Max-Planck-Institut
für Meteorologie

MAX-PLANCK-GESELLSCHAFT

International Max Planck Research School
on Earth System Modelling

Value of Information under Climate Targets: an Application of Cost-Risk Analysis

Delf Neubersch

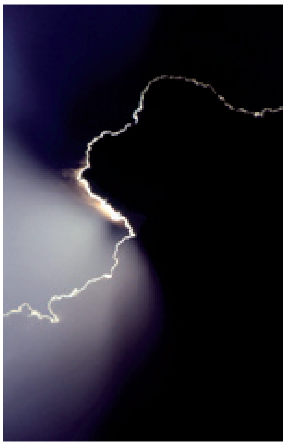

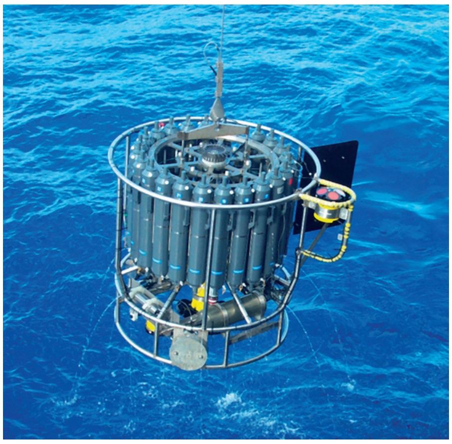

Berichte zur Erdsystemforschung 155 2014

Reports on Earth System Science

Hinweis

Die Berichte zur Erdsystemforschung werden vom Max-Planck-Institut für Meteorologie in Hamburg in unregelmäßiger Abfolge herausgegeben.

Sie enthalten wissenschaftliche und technische Beiträge, inklusive Dissertationen.

Die Beiträge geben nicht notwendigerweise die Auffassung des Instituts wieder.

Die "Berichte zur Erdsystemforschung" führen die vorherigen Reihen "Reports" und "Examensarbeiten" weiter.

Notice

The Reports on Earth System Science are published by the Max Planck Institute for Meteorology in Hamburg. They appear in irregular intervals.

They contain scientific and technical contributions, including Ph. D. theses.

The Reports do not necessarily reflect the opinion of the Institute.

The "Reports on Earth System Science" continue the former "Reports" and "Examensarbeiten" of the Max Planck Institute.

Anschrift / Address

Max-Planck-Institut für Meteorologie
Bundesstrasse 53
20146 Hamburg
Deutschland

Tel.: +49-(0)40-4 11 73-0
Fax: +49-(0)40-4 11 73-298
Web: www.mpimet.mpg.de

Layout:

Bettina Diallo, PR & Grafik

Titelfotos:
vorne:
Christian Klepp - Jochem Marotzke - Christian Klepp
hinten:
Clotilde Dubois - Christian Klepp - Katsumasa Tanaka

Berichte zur Erdsystemforschung 155 / 2014

Value of Information under Climate Targets: an Application of Cost-Risk Analysis

Delf Neubersch

Hamburg 2014

Reports on Earth System Science 155 / 2014

ISSN 1614-1199

Delf Neubersch
Max-Planck-Institut für Meteorologie
Bundesstrasse 53
20146 Hamburg

Als Dissertation angenommen
vom Fachbereich Geowissenschaften der Universität Hamburg

auf Grund der Gutachten von
Prof. Dr. Hermann Held
und
Prof. Dr. Uwe Schneider

Hamburg, den 27. Mai 2014
Professor Dr. Christian Betzler
Leiter des Fachbereichs für Geowissenschaften

ISSN 1614-1199

Value of Information under Climate Targets: an Application of Cost-Risk Analysis

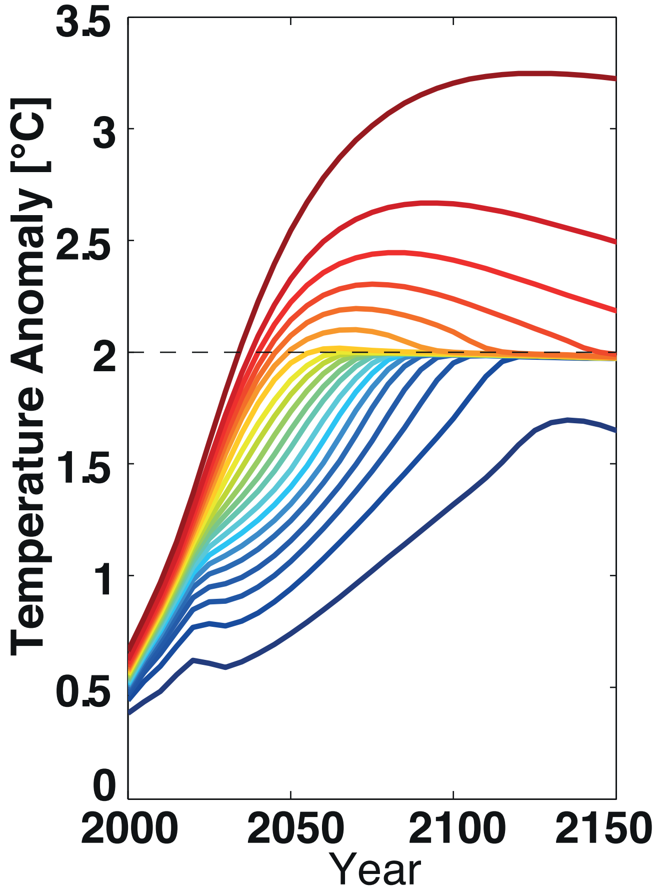

Delf Neubersch

Hamburg 2014

In memory of all the deceivingly great ideas

Eidesstattliche Versicherung

Hiermit erkläre ich an Eides statt, dass ich die vorliegende Dissertationsschrift selbst verfasst und keine anderen als die angegebenen Quellen und Hilfsmittel benutzt habe.

Hamburg, 8. April 2014

Author Attribution

My work on this thesis was done in close collaboration with Hermann Held and more loosely with Alexander Otto. About half of the material from Chapters 1-4 has been condensed into the form of an article and is under review for Climatic Change as of the time this thesis was printed as:

Neubersch D, Held H, Otto A: Operationalizing climate targets under uncertainty and learning, a Cost-Risk Analysis. Paper in review at *Climatic Change*.

All simulations and discoveries were made by myself whereas methods and interpretation where synthesized in collaboration. Hermann Held devised the calibration of CRA to a target of staying below 2°C with a likely chance, on the basis of discussions by the Conference of the Parties (UNFCCC). He further constructed the axiom of sacrifice inhibition and had the premonition that an analytical CRA can be written down concisely. Parts of Chapter 1 were reformulated from the introduction to the paper mentioned above written by Hermann Held, Alexander Otto and myself. Research for Chapter 5 was done only in collaboration with Hermann Held who derived and checked the analytical derivation with myself in collaboration.

To simplify the grammatical implication, this thesis sticks to first person plural throughout, although most parts of the thesis were created solely by myself. Apart from some formulations in Chapter 1, I authored all of the thesis' written contents and am responsible for most of the intellectual contents.

Abstract

The optimal policy that balances the cost of mitigation with the damages from climate change can be assessed by examining the interactions between the socio-economic system and the climate system. Traditionally, Cost-Benefit Analysis (CBA) is used for this problem but requires an approximate economic evaluation of the totality of climate damages which, in turn, calls for extensive impact assessments and ethical debates. Until such damage functions are agreed upon, another operational framework is essential for policy analysis. To fill the gap, climate (e.g. temperature) targets have been used in conjunction with Cost-Effectiveness Analysis (CEA). However, this analysis breaks down if a violation of the target is inevitable. This becomes especially salient when considering uncertainty and learning under temperature targets, given the presently infinite tail of climate sensitivity. A remedy was proposed that trades off the costs for mitigating climate change against the risk of exceeding climate targets: Cost-Risk Analysis (CRA). The implicitly defined preference order contained in the formulation of a climate target is absorbed into CRA by a calibration process, whereas it is an explicit constraint in CEA.

The UNFCCC's climate negotiations and many other organizations refer to climate targets in discussions, such as a 66% probability of keeping global warming below 2°C. Such a target includes an implicit assessment of the associated risk. This thesis explores, for the first time, the consequences of using such a target to calibrate CRA. We apply CRA to the climate problem including uncertainty and future learning and derive optimal mitigation paths. We calculate the value of future learning about the temperature response to be around 1/5 of total costs of climate protection (these costs include a monetized "risk" measure). The framework of CRA is based on expected utility maximization augmented by a risk-related term implemented by a risk metric which is subtracted from utility. A risk metric based on the probability of violating a temperature target leads to maximum emissions if it is learned that the temperature response is strong. We propose that such behavior is not in line with the preferences of the community supporting climate targets. Therefore, the thesis explores a risk metric based on the concept of degree years.

We develop a method of attributing the value of information and the cost of climate protection to their respective sources. A source can either be a change in consumption (i.e. economy related) or a risk-related utility change. Furthermore, an attribution to different time steps and states of the world is also presented. We find that about 2/3 of the value of information originates from the economy before 2050. The economic value of information offsets around 1/3 of the economic cost of climate protection (i.e. the cost of mitigation).

An advantage of calibrating CRA against a climate target is that the effect of normative parameters, such as the discount rate, on optimal policy is greatly reduced making CRA very robust. CRA remains operational even if the target is violated, and therefore, opening the possibility of investigating problems inaccessible to CEA. In summary, this thesis shows that CRA combines the benefits of CBA and CEA into a hybrid method that is well suited for policy advice in the next decades.

Zusammenfassung

Um die Abwägung zwischen den Kosten der Reduzierung von Treibhausgasemissionen und den Schäden des Klimawandels durchzuführen, müssen die Wechselwirkungen zwischen dem sozioökonomischen System und dem Klimasystem analysiert werden. Üblicherweise wurde Kosten-Nutzen Analyse (CBA) angewendet, um dieses Problem zu bearbeiten. Diese Analyse benötigt allerdings umfassendes Wissen bezüglich der Klimaschäden, was wiederum eine umfassende Klimafolgenforschung und ethische Debatten zur monetären Bewertung voraussetzt. Solange dieses Wissen nicht ausreichend vorhanden ist, wird eine alternative Methode benötigt, die eine Entscheidungshilfe ermöglicht. Bisher haben Klimaziele (Temperaturziele), zusammen mit der Kosten-Effektivitäts Analyse (CEA), diese Lücke zu schließen versucht. Diese Methode liefert allerdings keine Ergebnisse, sobald die Einhaltung des Ziels nicht mehr möglich ist. Dies ist besonders relevant bei Berücksichtigung von Unsicherheit und zukünftigem Lernen. Um ihre Funktionalität wieder herzustellen, wurde die „Kosten-Risiko-Analyse“ (CRA) vorgeschlagen. Sie beschreibt eine Abwägung zwischen Kosten und Risiko der Überschreitung. Die Formulierung eines Klimaziels impliziert eine Präferenz, die von CRA durch eine Kalibrierung der Abwägung absorbiert wird, wohingegen CEA das Klimaziel nur als Nebenbedingung implementiert.

Klimaziele werden unter anderem in den Klimaverhandlungen der UNFCCC diskutiert, wie zum Beispiel das Ziel den Anstieg der globalen Mitteltemperatur auf unter 2°C, mit einer Erfolgswahrscheinlichkeit von 66%, zu begrenzen. Ein Klimaziel impliziert eine Einschätzung des Risikos und diese Arbeit untersucht die Auswirkung der Kalibrierung von CRA an solch einem Ziel. In folge dessen wird CRA zum ersten Mal, unter Berücksichtigung von Unsicherheit und Lernen, auf das Klimaproblem angewendet. Es werden optimale Emissionspfade berechnet sowie der Wert von Informationen über die Temperaturantwort. Der Wert liegt ungefähr bei 1/5 der gesamten Kosten des Klimawandels (inkl. des monetarisierten Risikos). CRA basiert auf der Erwartungsnutzenmaximierung allerdings erweitert durch einen Term, der von der Risikometrik bestimmt und von dem Nutzen abgezogen wird. Wird die Wahrscheinlichkeit der Überschreitung als Risikometrik gewählt, kann das zu maximalen Emissionen führen, falls gelernt wird, dass eine substantielle Überschreitung unvermeidbar ist. Ein solches Verhalten ist nicht vereinbar mit den Präferenzen einer Gesellschaft, die Klimaziele unterstützt. Diese Arbeit untersucht deshalb eine Risikometrik, die auf Dauer und Ausmaß der Überschreitung basiert.

