a third step, we gave individuals the opportunity to self-insure against their heterogeneous damages. This turned out to be particularly effective for lax stabilization targets, because it allowed affected individuals to shift consumption from the short-term, where mitigation costs are low for these targets, to the long term, where damages are high.

We have also discussed qualitatively how these results will change if it is not known who bears what share of the climate damage risk. It will increase the effectiveness of insurance but decrease the effectiveness of self-insurance. It will not affect the results without insurance and self-insurance. We also shortly discussed income inequality and the interplay between income inequality and damage heterogeneity. First of all, current global income inequality is so large that it should be the primary concern for a utilitarian. Under the assumption of constant relative risk aversion and inequality aversion, income inequality argues for stricter climate policy if relative damages are negatively correlated with income, i.e. tend to be high for individuals with low income.

The analysis of Chapter 2 could be extended in various directions. First of all, the analysis will become more meaningful once detailed estimates of the global distribution of climate damages and risk are available. Secondly, our analysis did not take into account that damages will contain some variability over time and also accrue to different individuals in different years. It would be interesting to study the effect of this variability on optimal targets as well as on the potential of insurance and self-insurance. Thirdly, in the light of the large theoretical potential of insurance, it seems to be important to ask how this potential can be tapped to some extent despite the profound market failures and long time horizons involved.

Introducing agent heterogeneity and imperfect risk sharing amongst individuals has previously been used to explain the so called "equity premium puzzle". The equity premium puzzle (Mehra & Prescott, 1985; see Mehra, 2007, for a review) states that unrealistically high values of risk aversion are necessary to explain observed equity premia in representative agent models. The volatility of macro-consumption of about 2%/yr is simply too small to explain an equity premium of about 6%/yr. In reaction to this puzzle, two strands of literature have evolved. One strand argues that there isn't actually a puzzle, because the actual equity premium after controlling for selection bias, for instance, isn't that high. The second strand has developed a bunch of models and explanations for a high equity premium. Imperfect risk sharing as discussed in the climate change context in Chapter 2 is one of them (Campbell & Mankiw, 1990). These explanations might be highly relevant for climate policy and the risk premium associated with climate damages. Weitzman (2007) explains the premium by fat-tailed uncertainty about returns. This explanation has stirred a vivid and important discussion in the climate change context as well (Weitzman, 2009). Constantinides (1990) proposes habit formation as an explanation. See Ikefuji (2008) for a theoretical discussion in the environmental economics context. Abel (1990) uses the so called "catching up with the Joneses" effect, and Weil (1989) the separation of risk aversion from the elasticity of inter-temporal substitution. The validity of these explanations and their quantitative consequences for optimal climate policy under uncertainty are far from clear and constitute an interesting topic for future research.

6.2 Learning and Climate Policy

Learning and the resulting desire to remain flexible until uncertainty is reduced are crucial parts of the climate change problem. As highlighted in Sections 1.3 and 4.1, however, a number of integrated assessment studies have found that it has only a minor effect on optimal near-term climate policy. Stricter abatement prevents more warming and hence provides flexibility but at the same time demands sunk investments that reduce flexibility. An exception is learning about thresholds, or tipping elements, in the climate system, such as a shutdown of the Atlantic thermohaline circulation. In Chapter 4 we performed a detailed analysis of optimal climate policy under uncertainty and learning about the damages resulting from crossing such a threshold.

We started out by introducing some terminology including a novel concept, which we called the expected value of anticipation. The idea is to decompose the overall benefits from learning into the ones that stem from optimal anticipatory changes of decisions before learning and the ones that stem from adjusting decisions to what is learned. These two components contribute to answering different questions. The value of anticipation is crucial for deciding whether future learning has to be anticipated and incorporated into near-term decision making. The expected value of information, as its name indicates, is useful to identify uncertainties whose reduction is most valuable.

In accordance with the literature, we then confirmed in the integrated assessment model

MIND that learning about climate sensitivity and climate damages is valuable but does not demand substantial changes in near-term policy. The expected value of information is large but the expected value of anticipation is not.

When we introduced a climate threshold with uncertain resulting damages, we found that anticipation of learning is important if learning takes place in a narrow "anticipation window" in time. Inside this window, almost the entire benefits from learning about the threshold damages stem from anticipatory changes of decisions. More specifically, inside this window it is optimal to reduce emissions more aggressively to keep the option open to avoid the threshold if it turns out to be severe. Anticipation is not necessary if learning is expected to take place outside the window. The entire benefits from learning can then be reaped by solely adjusting decisions after learning.