In dieser Arbeit wird eine Methode ausgearbeitet, die es ermöglicht, gefundene Wohlfahrtsunterschiede, etwa induziert durch sofortiges Lernen, verschiedenen Ursachen

zuzuordnen. So ist es möglich, den Wert von Information einer Änderung der Kosten oder des Risikos zuzuschreiben. Des Weiteren ist es auch möglich den Wert über die Zeit oder den möglichen Zuständen der Welt aufzuteilen. Es wird gezeigt, dass 2/3 des Wertes von Information aus einer Reduzierung der erwarteten Kosten vor 2050 besteht. Dieser ökonomische Wert von Information macht 1/3 der gesamten Mitigationskosten aus.

Eine Stärke von CRA ist ferner die Robustheit der optimalen Entscheidungen gegenüber Veränderungen in den normativen Parametern, wie zum Beispiel der Diskontrate. Die Ursache dieser Robustheit liegt in der Kalibrierung anhand eines bestimmten Klimaziels. CRA bleibt außerdem auch bei Überschreitung des Klimaziels funktional und ist daher geeignet, einen weitaus größeren Problemkreis als CEA zu evaluieren. Zusammenfassend zeigt diese Arbeit, dass CRA entscheidende Vorteile von CBA und CEA vereint und deutlich besser als ihre Vorläufermodelle geeignet ist, in den nächsten Jahrzehnten Entscheidungshilfe zu leisten.

Contents

Nomenclature

ASI	Axiom of Sacrifice Inhibition
BAU	Business as Usual
CBA	Cost-Benefit Analysis
CBGE	Certainty and Balanced Growth Equivalents
CCP	Chance Constrained Programming
CEA	Cost-Effectiveness Analysis
CRA	Cost-Risk Analysis
CRRA	Constant Relative Risk Aversion
ECCP	Expected Cost of Climate Policy
EVPI	Expected Value of Perfect Information
GR	Guard rail
GWP	Gross World Product
IAM	Integrated Assessment Model
MIND-L	Model of INvestment and technological Development including Learning
PRTP	Pure Rate of Time Preference
Safety	Probability of staying below guard rail
SDR	Social Discount Rate
SOW	State of the World

1. Introduction

This thesis is about decision frameworks that are used to distill policy recommendations, mostly on the basis of Integrated Assessment Models (IAMs). IAMs are formal representations of the interconnected socio-economic and climate system. With differing levels of complexity in the representation of the economy, the energy system, the climate system, spatial resolution and decision making processes, IAMs are versatile tools to investigate the interactions between human activities, the environment and implications of anthropogenic climate change.

A subset of IAMs (e.g. MERGE (Manne, 2005), DICE (Nordhaus, 2008), PAGE09 (Hope, 2011), WITCH (Bosetti *et al.*, 2006), MESSAGE (Messner & Strubegger, 1995), REMIND (Luderer *et al.*, 2013), FUND (Tol, 1997)) apply decision theory to determine (welfare) optimal decisions for investments or taxes. Hence, they can be used to translate the call for stabilization of greenhouse gas concentrations to avoid dangerous anthropogenic interference with the climate system by the United Nations Framework Convention on Climate Change (UNFCCC) into concrete emission reduction targets or investment strategies.

The appeal of this type of analysis stems from the attractiveness of the concept of a (counter-factual) rational decision maker and from the analogy between maximizing overall welfare and internalizing the climate externality.

1.1. Decision Frameworks

A straight forward implementation of overall welfare optimization is the Cost-Benefit Analysis (CBA) (e.g. Nordhaus (2008)) that weighs the cost incurred by strong mitigation action against the benefits from avoiding climate change induced damage.

However, a series of challenges complicate the application of CBA to the climate problem: the required comparability of different types of damage, huge uncertainties underlying the choice of an appropriate damage function (e.g. see Azar & Lindgren (2003), Pindyck (2013)), and limits to the applicability of CBA in the case of fat-tailed uncertainties about the climate system response to greenhouse gas emissions

(Weitzman, 2009). More research into impact models is required to, at least, address the issue of finding an approximate aggregate damage function (even on a sectoral level for a single region). The scrutiny with which any proposed damage functions are regarded is exemplified in Ackerman & Munitz (2012) who take apart the damage function used in FUND.

The time to reach a global agreement on binding emission reduction targets (or cumulative emission) is dwindling, if the possibility of limiting the rise of global mean temperature to values around 2°C compared to the pre-industrial value is to be preserved (Kriegler *et al.* (2014) show that an agreement after 2030 results in a violation of the 2°C target in the majority of models). Hence, decision-aiding cannot wait for the advent of better impact models.

The difficulties with CBA and the time pressure drove scientists to simplify the normative and ethical discussions by talking about climate targets (e.g. temperature targets). A significant fraction of the community sees climate targets as a viable alternative in the light of deep ("Knightian") uncertainty and preferences in-line with the precautionary principle (Iverson & Perrings, 2012; Athanassoglou & Xepapadeas, 2012). Obviously, the choice of a specific climate target is influenced in part by intuitive expectations of potential climate-induced damage. Yet the construction of a climate target avoids having to formalize the totality of all global warming impacts and their respective uncertainties – a task that poses so far unresolved conceptual and formal problems when having to be expressed with Knightian uncertainty. Instead, this so far unresolved formal task is often replaced by a climate target (or "guard rail") that we interpret as an informal convolution of sparse impact information and decision preferences under Knightian uncertainty, as long as no more robust method is available.

Thus, Cost-Effectiveness Analysis (CEA) of (exogenous) climate targets emerged as a method that avoids the difficulties of defining a climate change damage function and separates the evaluation problem from the policy analysis (Patt, 1999; Luderer *et al.*, 2012). If policy makers agree on a climate guardrail, the most cost-effective policy can be determined. If the costs (monetary or otherwise) of a policy intervention are found "sufficiently small", societal action can be taken (Patt, 1999).

Another way of dealing with the shortcomings of IAMs is to switch focus to scenario frameworks and notions of robust decision making (Lempert *et al.*, 2006; Weaver *et al.*, 2013), thereby abandoning the appealing notion of optimality. The decisions are then based around the concept of being safe in all possible outcomes our adapting to each new observation in a specific way defined by scenarios. This approach is not further discussed in this thesis.

1.2. Why a New Decision Framework?

Apart from the need for an agreement on a climate target, CEA leads to other difficulties which are most evident if uncertainty and learning is included in the decision model. A comprehensive introduction to IAMs and uncertainty in general is given in Rotmans & van Asselt (2001) and a review of uncertainty in economic models of climate change can be found in Golub *et al.* (2013).

While CBA elegantly deals with uncertainty (and anticipated learning) as it is based on expected utility maximization (Gollier, 2004), CEA has to be modified in order to accommodate uncertainty. The modification is necessary to accomodate for the fact that uncertainty leads to high temperature - low probability combinations which, if forced to be under 2°C, can dominate the analysis or even render it infeasible (Held *et al.*, 2009). This issue is particularly dramatic if the uncertainty has a fat upper tail. By generalizing the target into a probabilistic guardrail, i.e. fixing a maximum probability of overshooting the temperature target (Kleinen, 2005; Meinshausen *et al.*, 2006; den Elzen *et al.*, 2007; den Elzen & van Vuuren, 2007; Meinshausen *et al.*, 2009) this problem can be solved.[1] The resulting welfare optimization problem can be solved by chance constrained programming (CCP) as discussed in Held *et al.* (2009). CCP is a general method to incorporate an unlikely catastrophic outcome in an analysis as a probabilistic target because it can not be avoided with certainty.

The uncertainty discussed in this thesis is assumed to lie in the conversion of a change in CO_2 concentrations in the atmosphere into a mean temperature response. In general two factors play a role: the climate sensitivity and the transient climate response. The climate sensitivity represents the equilibrium temperature of a doubling of pre-industrial concentrations whereas the transient climate response characterizes the short term response to CO_2 increase. To keep the analysis simple we make use of the correlation between the two, found by Frame (2005), and deduce the transient climate response by perfect correlation shown in Lorenz *et al.* (2012b). This enables us to regard only the climate sensitivity as an uncertain parameter adjusting the transient climate response accordingly.

The decision problem is complicated further by considering (and anticipating) future changes in our knowledge about the climate sensitivity due to new observations (Kelly & Kolstad, 1999), the assimilation of paleo information (Schneider von Deimling *et al.*, 2006; Lorenz *et al.*, 2009) and improvements in theory and modeling of sub-scale processes.

The possibility of receiving new information has been shown to change the resulting optimal policy substantially (e.g. O'Neill & Melnikov (2008); Webster *et al.* (2008); Webster (2000)). Allen & Frame (2007) even argue that a substantial fraction of climate response uncertainty could be compensated by learning after 2050 in order

[1] Probabilistic guardrails have also been explored in the "tolerable windows approach" (Bruckner & Zickfeld, 2008) that can be interpreted as "CEA without optimization".

to still be able to comply with ambitious climate targets. We recognize that it can not always be clearly stated in which direction the incorporation of learning affects the optimal policy (Webster, 2000), although evidence is accumulating that it does not support postponing mitigation (Ha-Duong *et al.*, 1997; Ha-Duong, 1998; O'Neill *et al.*, 2006; Lange & Treich, 2008). We do not want to elaborate on this discussion, but rather focus on the novel method of analysis.

It is important to be able to model future changes in uncertainty to find the effect on decision but also to calculate the economic value of new information about the climate response (e.g. Nordhaus & Popp (1997)). It is of vital importance for technologies and projects that have information as an output to be able to estimate its value and therefore justify investments (e.g. projects such as GEO-BENE (EU-FP6) (Fritz *et al.*, 2008) and EuroGEOSS (EU-FP7) (Pearlman *et al.*, 2011)).

It has been shown, however, that formulating CEA as a CCP, by defining a probabilistic target, still can not deal with future learning (Schmidt *et al.*, 2011). As it does not comply with two of the axioms of rational decision making (von Neumann-Morgenstern Axioms), or equivalently by not being structurally synonymous to an expected utility maximization (Gollier, 2004), this CCP formulation leads to a break down of the analysis. The following problems can occur:

1. After a perfect learning event, the probability of violation is either zero or unity for each state of the world. The only way to meet a probability target is to reduce the probability of violation to zero. This implies that all states of the world have to remain below the guard rail if learning is anticipated which forces the decision maker to do as much mitigation as needed to keep the state of the world with the highest climate sensitivity under the guard rail.
2. Due to the mechanism described in 1., learning leads to stronger mitigation without any added benefit (as no climate effects are included), therefore, CEA can lead to a negative value of information and reject learning, creating an incentive to prevent research. See Schmidt *et al.* (2011) for an analytical derivation of the negative value of information. That a formulation with CCP can lead to negative value of information was already known before and is discussed in Blau (1974), Jagannathan (1985) and Lavalle (1986).
3. Learning that climate sensitivity is very high and that the target cannot be met leaves the decision maker without any solution, i.e. the simulation becomes infeasible. Part of the literature "bypasses" this conceptual difficulty imposed on CEA by artificially truncating the upper tail of the distribution of climate sensitivity. While this generates a lot of academically relevant insight, a key property of the climate problem has been chopped off. It implies that the crucial feature Weitzman (2009) pointed to as a weak point of CBA, has been ignored for CEA as well to save its functionality – hence, not providing a conceptually satisfying solution.

Of course the above problems are based on the same underlying reason: CEA with CCP violates the axioms of rational decision making. We argue, that a decision tool that features a negative value of information is inadequate for policy advice and favor an expected utility maximization approach. However, the necessary tools are not yet available to do a full fledged CBA of the climate problem and until such time an interim method is needed. This method should be based on the expected utility maximization framework but incorporate the concept of a climate target. A new decision tool is needed!

1.3. Cost-Risk Analysis in Short

Circumventing these conceptual and axiomatic issues, Schmidt *et al.* (2011) proposed Cost-Risk Analysis (CRA) in which a trade-off is made between the risk of overshooting climate targets and the economic utility from fossil-fuel based consumption, thereby providing a normative decision criterion that does not explicitly require a climate damage function, but builds on a consensus-based climate target transformed into a penalty function operating additively on the utility level, thus preserving the equivalence to the expected utility framework.

This method is one step away from purely quantitative CBAs towards a more qualitative assessment. CRA can be interpreted as a consequence of Morgan *et al.* (1992), who suggest to model a complex system the simpler, the sparser the level of the modeler's understanding of the system under consideration. Accordingly, CRA models the preference order of the community supporting a temperature target in view of poorly understood damage functions in the simplest way.

In contrast to CBA, CRA does not rely on a formal aggregation of possible impacts of climate change into a damage function but rather divides the problem into two parts. The first is to find a climate target which is in itself an assessment of the impacts on a meta level and is easy to understand. The second part is to find a risk metric and implementation which reflects the preferences intrinsic to the formulation of the target. This difference is visualized in a flow chart to clarify the process involved in the evaluation of the effects of climate change (Figure 1.1).

The idea to incorporate a target as a trade-off instead of a hard constraint originated in other fields which have similar problems with negative value of information. Bordley & Pollock (2009) suggest to incorporate given cost targets in engineering design problems into the model formulation as a trade-off. Also Jagannathan (1985) uses the probability of violating a target in a trade-off rather than a chance constraint. A related approach is positive mathematical programming (Howitt, 1995), in which the utility function is extended by a calibrated term to force the model to reflect a pre-defined condition. This condition could be, for example, a probabilistic target. CRA is put into context with respect to other decision frameworks more thoroughly in Section 2.3.4, after a detailed explanation.

Addressing the part within the climate economics community that currently prefers CEA to CBA, this study builds upon the literature described above and presents

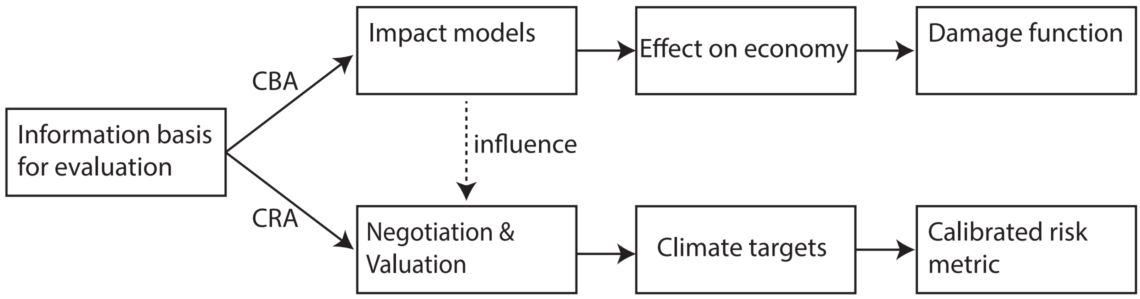

Figure 1.1. – Comparison of the process that leads to the representation of the effects of climate change in CBA and CRA. The impact modelling influences the negotiation and valuation which leads to climate targets. The negotiation process encompasses all conveyed uncertainties.

the first application of CRA to the mitigation of climate change as a generalization of CEA to include uncertainty and anticipated learning. To implement CRA it is necessary to define a risk metric and tune it. The risk metric defines the preference order of temperature paths, i.e. which temperature paths are preferred over others. The risk metric is then tuned by a calibration process to internalize a given target. This process is thoroughly described in this thesis.

In summary, the main innovations of the approach described in this thesis consist of (i) the choice of risk metric, (ii) the reasoning for it and, (iii) the calibration against existing discussions and targets. To the authors knowledge, there is no study that uses an expected utility maximization framework and internalizes the preferences implied by supporting a climate target. Such an approach enables the climate target community to analyze aspects such as the value of information.

1.4. Outline of the Thesis

Most of the analysis is based on the IAM MIND-L (Model of Investment and Technological Development including Learning) (Edenhofer *et al.*, 2005; Held *et al.*, 2009; Lorenz *et al.*, 2012b) described in more detail in Chapter 2. The model is made up of a simple one-box climate model, a representation of the renewable energy sector, the fossil fuel extraction sector together with the fossil energy production. Furthermore, the model includes endogenous learning for energy efficiency, labor efficiency, fossil energy production and renewable energy production. The underlying economic model is a Ramsey type growth model.

The overall question this thesis set out to answer is: How can the value of information be calculated if climate targets are included in the analysis? The following chapters go into different aspects of the answer revolving around CRA.

Chapter 2 develops the framework of a CRA and motivates our choice of risk metric: expected discounted degree years. This resonates well with Schneider & Mastrandrea

(2005) who also proposed degree years as a sensible way to measure climate risk. Degree years are calculated by finding the area between a temperature path and the guard rail. We show why the original formulation of the risk metric as the probability of violating the target is unacceptable as a proper representation of environmentalists' preference order. A calibration of the risk metric is proposed by referring to the preference implicit in the Cancun accord (UNFCCC, 2010) of staying below 2°C with a (likely) 66% probability. The last part of the chapter applies CRA to an IAM for the first time and calculates the value of information. We also discuss similarities to other approaches and compare our results with a traditional CEA.

Chapter 3 investigates the origin of the value of information by making use of the additive structure of CRA. We find it is possible to divide any welfare change into parts not only according to their origin in time and state of the world but also with respect to their type: economic and risk-related. To the authors knowledge, such a division has not been presented before. The effect of changing the target is analyzed and it is clarified that adjusting the guard rail has a different effect than adjusting the target probability if learning is considered. The risk metric is also varied and its effect on the calibration and the value of information is determined. Lastly, the underlying assumption of an anticipated learning event is analyzed by finding the welfare gain of anticipation.

Normative parameters in MIND-L define the risk aversion of the decision maker as well as the time preference. The climate target itself is also a normative decision of sorts and it is shown that it reduces the effect of risk aversion and time preference on optimal decisions. In Chapter 4 we combine risk aversion and time preference into a social discount rate using the Ramsey equation to facilitate an analysis of the effect of discounting on the value of information. We show that discounting only affects optimal decisions marginally but produces very different valuations of the same situation. The chapter concludes with an analysis of the loss of welfare that is to be expected if optimal decisions are made based on a discount rate that does not reflect the preferences of society.

Chapter 5 develops the basis of a simplified model of CRA and IAM with the goal of facilitating the implementation in areas where a less computationally intensive model is needed or a simpler model is necessary to improve understanding. A full-fledged IAM is difficult to use for highlighting the driving mechanisms. We derive a static model of CRA and calculate the optimal emissions depending on the climate sensitivity that is learned. In a second part, the optimum of MIND-L for a case without learning is analyzed with respect to the control parameters. A regression is done to fit a second order function to the welfare equation thereby finding the important control parameters. This analysis also allows us to find the welfare for small deviations from the optimum as well as the amount of variation in decision variables that produces an acceptable loss of welfare.

The thesis closes with a summary of the findings and an outlook on the future of CRA and potential future research topics.

2. Cost-Risk Framework applied to MIND-L

2.1. Introduction

The Cost-Risk Analysis (CRA) is a trade-off analysis, in principle like a standard Cost-Benefit Analysis (CBA), with the key difference in determination, application and interpretation of the “benefits”. The benefits in CRA come from a relative reduction of potential for danger due to a reduction in temperatures. The starting point is a climate target. The analysis only makes sense for a community that supports the notion of climate targets. A conclusive motivation for climate targets can be found elsewhere (WBGU, 1997; Ott *et al.*, 2004; Oppenheimer & Petsonk, 2005), we only point out that it is a way to express a preference about the future as long as full-blown impact models are unavailable. A discussion of modifications and re-interpretations to the 2°C target can be found in Geden (2013).

Once a climate target is established, the most cost efficient solution can be found that reaches this target with Cost-Effectiveness Analysis (CEA). One method to implement CEA including uncertainty is chance constrained programming (CCP) which we imply when referring to CEA. Chapter 1 stressed that this produces problems when considering learning. Therefore, a utility penalty function is implemented that leads the optimal solution to reach the equivalent target. This utility penalty function will be referred to as the risk metric which is calculated from the temperature profile by a penalty function. In essence the risk metric defines a preference order on temperature paths.

It is theoretically possible to base the risk metric on a different variable like emissions or atmospheric CO_2 concentrations but as targets are often formulated in terms of temperature we also formulate the risk metric in terms of temperature. We call this process of finding the equivalent penalty function “calibration”. This process is similar to what is used in positive mathematical programming (Howitt, 1995) and also has the same drawbacks of non-uniqueness. Mathematically speaking, a potential is added to the optimization that forces the optimum to the desired result.

After the calibration of the risk metric to a specific climate target, parameters can be varied and effects calculated, always on the basis of the preferences of a climate target community. Resolving the uncertainty in the model allows the calculation of the value of research into the uncertain parameter. In our analysis the uncertain parameter is a combination of climate sensitivity and the transient climate response, as we assume perfect correlation (Lorenz *et al.*, 2012b).

This chapter recaps the propositions for CRA by Schmidt *et al.* (2011) and demonstrate drawbacks for the first time. As a solution, we introduce expected discounted degree years as a viable risk metric and talk about the implications of this choice. We calibrate against a climate target taken from discussions of the UNFCCC and then apply the framework to an integrated assessment model (IAM) for the first time. We compare the results to CEA and discuss their implications. The calculation of the value of information about the climate response exemplifies the use of CRA and we close with a discussion of the results.

2.2. Static Model as Guide

At the heart of CRA lies a trade-off between the cost of mitigation and the perceived risk of high temperatures. Note that the term “risk” is used in layman terms and not in the strict sense (not as product of loss and probability). This section introduces the formalism needed for the discussion. First a static model is introduced to guide intuition followed by the addition of uncertainty. The next section expands the problem into the time dimension.

We start with the formulation of the minimization problem for a CRA with mitigation costs $C(E)$ and a penalty function $R(E)$ (representing the risk) both depending on cumulative emissions E that are chosen by the decision maker. In order to motivate why further development of CRA compared to the suggested scheme of Schmidt *et al.* (2011) is regarded as necessary and for the sake of conceptual clarity in this sub-section we restrict the discussion to a static picture, using cumulative emissions as the crucial control variable (Lorenz *et al.*, 2012a). Using β as the trade-off parameter between costs and risk the deterministic minimization problem reads:

$$\min_{E} \{C(E) + \beta R(E)\} . \tag{2.1}$$

This is the simplest way to formulate CRA and it also reveals that CRA is just a CBA with a different interpretation of damages. The difference to CBA is further discussed in Section 2.3.4. There is no uncertainty, no time and no guard rail or target (yet). If society agrees on a certain emissions budget E_{g} that complies with its normative preferences by fulfilling the climate target, we can tune β to make E_{g}

the optimal strategy by setting β to:

$$\beta_{\text{cal}} = -\frac{C'(E_{\text{g}})}{R'(E_{\text{g}})}. \tag{2.2}$$

We call this process "calibration" because the trade-off is made in such a way that a specific target is met (in the case above, an emission target). In practice, policy makers usually talk about temperature targets, therefore, we introduce a simple relationship between temperature and emissions using the sensitivity γ. This assumes an approximately linear relation between maximum temperature and cumulative emissions which was shown in Allen *et al.* (2009). We further assume that γ is uncertain and is distributed by $f(\gamma)$. The penalty function $R(E)$ from above is now replaced with $R(\gamma E)$ implying that the penalty is applied to the temperature $T = \gamma E$. The optimization functional then reads:

$$\min_{E} \left\{ C(E) + \beta \mathbb{E}\left[R(\gamma E) \right] \right\}. \tag{2.3}$$

Here $\mathbb{E}[..]$ indicates the expected value operator. The goal is to calibrate this model to a policy target of staying below a guard rail $T_{\text{g}} = \gamma_{\text{g}} E_{\text{g}}$ with a probability of p_g. Due to the simple structure, we can invert the cumulative distribution function F of γ to find the sensitivity γ_{g} for which the probability to lie below is p_{g}:

$$\gamma_{\text{g}} = F^{-1}(p_{\text{g}}), \tag{2.4}$$

$$E_{\text{g}} = \frac{T_{\text{g}}}{F^{-1}(p_{\text{g}})}. \tag{2.5}$$

Again we can calculate the calibrated trade-off parameter, including the expected utility operator:

$$\beta_{\text{cal}} = -\frac{C'(E_{\text{g}})}{\frac{d}{dE}\mathbb{E}\left[R(\gamma E_{\text{g}}) \right]}. \tag{2.6}$$

After a successful calibration, the risk metric R reflects the preferences implied by supporting the climate target. In a next step, parameters can be changed and the reactions of the decision makers can be investigated. The change we consider here, is a change in the available information, and to simplify the discussion we consider perfect learning in the following. As there is no time dimension yet, information is available immediately. Before the optimization the policy maker receives a message providing the correct value of γ. The probability of receiving a certain message is taken from the prior distribution $f(\gamma)$ which ensures that the overall expected value of the stochastic parameter is unchanged by the learning event. The calibration of the trade-off is not changed and remains at β_{cal} because we regard the case without learning as the baseline calibration case. In the case of perfect information the problem reads:

$$\mathbb{E}\left[\min_{E(\gamma)}\left\{C(E(\gamma)) + \beta_{\text{cal}} R(\gamma E(\gamma))\right\}\right]. \tag{2.7}$$

Now an emission budget can be chosen for each value of γ because its value is known perfectly. The goal variable of the optimization is calculated by taking the expected value across the distribution $f(\gamma)$. The outcome of such a simulation can be compared to the outcome of the case without learning to find the effect of resolving uncertainty.