We also showed that the location of the anticipation window and its extent are very sensitive to the flexibility with which emissions can be reduced. A lower flexibility, e.g. due to technological or political constraints, broadens the anticipation window and shifts it towards the present. It also lowers the value of information because it prevents an optimal adjustment of decisions.

The analysis in Chapter 4 could be extended in various directions. First, multiple thresholds could be considered simultaneously and in combination with uncertainty about the climate system. However, this kind of analysis is computationally very intensive and might not provide substantial new insights. Second, better estimates of the uncertainty and especially the reduction of uncertainty about climate threshold damages would be necessary to decide whether we are actually in an anticipation window. Third, what and when was learned about the threshold was exogenous in Chapter 4 and could be endogenized. The closer one gets to the threshold, the better it will be known. Solving a model with endogenous learning demands dynamic programming, though, which is only manageable for very simple models, at least up to now (see Chapter 5).

Endogenous learning about uncertainty is of more general interest. It certainly plays a crucial role in learning about the costs of different mitigation technologies. The extent and trend of cost reductions will only become known by applying these technologies. This is one aspect of the interesting portfolio problem that mitigation constitutes. There is uncertainty about the costs of different technologies, which is an argument for diversification, i.e. spreading investments over multiple technologies fulfilling the same purpose. At the same time many technologies show increasing returns to scale mainly due to learning-by-doing but potentially also due to scale effects in manufacturing. This is an argument for focusing on a single technology. Furthermore, we have the above mentioned argument that uncertainty is resolved by investing, which is an argument for at least trying out multiple technologies, but possibly only one at a time. How these three aspects play out in determining optimal mitigation portfolios is an interesting question for future research.

Another interesting question in this context is how the anticipation of future learning may be an argument to limit the lock-in of developing countries in carbon intensive technologies. Even if developing countries are not supposed to reduce emissions today, there is

a chance they will reduce emissions in the future. This is uncertain, however, and will only become clear over the coming years to decades. Should this be taken into account in current investment decisions? Since investments have an impact on bargaining power, though, this question should probably be analyzed jointly with the international negotiations under uncertainty (see e.g. Kolstad & Ulph, 2008).

A further question that needs more research is how learning affects the implications of tails in the probability distributions of uncertain parameters. Weitzman (2009) argues that "fat" tails, for which probability decreases slowly for very high values of climate damages, cannot be "slimmed" through learning. However, learning might still be very valuable and substantially change near-term policy. Studying this, however, demands innovative approaches that can combine tails and learning, such as the variation of real options analysis discussed in Chapter 5.

6.3 A Decision Criterion for Climate Policy

To obtain the results summarized above, we used expected welfare maximization, often simply called cost-benefit analysis, as decision criterion. This is the most widely used decision criterion in economics. However, it has numerous critics in the climate change context. A first criticism concerns the use of unique probabilities in the face of deep uncertainty particularly about climate damages (see Section 1.3). Methods for taking deep uncertainty into account are rapidly advancing but still too involved to allow a satisfactory application to climate change. This is a very promising and demanding area for future research. A second criticism of cost-benefit analysis concerns the monetization and aggregation of all climate damages including loss of life and biodiversity amongst others. It was shortly discussed in Section 1.2.

As a result of these criticisms, an increasing number of studies have limited themselves to finding cost-effective mitigation strategies that achieve politically given climate targets such as the 2°C target. Chapter 3 has shown that this implies major conceptual problems if uncertainty about global warming is taken into account.

The chapter started from the observation that climate targets should not be interpreted as strict targets that are to be met with certainty once uncertainty is taken into account. A strict interpretation of the 2°C target, for instance, would imply excessively high mitigation costs or would even be impossible to fulfill. Climate targets should be met with a certain probability instead. Chapter 3 then showed that such probabilistic targets have normatively undesirable properties if learning about uncertainty is taken into account. This part of the argument was made both via a simple and intuitive example and by resorting to results from the decision theoretic literature.

The first undesirable property was the possibility of a negative expected value of information. Better information about the climate system could be undesirable and thus be rejected. This should arguably not be possible for a rational decision criterion. The reason was that learning changes the requirements imposed by the probabilistic target without

taking the associated benefits into account.