2.2.1. Probability of Violation as Risk Metric

To gain more insight into CRA we now look at explicit risk metrics. A risk metric converts temperature into a penalty to the optimization which we call “risk”. First, we define a term to simplify the discussion:

Safety is given by the probability, in percent, that the temperature increase will remain below a defined guard rail. It is the complement to the probability of violation.

For the static case a utility penalty function it is analogue to the formulation of climate damages and a damage function. If we use a risk metric that equals the probability of passing some defined guard rail (i.e. the complement to the safety: Risk = 100% - Safety) a problem can arise after learning that we want to elaborate on in this section. Such a risk metric was one of the proposed methods by Schmidt *et al.* (2011). Such a risk metric for a scenario without learning can be represented by the expected value of a Heaviside function Θ:

$$\begin{aligned}
\mathbb{E}\left[R(\gamma E)\right] &= (1 - \text{Safety}) \qquad (2.8)\\
&= \mathbb{E}\left[\Theta\left(\gamma E - T_{\text{g}}\right)\right] \qquad (2.9)\\
&= \int_0^\infty \Theta\left(\gamma E - T_{\text{g}}\right) f\left(\gamma\right) \mathrm{d}\gamma \qquad (2.10)\\
&= 1 - F\left(\frac{T_g}{E}\right). \qquad (2.11)
\end{aligned}$$

After the perfect learning event however, we have Equation 2.7 as the functional. The minimization is inside the expected value operator, meaning that we are looking at an ensemble of minimizations that have no interaction. We can therefore investigate the optimizations separate from each other:

$$\min_{E(\gamma)} \{C(E(\gamma)) + \beta_{\mathrm{cal}} R(\gamma E(\gamma))\} \,. \tag{2.12}$$

The safety is either 0% or 100% depending on if the guard rail is passed or not. If the risk metric is based on safety, the risk is either 1 or 0 respectively. After perfect learning $f(\gamma)$ becomes a Dirac delta function at the value that was learned and so the risk function is simplified to:

$$R\left(\gamma E\right) = \Theta\left(\gamma E - T_{\mathrm{g}}\right). \tag{2.13}$$

If an emissions budget E_{crit} that would lead to the reduction in risk from 1 to 0 is extremely costly (as would happen if we learn climate sensitivity is high) then it would be optimal to do no mitigation. This situation is shown in Figure 2.1.

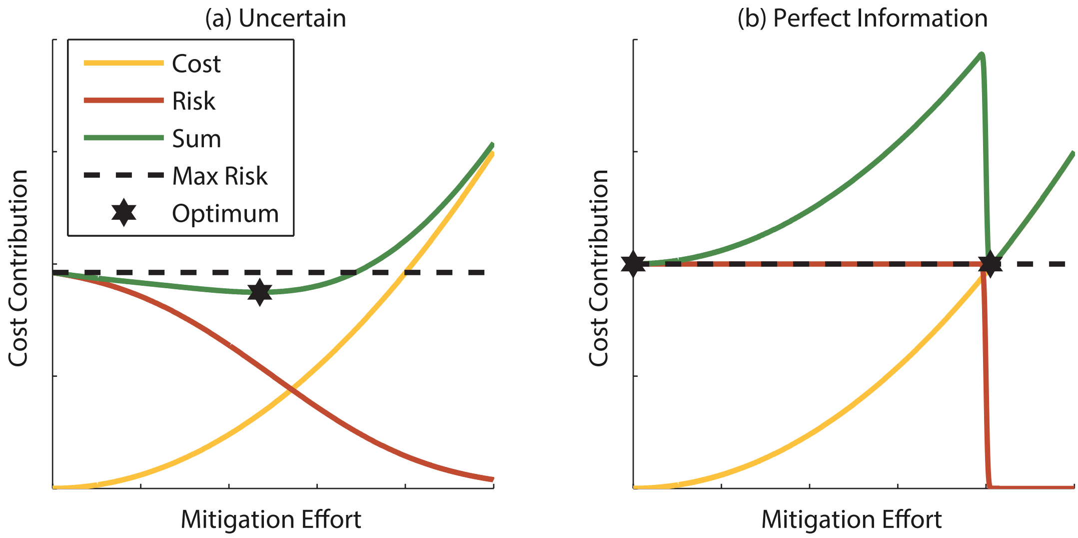

Figure 2.1. – Cost and risk curves for a case where the risk measure is defined by the probability of violation. The abscissa is defined as mitigation effort, i.e. the difference between maximum and chosen emission budget. (a) features a risk function calculated as the expected risk over a distribution of sensitivities, as in equation 2.8. (b) demonstrates the situation with a risk function based on the same equation but with a certain value for γ representing a specific state of the world after perfect learning.

It illustrates two different situations with the same cost function but different risk functions due to different uncertainties while using the same risk metric. If the variance of $f(\gamma)$ is sufficiently large, the optimum is clearly visible but if the probability of violation drops suddenly (as would happen if $f(\gamma)$ were a Dirac delta function, i.e. after learning a specific value for γ) at high mitigation effort, a local optimum is created that might or might not lie below the optima for no mitigation. This situation would mean that if policy makers learn that climate sensitivity is high, they would abandon mitigation.

2.2.2. Linear Risk Metric

We develop an axiom to exclude the behavior described in the case above, as we do not believe it to be a behavior that is in line with preferences of the community supporting climate targets.

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

As we can see from Figure 2.1, the problem of climate sacrifice originates from multiple optima in the optimization. A little change in emissions (or cost of mitigation) causes the solution to jump to a different regime. If multiple optima are to be avoided, the sum of cost and risk has to be convex to have a single minimum. Assuming that the cost function can be any convex function, implies that the function $R(\gamma E)$ has to be at least linear to guarantee that the sum is again a convex function in E.

We chose to use a linear function because it is the simplest solution (in accordance with Occam's razor) and the limiting case of a convex function. A risk metric based on a linear dependency on mitigation is considered from this point onwards.

To preserve the key ingredient of CEA we assume that there is no risk for temperatures below the guard rail. Therefore, the resulting risk metric is a function which penalizes transgressions of temperature beyond the guard rail in a linear way:

$$R(T) = \Theta(T - T_{\mathrm{g}}) * (T - T_{\mathrm{g}}). \tag{2.14}$$

We express the metric only in terms of temperature as that is the relevant variable in the upcoming analysis. Using this metric, the problem of sacrificing cannot occur because it will always be optimal to mitigate as much as possible until the marginal benefit is equal to the marginal reduction in "risk". Due to the Heaviside function Θ, the risk function includes a kink at the guard rail making it convex which means we can still guarantee a convex optimization. In the following we explain how this concept can be applied to the dynamic setting.

2.3. Dynamic Framework

To apply this framework to a dynamic optimization the following choices have to be made:

1. is the penalty calculated for each time slice and then aggregated or is it calculated for the maximum temperature so that a decrease in temperature is never rewarded?
2. if the penalty is calculated per time slice, is it discounted?

This is a normative decision and there is little guidance to help. Out of many possible interpretations, we regard the climate problem in its basic structure as reversible and decide that a decrease in temperature should be rewarded as well as an earlier decrease. Even if irreversible tipping points exist, they will have a certain inertia which may allow for a degree of reversibility. Therefore, we calculate the penalty in each time slice and then aggregate. We transform the optimization problem into a maximization of welfare and include a utility function U that depends on decisions X and time t. To ensure time consistent solutions we choose to apply the same exponential discounting with a discount rate δ to both the economic and the risk part of the trade-off.

To test this framework with a dynamic IAM we use MIND-L which is a climate-economy-energy model. This model is used throughout the thesis and is introduced more thoroughly in the next section. MIND-L includes an uncertainty in the climate response, i.e. for a deterministic emission path there are many possible temperature paths. This uncertainty is represented in equation 2.15 by summing over the states of the world (SOW) s and multiplying with the probability p_s of each SOW.

Formulated as a maximization problem we arrive at the discounted utility welfare equation for a CRA in a probabilistic setting without learning as follows:

$$W = \max_X \sum_{t=0}^{t_{\text{end}}} \sum_{s=1}^{S} p_s \left\{ \underbrace{U(X,t)}_{\text{economic}} - \underbrace{\beta R(T(X,t,s))}_{\text{risk-related}} \right\} e^{-\delta t} \tag{2.15}$$

To include learning we have to divide up the decisions into before and after learning. Let X_0 consist of all decision variables for all time steps before learning and $X_{1,s}$ is made up of all decisions variables for all time steps after learning that s is the true SOW. The optimal welfare is then calculated by the following equation:

$$\begin{aligned} W = \max_{X_0, X_1} \sum_{s=1}^{S} p_s & \left\{ \sum_{t=0}^{t_{\text{learn}}} \{U(X_0,t) - \beta R(T(X_0,t,s))\} e^{-\delta t} \right. \\ & \left. + \sum_{t=t_{\text{learn}}+1}^{t_{\text{end}}} \{U(X_0, X_{1,s}, t) - \beta R(T(X_0, X_{1,s}, t, s))\} e^{-\delta t} \right\}. \end{aligned} \tag{2.16}$$

2.3.1. “Giving up”

Before implementing the linear risk function we show, by example, that the effect of giving up on mitigation, i.e. sacrificing, can also occur in the dynamic model. We use a value for β that would reach a safety of 66%[1] for a 2°C target without considering future learning (the choice of target is discussed in Section 2.3.2). The risk metric is just the complementary value to safety and is calculated per time slice analogue to equation 2.10:

$$R(t) = \sum_{s=1}^{S} p_s \Theta(T(t,s) - T_{\mathrm{g}}) \tag{2.17}$$

Figure 2.2a shows the situation without learning. The optimal strategy is to mitigate just enough to keep 66% of the possible futures below the guard rail. In 2.2b a learning scenario is shown with exactly the same parameters but this time the decision maker has perfect information about the climate response and can adjust the decisions accordingly from 2015 onwards. The colored paths represent the SOWs with low (blue, lower curves) to high (red, upper curves) climate sensitivity.

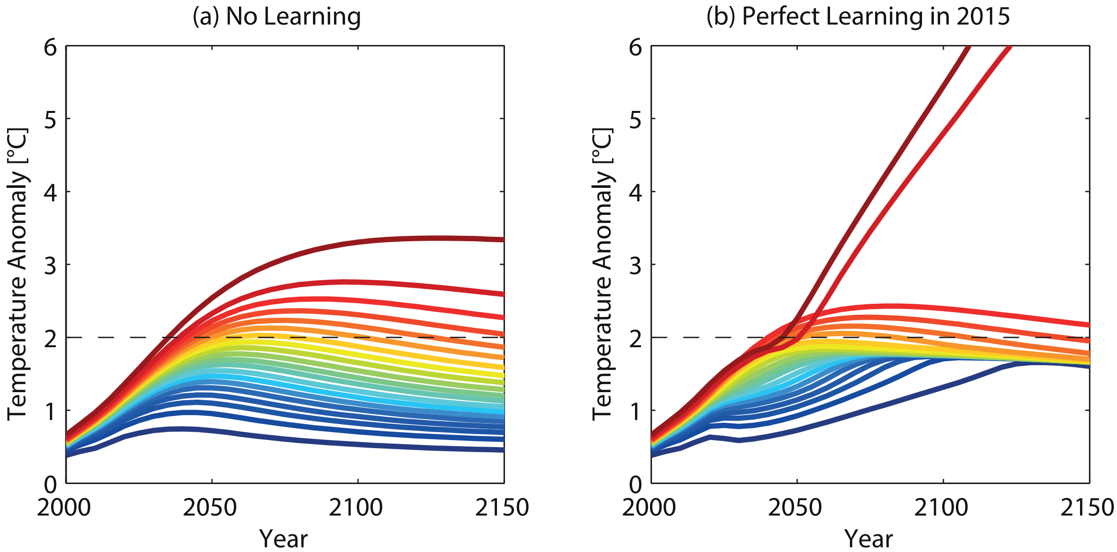

Figure 2.2. – Demonstration of “giving up” in MIND-L. The model is based on uncertainty in the climate response producing different temperature paths. The calibration was done to ensure that 66% of the paths stay below 2°C temperature anomaly (by interpolation). Panel (b) shows clearly that mitigation is abandoned if a strong temperature response (red, upper lines) is learned to be true.

The decision maker decides to drive up the temperature drastically if the climate sensitivity is found to be in the upper 10% quantile. As shown in the previous

[1]As the sampling used in the numerical simulation results in 5% steps of probability, an interpolation is necessary in the calibration to be able to tune the trade-off to 66%.

section, this is because the cost of reducing it below the guard rail is higher than the constant risk that a trespassing of the guard rail incurs. In other words, once above the guard rail, increasing the temperature further does not come with added risk.

In the static model we found that using a linear risk metric eliminates the sacrificing behavior and therefore we translate the risk metric into the dynamic form by applying a penalty function in each time slice and each SOW.

2.3.2. Degree Years as a Risk Metric

Applying a linear risk metric together with a threshold temperature in a dynamic setting has an alternative interpretation: degree years. Mastrandrea & Schneider (2004) use the concept of degree years (DY) in their article to assess the climate by calculating time integral of the temperature path above a specified threshold. Using this terminology we can redefine our risk measure as measuring "expected discounted degree years". To clarify this choice in definition, we compose the total aggregated welfare from risk:

$$W_{\mathrm{R}}(X) = -\beta \sum_{t=0}^{t_{\mathrm{end}}} \sum_{s=1}^{S} p_s \Theta\left(T(X,t,s) - T_{\mathrm{g}}\right)\left(T(X,t,s) - T_{\mathrm{g}}\right) e^{-\delta t}. \tag{2.18}$$

The equation calculates the temperature above the guard rail T_{g} for each SOW and time instant (the penalty function). Then the expected value is calculated, which is discounted and summed over time. Therefore, the risk is composed of the trade-off parameter β and the calculated discounted expected degree years. Because of discounting, a transgression of the guard rail has less negative effect if it happens further in the future.

Mastrandrea & Schneider (2004) further give a reasoning for such a form of risk metric by identifying five reasons for concern with equal weight. As temperature increases, gradually more and more of the reasons for concern become important. This leads to a linear risk increase and supports using degree years as a risk metric. They also suggest additionally using the maximum exceedance amplitude as a further risk metric but we found them to correlate greatly in our simulation due to inertia in the energy and climate system and therefore only consider degree years. It should be noted that Schmidt *et al.* (2011) also mentioned the possibility of using a function of degree years and maximum exceedance amplitude for the risk but did not further elaborate.

2.3.3. Calibration

Now that the risk metric is defined, the choice of the trade-off parameter β is still open. We discussed in previous sections that it can be calibrated against a specific desired solution. The target of constraining global mean temperature increase to 2°C above pre-industrial conditions has been discussed prominently in the climate policy debate. The 17$^{\text{th}}$ Conference of the Parties (COP) to the UNFCCC refer in their decisions to "aggregate emissions pathways consistent with having a likely chance of holding the increase in global average temperature below 2°C or 1.5°C above pre-industrial levels" (UNFCCC, 2011). We use the interpretation of "likely" as implying a probability of observing the target of at least 66%, as recommended in Mastrandrea *et al.* (2010). Assessment of these targets with respect to emission budgets can be found in Rogelj *et al.* (2011, 2012). Building on this assumption and the statement from COP17, we formulate the climate target in our analysis as follows:

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

By tuning β in Equation 2.15 until the solution satisfies the target above, we derive the perceived risk that policy makers and scientists, that support a 2°C target, imply. The information structure chosen to calibrate the trade-off parameter should be the one that is subjectively present in the minds of the political actors that are discussing this target. Furthermore, such a target becomes ill-posed after perfect learning so the calibration necessarily has to be done in the scenario without learning.

The method used to calculate the safety in the calibration process uses the maximum temperature of each temperature path to interpolate the probability to stay below the guard rail. See Appendix A.1 for the equation and a discussion.

After a short review of similarities to other approaches, the rest of the chapter deals with the application of this framework to MIND-L. The method is compared to CEA by looking at key variables of the simulation and the value of information is calculated.

2.3.4. CRA in Context

For the static setting, it is clear that CRA and CBA are indistinguishable from each other, if the risk function is interpreted as a damage function. In a dynamic setting one can differentiate between economic, non-economic or hybrid impacts of climate change. Purely economic impacts on the production (Gross World Product, GWP) are most common due to the plausibility of a damage function (the more production there is, the more can be damaged) and the popularity of DICE (Nordhaus, 2008). The non-economic impacts can be either applied to the consumption or utility directly (i.e. in- or outside of the utility function). In models such as MERGE (Manne, 2005), PAGE09 (Hope, 2011) and in Acemoglu *et al.* (2012) the non-economic damages are applied to the consumption multiplicatively. The model FUND (Tol, 1997) and theoretical discussions by Weitzman (2009) consider additive damages to the utility. Many of these models become increasingly complex by including regions and many different damage sources making it difficult to understand what is happening.

We apply CRA to an IAM in the simplest way possible and work towards qualitative understanding of such a framework.

A discussion about the effect of varying the distribution of damages between additive utility damages and production damages can be found in Barrage (2012). She shows that if all damages are applied to production then the optimal carbon tax is overestimated by 5%, whereas if all damages are applied to utility then there is an underestimation of 20%. The CRA presented in this thesis uses only additive utility "damages" formally but argues the functional form of the impact differently. Where past studies always considered studies about loss of life, migration or willingness to pay for species conservation to find damage functions on utility, our analysis presumes that the climate debates have already aggregated this information in the form of climate targets. We are not aware of any literature that uses additive utility damages in an inter-temporal IAM to absorb preferences of decision makers. Note that we acknowledge that a climate target is in itself subject to normative preferences of the supporting community. If policy makers do not find climate targets a useful tool, then CRA is not applicable in the form presented here.

Positive mathematical programming (Howitt, 1995) is also related to this topic. It is the process of adding a function to the objective function and tuning it in a way that a predetermined goal is forced to be optimal without implying explicit constraints. In other fields, especially in engineering (Bordley & Pollock, 2009), similar techniques have been proposed to include a desirable target in an optimization without chance constraints or leaving the expected utility framework.

2.4. Application to MIND-L

For the numerical analysis, we use the Model of Investment and Technological Development (MIND-L) in the form presented by Lorenz *et al.* (2012b). MIND-L is an extension, through the addition of learning, to the stochastic model MIND-H ("H" for "hedging") by Held *et al.* (2009), which itself was developed from the deterministic model presented by Edenhofer *et al.* (2005). MIND-L is an IAM consisting of three parts: economy, energy, and climate. Other models falling into this category are: MERGE (Manne, 2005), DICE (Nordhaus, 2008), PAGE09 (Hope, 2011), WITCH (Bosetti *et al.*, 2006) and FUND (Tol, 1997). MIND-L is a forward-looking Ramsey-type macro-economic growth model[2] that comprises induced technological change in the energy sector, the latter consisting of a renewable and a fossil sector. For a summary of other models with endogenous technical change see Baker & Shittu (2008).

An energy balance model (Kriegler & Bruckner, 2004) represents the climate module within MIND-L that links emissions to global mean temperature change. That

[2] We would like to mention literature that points to problems with this type of growth: Cooke (2013)

energy balance models can reproduce the average temperature trend of more complex climate models adequately is shown in van Vuuren *et al.* (2011).

MIND-L is implemented in the modeling language GAMS with the numerical solver CONOPT and all evaluation and post-processing is conducted in MATLAB.

2.4.1. Setup

MIND-L is based on a Ramsey type growth model with a constant elasticity of substitution production function and a utility function with constant relative risk aversion:

$$U\left(C\left(t,X\right)\right)=L(t)\frac{1}{1-\eta}\left(\frac{C\left(t,X\right)}{L(t)}\right)^{1-\eta}. \tag{2.19}$$

The consumption C depending on time t and controls X are inputs to a utility function with decreasing marginal utility to catch the risk aversion of society. The population development is given exogenously by $L(t)$. The constant relative risk aversion is set to $\eta = 2$ in this chapter. In Chapter 4 this assumption is analyzed. The controls X are made up of investments into fossil fuel extraction, fossil energy production, renewable energy production, energy efficiency, labor efficiency and aggregated capital representing the rest of the economy.

Another aspect of the economic model is that the future is discounted by the pure rate of time preference δ. There is plenty of discussion in literature what exactly this parameter does and how it should be chosen. We set it to 1%/yr and use exponential discounting which has the advantage of producing time consistent decisions. See Chapter 4 for more details.

The usage of energy from fossil fuels in the production function creates emissions of carbon dioxide into the atmosphere. Through a calibrated energy balance model a temperature increase is calculated on the basis of climate sensitivity (see Andronova *et al.* (2007) for a good introduction to the concept of climate sensitivity) and transient climate response. We assume them to be perfectly correlated, an approximation as in Lorenz *et al.* (2012b) on the basis of analysis of the temperature record of the last two centuries by Frame (2005) and further aggregated in Held *et al.* (2009). Otto *et al.* (2013) argues that it is more beneficial to learn about transient climate response than climate sensitivity, but we consider them to be correlated which implies that the learning studied in this thesis can be considered as learning of both parameters. The uncertainty on climate sensitivity is set to the log-normal distribution: $\mathcal{LN}(0.973, 0.4748)$ (Wigley & Raper, 2001). We take 20 samples by dividing the distribution into 5% quantiles and taking the expected value of each quantile. The formula for this and the resulting values for climate sensitivity can be found in Appendix A.2.

In the following analysis we use three different types of simulations:

- BAU (business as usual) is a scenario in which climate change is completely ignored and there is no loss of utility from it. This counter-factual reference case is associated with the highest welfare, as the existence of climate change represents a net loss in all other scenarios.
- NOLEARN is a scenario that assumes that science is unable to find better information about the climate system in time and our decisions on how strongly to invest in mitigation have to be made with the information we have today. There is a climate problem and the decision maker performs a trade-off between cost and risk but one decision has to be made for all SOWs. The trade-off parameter β is calibrated so that this scenario exactly matches the given target.
- LEARN(2010...2075) assumes that at a specific learning point the decision maker learns the true SOW and can act accordingly. This stands for the case that research has a breakthrough in climate science and/or new information is gained from observations.

In order to compare the results from different scenarios and calculate the expected value of information, the concept of Certainty and Balanced Growth Equivalents (CBGE) is employed. This idea was originally proposed by Mirrlees & Stern (1972) and later employed by Anthoff & Tol (2009) who give a clear derivation of its application to the probabilistic case.

In general, welfare is only defined up to an affine transformation. To asses a difference in welfare a method is necessary that is invariant under an affine transformation. The CBGE method converts the values back from welfare W to consumption and thereby circumvents this problem. To compare two scenarios in which the welfare has been calculated by a CRRA utility function, Anthoff & Tol (2009) derive the following formula:

$$\Delta\mathrm{CBGE}(W, W_{\mathrm{ref}}) = \left(\frac{W}{W_{\mathrm{ref}}}\right)^{\frac{1}{1-\eta}} - 1 \qquad \text{for}\, \eta \neq 1. \tag{2.20}$$

This is calculated in the stochastic regime, i.e. W is seen as the expected welfare. In summary, ΔCBGE gives the change in initial consumption that is necessary to reach the difference in welfare assuming equal consumption growth in both scenarios.

Let LP represent the time at which information arrives (which could also be "never", i.e. NOLEARN). Then we can define two ΔCBGE quantities:

- **Expected Cost of Climate Policy** ECCP(LP): the cost created by including the calibrated penalty function and assuming perfect information in year LP. This cost includes the cost of mitigation as well as the utility loss from the risk function, i.e. the losses due to the mere existence of the climate problem. It is calculated by the percentage change of CBGE between LEARN(LP) scenario and the BAU scenario. Note that if the reference is the BAU case, the resulting value is be negative. Therefore, ECCP is defined as the negative change in CBGE:

$$\mathrm{ECCP}(LP) = -\left[\left(\frac{W_{\mathrm{LP}}}{W_{\mathrm{BAU}}}\right)^{\frac{1}{1-\eta}} - 1\right] \qquad \text{for } \eta \neq 1. \tag{2.21}$$

- **Expected Value of Perfect Information** $\mathrm{EVPI}(\mathrm{LP}_1, \mathrm{LP}_2)$: calculated by the percentage change of CBGE between $\mathrm{LEARN}(\mathrm{LP}_1)$ and $\mathrm{LEARN}(\mathrm{LP}_2)$ scenarios. Gives the value created by receiving the information in year LP_1 instead of year LP_2. The learn scenario with the latest learning point is chosen as the reference.