The second undesirable property was that probabilistic targets can become infeasible due to learning. This is the case if there are values of climate sensitivity, for instance, that cannot be excluded based on current information and for which the target threshold cannot be avoided. The target will be infeasible as soon as one of those values is learned to be the true one. It is then unclear how to perform cost-effectiveness analysis before learning, as there is no (contingent) strategy available that meets the given target no matter what is learned.

In consequence of the conceptual problems of cost-effectiveness analysis, we proposed an alternative decision criterion that allows for a trade-off between mitigation costs and the probability of crossing a given target threshold instead of limiting this probability. We called this criterion "cost-risk analysis". It avoids the conceptual problems of cost-effectiveness analysis but remains to some extent based on given climate targets. Whether cost-risk analysis describes the preferences of actual decision makers, how the parameters of the trade-off should be chosen, and what the resulting policy implications are remains for future work.

Both cost-effectiveness analysis and cost-risk analysis avoid a detailed and explicit trade-off between mitigation costs and climate damages. The target and the simple trade-off between mitigation costs and target compliance probability, respectively, are presumed and not derived. In order to derive targets in a formal analysis and to make the underlying normative assumptions explicit, the problems of cost-benefit analysis mentioned above will have to be addressed. In Chapters 2 and 4, however, we already used cost-benefit analysis, because we were explicitly interested in the trade-off between damages and mitigation costs. The results obtained there are conceptual to the same extent as the probabilistic, aggregate and monetized description of climate damages.

How should a decision criterion and normative parameters be chosen? Can the criterion be derived from the observation of markets or experiments as advocated amongst others by Nordhaus (2007)? Can normative parameters thus be uncertain (e.g. Pizer, 1999)? Or should the decision criterion emerge, at least in part, from an ethical discussion as advocated amongst others by Stern (2007)? These are crucial questions not only for climate policy but for public policy more generally. The author of this thesis clearly favors the latter position: If there is a strong moral argument for a certain parameter choice to which most people whould agree when asked then it should be used even if the resulting decision criterion does not do a good job in describing actual behavior. A good example is the choice of a very small pure rate of time preference.

6.4 Final Remarks

Science has made great progress in identifying and quantifying the uncertainties surrounding climate change. Economics has picked up the challenge and laid out the main implications for climate policy. Uncertainty is generally found to be an argument for stronger emissions reductions but opinions still notedly differ on how much. Chapter 2 has shown

that social preferences concerning inequality and the heterogeneous distribution of climate damages play a crucial role in this. Future learning about uncertainty is generally found to have little impact on optimal policy and is thus not an argument for waiting for better information before reducing emissions. Chapter 4 has shown that learning about tipping-elements in the climate system can even be an argument for stronger near-term emissions reductions. All these results crucially depend on the chosen decision criterion at which point Chapter 3 made its contribution.

Further effort will be needed to refine the quantification of optimal policy under uncertainty and learning. Priority should be given to the agreement on an adequate and fair decision criterion, to the formalization of deep uncertainty, and the study of technology uncertainty. While important and exciting research questions remain, there is no doubt that timely, determined, and coordinated action against climate change is mandatory.

6.5 References

A. B. Abel. Asset prices under habit formation and catching up with the joneses. *American Economic Review*, 80(2):38–42, 1990.

J. Y. Campbell and N. G. Mankiw. Permanent income, current income, and consumption. *Journal of Business & Economic Statistics*, 8(3):265, 1990.

G. M. Constantinides. Habit formation: A resolution of the equity premium puzzle. *The Journal of Political Economy*, 98(3):519–543, 1990.

M. Ikefuji. Habit formation in an endogenous growth model with pollution abatement activities. *Journal of Economics*, 94(3):241–259, 2008.

C. Kolstad and A. Ulph. Learning and international environmental agreements. *Climatic Change*, 89(1):125–141, July 2008.

R. Mehra. *The equity premium puzzle: a review*. Now Publishers Inc, 2007.

R. Mehra and E. C. Prescott. The equity premium: A puzzle. *Journal of Monetary Economics*, 15(2):145–161, March 1985.

W. Nordhaus. Critical assumptions in the stern review on climate change. *Science*, 317 (5835):201 –202, July 2007.

W. A. Pizer. The optimal choice of climate change policy in the presence of uncertainty. *Resource and Energy Economics*, 21(3-4):255–287, August 1999.

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

P. Weil. The equity premium puzzle and the risk-free rate puzzle. *Journal of Monetary Economics*, 24(3):401–421, November 1989.

M. L. Weitzman. Subjective expectations and Asset-Return puzzles. *American Economic Review*, 97(4):1102–1130, 2007.

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