We point out that the equality $\mathrm{ECCP}(\mathrm{LP}_1) = \mathrm{ECCP}(\mathrm{LP}_2)$ - $\mathrm{EVPI}(\mathrm{LP}_1, \mathrm{LP}_2)$, i.e. the cost when learning at LP_1 is equal to the cost when learning at LP_2 diminished by the value of information that is created by learning earlier, only holds in approximation due to the curvature of the utility function.

2.4.2. Comparison to CEA

CEA is the standard tool that the climate target community uses to date. We offer a comparison between the CEA and CRA to clarify similarities and differences. This comparison can only be done for the NOLEARN case as learning is ill-defined in CEA.

Figure 2.3a shows temperature paths for each SOW (the CRA part repeats Figure 2.2a). We can see that CEA deviates from the path only after 2060 and in a way that would pose a considerable threat to society: in 2150 there is a 10% chance of temperatures above 3°C. CRA only shows a 5% chance with decreasing trends. Figures 2.3b-f show the important intermediate system or control variables and their effect on the carbon budget. The agreement up to 2050 is considerable which supports the usefulness of CEA simulation for near term decisions.

Figure 2.3c shows a clear difference in the emissions after 2050 although the investment into renewable energy looks to be similar in (f). The total energy produced in both simulations is the same but the balance is shifted towards fossil fuels in CEA slightly.

Figure 2.4 shows the expected discounted degree years explicitly. It becomes clear that CEA suggests a greater risk to society than CRA although they both exactly reach the calibration target. This raises the question how much worse the CEA solution is if it were evaluated in a CRA setting and conversely. This can be done in post processing and is shown in Table 2.1.

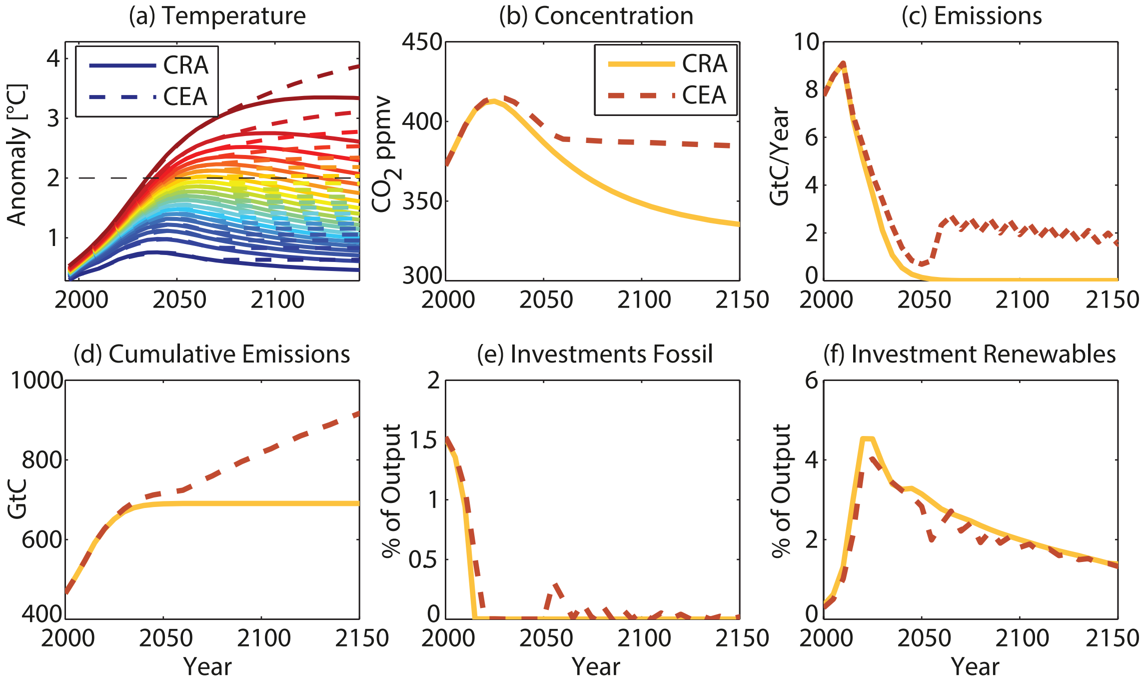

Figure 2.3. – A comparison of CEA and CRA with an equal safety (probability of staying below the target) of 66% w.r.t. a target of 2°C. The plot shows (a) the temperature paths (color indicates the climate sensitivity: low (blue) to high (red)), (b) CO_2 concentration in the atmosphere, (c) carbon emissions per year, (d) cumulative carbon emissions, (e) investment share of output into fossil fuel energy production and, (f) investment share of output into renewable energy production.

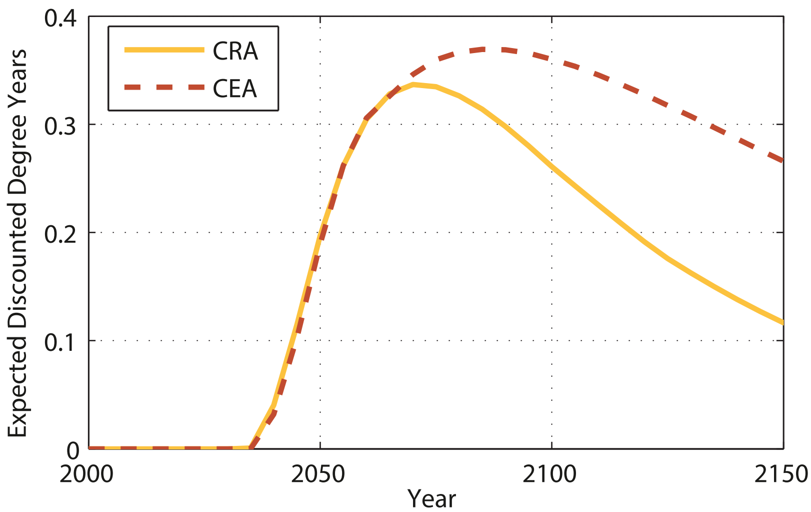

Figure 2.4. – Expected discounted degree years for CEA and CRA. We see that the risk only deviates towards the end of the simulation. Note that in both cases, the calibration target is not violated (not shown).

Scenario	Reference	CBGE change	Description
X_{CEA} in CEA	X_{BAU} in BAU	**1.3 %**	Mitigation cost for CEA
X_{CRA} in CRA	X_{BAU} in BAU	3.36 %	Total ECCP for CRA
X_{CRA} in CEA	X_{BAU} in BAU	**1.52 %**	Mitigation cost for CRA
X_{CRA} in CEA	X_{CEA} in CEA	-0.22 %	CRA decisions as sub-optimal decisions in CEA
X_{CEA} in CRA	X_{CEA} in CEA	-1.24 %	CEA decisions as sub-optimal decisions in CRA

Table 2.1. – Comparing CEA and CRA by costs and non-optimality of decisions. The pure mitigation costs (bold face) are similar for CEA and CRA. The solution found with CRA creates a 0.22% loss in CBGE if it were to be evaluated in a CEA framework compared to the CEA optimal solution. However, if CRA is taken as the reference then the CEA solution comes with a loss of 1.24% due to non-optimality with respect to the optimal CRA solution.

Assuming three different sets of optimal solutions X_{BAU}, X_{CEA} and X_{CRA} originating from the three frameworks BAU, CEA and CRA. Welfare can be calculated from any combination of solution and framework. The welfare that is calculated can be compared to any other welfare to give a CBGE percentage change. Theoretically there are 9 welfare values and therefore 36 possible CBGE changes. In Table 2.1 only the five most interesting ones are shown. It seems that the decisions found with CRA, if evaluated in a CEA framework, produce a moderate loss of -0.22% CBGE. However, if the CEA optimal solution is evaluated in a CRA framework it performs poorly in contrast to the optimal CRA solution (-1.24% CBGE). If it is unclear which framework to use to find decisions, this shows that it is better to use CRA optimal decisions as they also perform well in a CEA framework.

We calculate ECCP(never) for CEA and CRA but note that it has slightly different interpretations in each case. In a CEA it is equal to the cost of mitigation, whereas in CRA it is the cost of mitigation combined with the utility loss from risk. It is possible to estimate the pure cost of mitigation for CRA by evaluating the solution in a CEA framework (X_{CRA} in CEA compared to X_{BAU} in BAU). This results in a cost of mitigation of 1.52% CBGE. Again it is evident that CRA mitigates slightly more than CEA, but both lie in the mid range of literature mitigation costs for a 2°C target (Edenhofer *et al.*, 2010).

Recalling that the overall goal of IAMs is to advise policy makers we investigate to what extent CRA and CEA reflect the preferences of the community that supports climate targets. Most argue with a precautionary approach (Athanassoglou & Xepapadeas, 2012) or catastrophic damages approach and we infer that a resulting preference would decrease temperature in the future if possible and not too costly. If the target is reached it seems unlikely that increasing emissions around 2050, which happens in CEA, would be a supported preference. This is one of the reasons why we argue that the CRA is more adequate to represent the preferences of a climate target community.

The origin of the discrepancy between CEA and the preferences of the 2°C community probably lies in the fact that the target community usually does not mean to imply that above 2°C the world suddenly becomes unlivable but rather that the gut-feeling is that it would be beneficial to keep temperatures below this level, i.e. the feeling that if better impact models were available, 2°C maximum temperature rise could be optimal.

2.4.3. Value of Information

We originally set out to find the value of information about the climate response. Put differently: how much would we pay to resolve the uncertainty on climate sensitivity? First, a reference case has to be established and it is reasonable to assume that learning about climate sensitivity from temperature observations is always an option. It is a gradual process and Kelly & Kolstad (1999) have estimated it to take at least 90 years and Webster *et al.* (2008) find that it will take two to five decades to reduce uncertainty by 20-40%.

The necessary investment to benefit from climate observations is comparably small as measurement infrastructure is already widely spread. In a first approximation it can be seen as free information. To calculate the value of even better information about climate sensitivity we compare learning at an early stage with learning at a late stage (the reference case), where the late stage represents the learning scenario where information only originates from new observations. To test the sensitivity of the late learning assumption we conduct simulations with two reference years for late learning, 2050 and 2075, as approximations to the gradual learning over 90 years.

It appears plausible, that an un-accelerated scientific progress will reduce climate response uncertainty in a couple of decades. Hence, we assume a much smaller uncertainty in 2050 or 2075 compared to today. If the reader finds this overly optimistic, he or she may interpret the so derived EVPI as lower bounds. Furthermore, for continuous learning from observation of the temperature time series, learning in 2050 to 2075 might also be regarded to be a good approximation as an intermediate value. Because of the correlation between climate sensitivity and transient response, low values will be learned in a couple of decades due to the faster response times of the climate system, whereas, a high climate sensitivity will require at least centennial scale due to slower response times. In total, perfect learning in 2050 or 2075 seems a good reference point.

To calculate EVPI we do 13 simulations by varying LP from 2015 to 2075 and then compare it to LEARN(2050) and LEARN(2075). The results are shown in Figure 2.5. The EVPI, and hence the willingness to pay to resolve the uncertainty in 2015

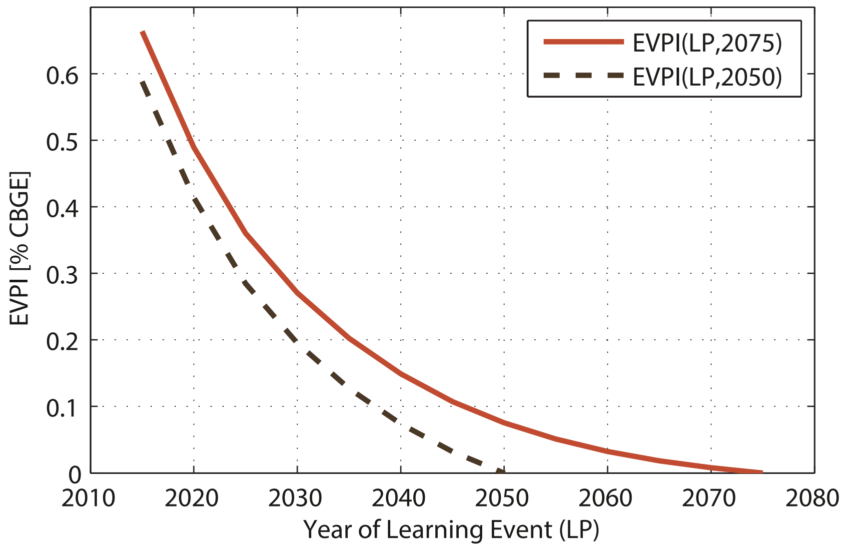

Figure 2.5. – Expected Value of Perfect Information (EVPI) given in percentage change of CBGE for different timings of the arrival of information (LP) in relation to two different reference scenarios: LEARN(2075) and LEARN(2050). The EVPI drops steeply in the first decades. Setup: guard rail 2°C, safety 66%, risk aversion 2, time preference 1%/yr.

compared to 2075, is around 0.66% CBGE. If 2050 is chosen as the reference, EVPI is around 0.60% CBGE. The figure also shows that the value of information will halve before 2030.

2.4.4. Effect of Learning

The calibration was conducted with the goal of reaching 2°C with 66% safety. This does, however, not constrain the safety in the LEARN case. By finding the number of SOWs that cross the guard rail we can imply an overall safety for the LEARN cases which lie around 70%. Due to the sampling constraint to steps of 5%, this can be seen as virtually equal to the 66% of the calibration. The equality is pleasing as it means that the result is independent of the assumption about learning for calibration, making the approach more robust. A calibration using the LEARN case instead of NOLEARN poses numerical challenges and is therefore impractical.

A plot of the cumulative distribution function for the maximum temperature in NOLEARN and LEARN(2015) cases can be found in Figure 2.6. Most of the SOWs cluster around 2°C and only one SOW does not reach 2°C. Maximum temperature above the guard rail is only reduced slightly. This hints at the fact that the NOLEARN case is already close to the maximum possible mitigation.

In Figure 2.7 the effect of a learning event on temperature, emissions and the investment into renewable energy is shown by example of LEARN(2015). Figure 2.7b shows the yearly carbon emissions and the sharp increase that is associated with learning that climate sensitivity is very low. For high climate sensitivities the emissions drop to zero in 2040 just as in the NOLEARN case shown in Figure 2.3. The investment into renewables is delayed if information arrives that climate sensitivity is low but peaks strongly at 8% if a high climate sensitivity is learned.

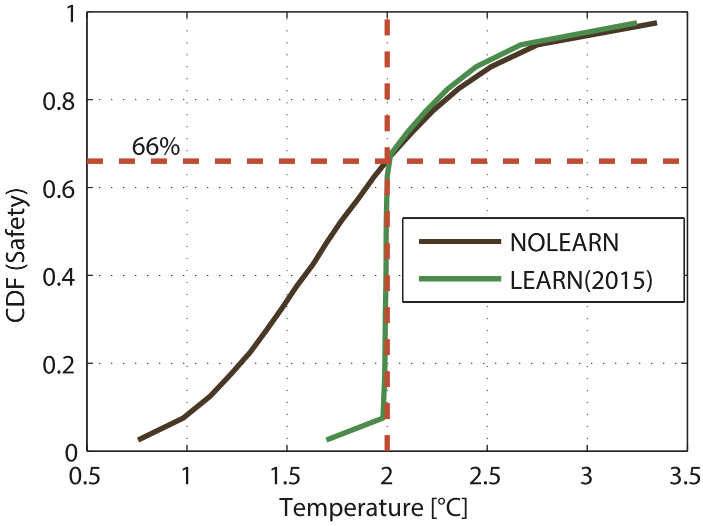

Figure 2.6. – The cumulative density function (or safety) for NOLEARN and LEARN(2015) scenarios. After learning, most SOWs cluster around 2°C which is the target that was used for calibration.

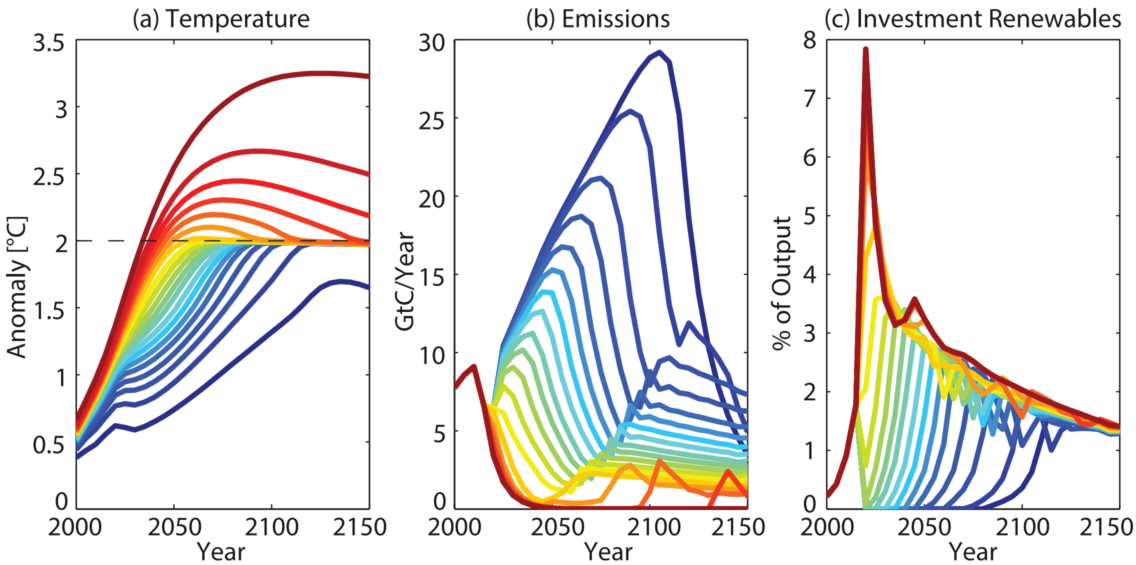

Figure 2.7. – (a) Temperature, (b) emissions and (c) investment into renewables for LEARN(2015). The colors indicate climate sensitivity going from low (blue) to high (red). Each line has a probability of 5%.

2.5. Discussion

It is unlikely that impact models will improve sufficiently within the next decade (even more so considering the ethical problems) to allow the climate target community to accept CBA as *the* tool for policy analysis. However, the alternative tool, CEA, has serious drawbacks as it does not represent the true preferences of a proponent of climate targets under anticipated learning. In CEA there are incentives to prevent research as the value of information about the climate response can be negative. Furthermore, there are incentives to fire up the fossil fuel industry once the target is reached (around 2050), also not inline with preferences.

As a remedy we apply a cost-benefit trade-off but instead of the classic production damages we include a penalty that affects the utility directly and depends on the expected transgression of the guard rail, an idea that has been proposed by Schmidt *et al.* (2011) and in other contexts as positive mathematical programming (Howitt, 1995) or internalization of cost targets (Bordley & Pollock, 2009). We construct the trade-off in a way that the preferences of the climate target community is reflected without imposing a constraint but by internalizing them. The resulting solution diverges w.r.t. CEA only after 2050 allowing CRA to substitute CEA without invalidating previous research with CEA.

By applying CRA we return to the expected utility framework and are able to calculate the value of information about the climate response. We found the value of information to be around 0.66% of consumption per year if learning in 2015 is compared to learning in 2075. This value makes up around 1/5 of the total expected costs of climate protection. This is a substantial value that can potentially be reaped with adequate investments into climate response research.

To convey the size of 0.66% CBGE we recall that it represents a change in consumption per year. In 2012 the world had a GWP of around 70 trillion USD and a consumption share of 75% (this coincides roughly with the consumption share in MIND for the reference case as well as for the world wide consumption share estimated by The World Bank (2013)). This leads to an absolute value of information of 347 billion USD. Even a reduction by an order of magnitude still leaves a considerable value that is by far larger than any climate science research budget.

Previous publications have calculated the value of information about the climate in different ways and with different metrics, e.g. using net present value. Using a CBA with a damage function, Nordhaus & Popp (1997) find the value of learning in 2015 instead of 2045 about the climate feedback to be 7.81 billion 1990 U.S. dollars. The optimal strategy found by Nordhaus & Popp (1997) yields a temperature rise of about 3°C with a probability of 66% and using a similar target for calibration we find around 0.086% CBGE for the EVPI(2015,2045) (not shown in the analysis above).

To convert a change in CBGE into a change in net present value we use a finding from Held *et al.* (2009) and divide the value by 1.5. In 1995, which is the reference year in both studies, the GWP was 33.6 trillion 1990 USD. This yields a net present value of 19.4 billion 1990 USD for the CRA, which is on the same order of magnitude as the value from Nordhaus & Popp (1997). This underlines the fact than an inclusion of the climate effects as an additive part to utility results in a similar valuation of information. Keller *et al.* (2007) come to the conclusion that the value Nordhaus & Popp (1997) found is increased by more than an order of magnitude if an uncertainty about the existence of a temperature threshold of 2.5°C is assumed. Along similar lines, Lempert *et al.* (2000) find that it is worth 15-22 billion USD a year to know if society has to mitigate a lot or a little.

In the context of global observation systems, Cooke *et al.* (2013) calculates the avoided damages if a better observation system is put in place, allowing to reduce the uncertainty on climate sensitivity. The authors find an extreme value of ca. 3 trillion USD which stems from the fact that they assume a business as usual case until uncertainty has reduced substantially. Another study that calculates the value of information about climate sensitivity gives a net present value of 23.9 billion 2005 USD for learning in 2020 compared to never (Webster *et al.*, 2008). This is about an order of magnitude less than the value calculated in this thesis due to the fact that the CBA used in their work, produces higher temperatures than CRA with a 2°C and 66% target.

The next chapter delves deeper into the workings of CRA with some sensitivity studies and a closer look at the origin of EVPI. Questions about the shape of the risk metric, the effect of anticipation or different calibration targets are analyzed. The main focus is on the origin of the value of information.

3. Disentangling Contributions to the Value of Information

3.1. Introduction

Cost-Risk Analysis (CRA), as presented in this thesis, has not been explored thoroughly in literature yet. As it is the first time it is applied to an integrated assessment model (IAM), it is of special benefit to do a thorough sensitivity analysis and investigative simulations to understand the dynamics behind such a trade-off. The crucial difference to Cost-Benefit Analysis (CBA) is that – due to the calibration to a desired temperature target – the objective function is adjusted anytime a parameter that influences the optimal solution in the NOLEARN scenario is changed. How and why will be explored in this chapter.

The welfare equation in CRA consists of two terms that are added, one for the economy and one for climate effects. The fact that this is done in an additive way, creates a separability of utility changes into economy or risk-related utility changes. This in turn allows for an attribution of effects to one of the two mechanisms. However, the situation is complicated by the non-linear utility function making the influence of a consumption change depend on the level of consumption. Because the unit of measurement is given in certainty and balanced growth equivalents (CBGE, see Section 2.4.1), which is a unit of consumption change, one has to be careful when splitting CBGE quantities. This division is explored in this chapter and used as a tool to analyze the results of CRA.

The goal of this chapter is to gain insight into CRA with MIND-L by changing a variety of parameters:

- Calibration – previously we chose 2°C with a safety of 66% as a calibration reference. The effect of this assumption is tested by varying guard rail and safety. These two variations produce similar results but have a key difference when learning is considered.

- Risk metric – in the previous chapter we found that a threshold-linear risk metric is identical to the concept of degree years and represents the boundary case at which a sacrificing of the climate is prevented. We test the effect of this choice by comparing the results with simulations that use different risk metrics. We represent risk metrics by a penalty function per time slice on temperature and vary the functional form from logarithmic to fourth order functions.

The last analysis considers the assumption that the decision maker knows the point of learning (anticipated learning). In literature the anticipation effect is calculated to be small (Lorenz *et al.*, 2012b) and we find this to be true for CRA and MIND-L as well.

There are of course many more parameters that could be changed, for example: the class of utility functions, productivity factors, exogenous population, other greenhouse gases, learning rates for renewables, resource base for fossils, learning rates for efficiency and many more. The focus of this thesis is on CRA and so aspects are analyzed that are of interest with respect to the decision framework.

To be able to assess the differences that occur by the variations of the above points we use the fact that the utility is separable in time, states of the world (SOWs) and also in risk and cost (referred to by "type"). The first part of the chapter assesses the origin of expected value of perfect information (EVPI) by looking at the reference case: EVPI(2015,2075). In Section 3.3.1 the calibration is varied and the effects are studied. In Section 3.3.2 we move on to explore the effect of using other penalty functions for the risk metric. The last analysis concentrates on the effect of anticipating the learning event. We close with a summary and discussion of the results.

One important aspect that has been left out in this chapter is the effect of other normative choices apart from calibration, such as risk aversion and discounting. The next chapter is devoted entirely to these choices on which there is a lot of controversy and discussion in the literature.

3.2. The Reference Case

Our starting point for this chapter is the model setup that was described in Chapter 2. The key parameters of the setup are summarized in Table 3.1. The goal of this section is to dissect the expected cost of climate protection (ECCP) and the EVPI that was calculated in the previous chapter and find the origin and reasons for their values.

A first step into a deeper analysis of this reference case can be made by plotting the risk and the cumulative emissions over time. We represent the risk by plotting expected discounted penalties (which, aggregated over time, form the expected

Parameter	Value
Constant relative risk aversion η	2
Pure rate of time preference δ	1%/yr
Penalty function	$\Theta\left(T(t)-T_{\mathrm{g}}\right)\left(T(t)-T_{\mathrm{g}}\right)$
Calibration safety	66%
Calibration guard rail	2°C
Calibration information structure	no information
EVPI reference year	2075
Learning type	anticipated

Table 3.1. – Setup of the reference model that serves as a basis for comparison to the sensitivity analysis in this chapter.

discounted degree year measure). Figure 3.1 shows a plot of these variables and illustrates that there is a benefit on the risk side as well as on the economic side from learning earlier. The risk is only slightly reduced by learning and most of the reduction is after 2050. Cumulative emissions start to differ earlier, this is of course due to the effect of the inertia of the climate system.

Figure 3.1b can be seen as a proxy for mitigation costs. If more emissions are allowed it will cost less for the economy because fossil fuels are cheaper. Therefore, the source of part of the value of information is the possibility of emitting more for SOWs with low climate sensitivity if information arrives early. SOWs with high climate sensitivity will stick to the lower bound of cumulative emissions, bringing almost no economic value of information but instead reducing the risk slightly as shown in Figure 3.1a. It is not clear how much of the EVPI can be attributed to the risk reduction and how much to reduction in mitigation costs. Furthermore, it is not easy to see which of the SOWs contribute how much.

To gain more insight into the reasons for EVPI and the workings of the model, the aim is to disentangle the contributions to the EVPI in three dimensions: type (risk or economic), SOW and time.

3.2.1. Dividing Welfare Changes

The biggest challenge when dividing welfare a change into its components is that any calculation of ratios is not sensible due to welfare being only defined up to affine transformations. If the welfare measure is converted back into an equivalent consumption this problem is avoided. The EVPI (as well as the ECCP, or any CBGE changes) is calculated by comparing welfare, or more specifically, their equivalent initial consumptions. To find a meaningful disentangling of the components that make up the calculated CBGE change, we first argue that the calculation of CBGE for CRA is at all sensible.

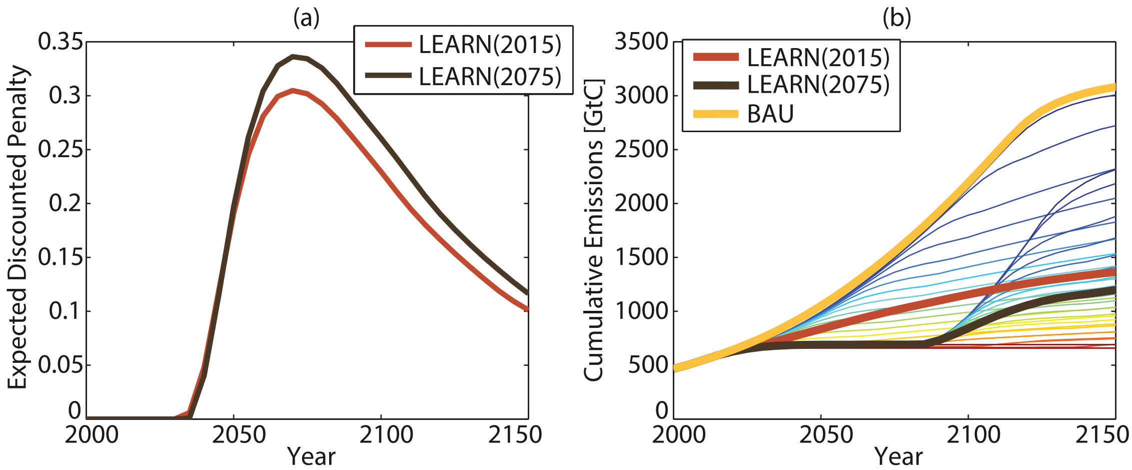

Figure 3.1. – (a) Expected discounted penalty and (b) cumulative emissions for relevant scenarios of the reference CRA setup. Expected discounted degree years are calculated by the area underneath the curves in (a). Panel (b) also shows cumulative emissions for each state of the world: red to blue thin lines represent decreasing climate sensitivity. If perfect information arrives in 2015 instead of 2075 the risk is decreased as shown in (a) and the emission can be increased on average (see b). If very low values of climate sensitivity are learned then emissions can roughly follow the BAU case if learning is in 2015 which produces a large economic benefit.

Calculating the CBGE for a CRA simulation assumes that the utility loss created by the risk-related part can be expressed as a consumption reduction. Therefore, the calculation of a balanced growth equivalent does not just balance ragged growth rates but also includes a monetary loss from the climate risk. The CBGE consumption level for a simulation (calculated from its welfare) will be smaller than the initial consumption level calculated from balancing the consumption path in CRA. This process therefore incorporates the risk as an equivalent reduction in initial consumption.

Once a CBGE value for CRA is accepted as a sensible measure, the next step is to compare CBGE values. In this analysis we only compare scenarios which are very similar. The CBGE change is usually below 5% and the absolute value for the consumption streams are very similar across scenarios. This allows a linearization of the CBGE formula.

The CBGE change between two scenarios a and b with a separable welfare $W_a = \sum_j W_{a,j}$ can be calculated by the CBGE formula for CRRA utility functions:

$$\Delta = \left(\frac{W_a}{W_b}\right)^{\frac{1}{1-\eta}} - 1 = \left(\frac{\sum_j W_{a,j}}{\sum_j W_{b,j}}\right)^{\frac{1}{1-\eta}} - 1. \tag{3.1}$$

Here the index j stands for any division in welfare parts that makes sense, for example: time slices, SOWs or economic and risk related parts. To find the contribution of each summand j to Δ, the equation is linearized at $W_a/W_b = 1$. The linearization reads:

$$(x)^{\frac{1}{1-\eta}} - 1 \approx \frac{1}{1-\eta}(x-1). \tag{3.2}$$

The CBGE change calculated by the linearized equation is denoted by $\tilde{\Delta}$:

$$\tilde{\Delta} = \frac{1}{1-\eta}\left(\frac{W_a - W_b}{W_b}\right) \tag{3.3}$$

This can be split into the components denoted by the index j:

$$\tilde{\Delta} = \frac{1}{1-\eta}\left(\frac{\sum_j W_{a,j} - W_{b,j}}{W_b}\right), \tag{3.4}$$

and for each component we have:

$$\tilde{\Delta}_j = \frac{1}{1-\eta}\left(\frac{W_{a,j} - W_{b,j}}{W_b}\right). \tag{3.5}$$

When comparing two scenarios there is a difference between $\tilde{\Delta}$ and Δ. To ensure that the overall CBGE change continues to be Δ, a re-scaling is done:

$$\Delta_j = \Delta\frac{\tilde{\Delta}_j}{\tilde{\Delta}} \tag{3.6}$$

Inserting equation 3.3 and 3.5 into the above equation leads to the split of the CBGE change in proportion to the differences of the welfare summands:

$$\Delta_j = \Delta\frac{W_{a,j} - W_{b,j}}{W_a - W_b}. \tag{3.7}$$

This can be applied to any sensible division of welfare as long as W_a and W_b are similar. To assess the error we make by the linearization we look at the difference between $\tilde{\Delta}$ and Δ. This difference points to the order of magnitude of the error. Over all learning scenario comparisons and applied normative parameters (see Chapter 4) the maximum error is 1%, i.e. 0.0066 percentage points for a 0.66% CBGE change.

The error is always negative due to the curvature of the utility function but is compensated by the rescaling. See the Figure A.3 in the Appendix for a plot of the error. In summary, it is safe to use this approximation for the cases analyzed in this thesis.

3.2.2. Origin of EVPI and ECCP

With this tool of decomposition it becomes possible to assess if the EVPI originates mostly from increased consumption or from decreased risk. To do this we divide the welfare calculation into two terms, one is calculated by the utility function from the consumption stream (W_{U}) and the other is calculated by the risk metric and multiplied with the trade-off parameter (W_{R})[1]:

$$W = W_{\mathrm{U}} + W_{\mathrm{R}} \tag{3.8}$$

Once this separation is established, equation 3.7 can be applied to divide the EVPI into economic and risk-related terms. Figure 3.2 shows the division of EVPI(LP,2075) for different LP. The suspicion that the economic part is more important is confirmed as it contributes most to total EVPI. After 2040, the risk-related contribution can even be negative, implying that temperatures were increased after learning which then allowed for less mitigation and a higher contribution from the economic part. Note that if information arrives in 2020 the risk-related value drops considerably and is less than a fourth of the economic value. This points towards the fact that temperatures can only be influenced effectively now.

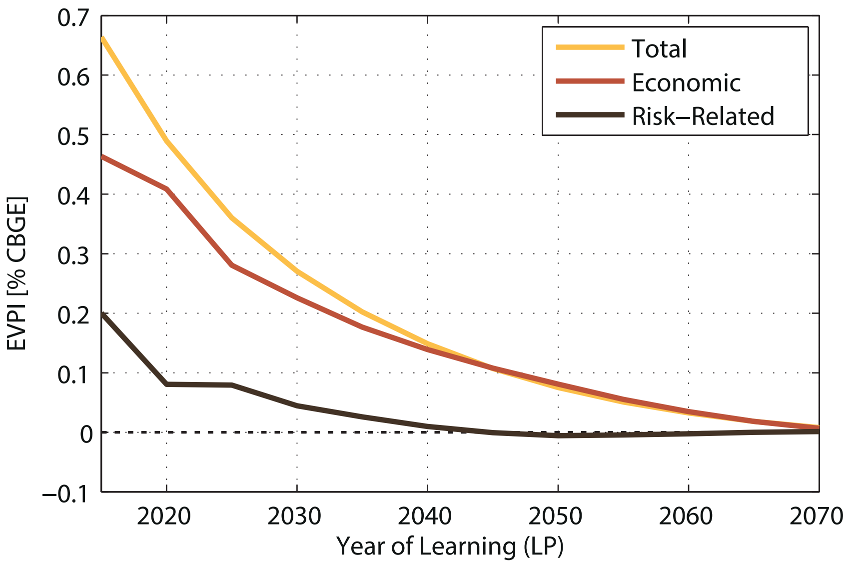

Figure 3.2. – EVPI(LP,2075) split into economic (red) and risk-related (black) parts for different learning points LP.

In addition to the division into economic and risk-related values, we further divide the EVPI along time (Figure 3.3 a, b, c) and SOW (Figure 3.3 d, e, f). The color coding indicates the learning time point, where darker colors stand for earlier learning. We make the following observations in Figure 3.3:

[1] The equations for the parts can be found by referring to equations 2.15, 2.18 and 2.19.

- As learning occurs later, the time at which EVPI originates also occurs later (see (a)).
- Generally, the risk related EVPI originates towards the end of the simulation whereas the economic EVPI is mostly created in the beginning of the simulation. (compare (b) and (c)).
- SOW 14 to 16 (of a total of 20) do not contribute to the EVPI. If we learn that we are in these SOWs, we would not adjust our emission strategy compared to the NOLEARN case (see (d)).
- For low climate sensitivities (SOW smaller than 14) EVPI results only from the economy. For high climate sensitivities, the trade-off is made between cost and risk: any increase in emissions would reduce the loss in (e) but would be overcompensated by a reduction of the gain in (f).

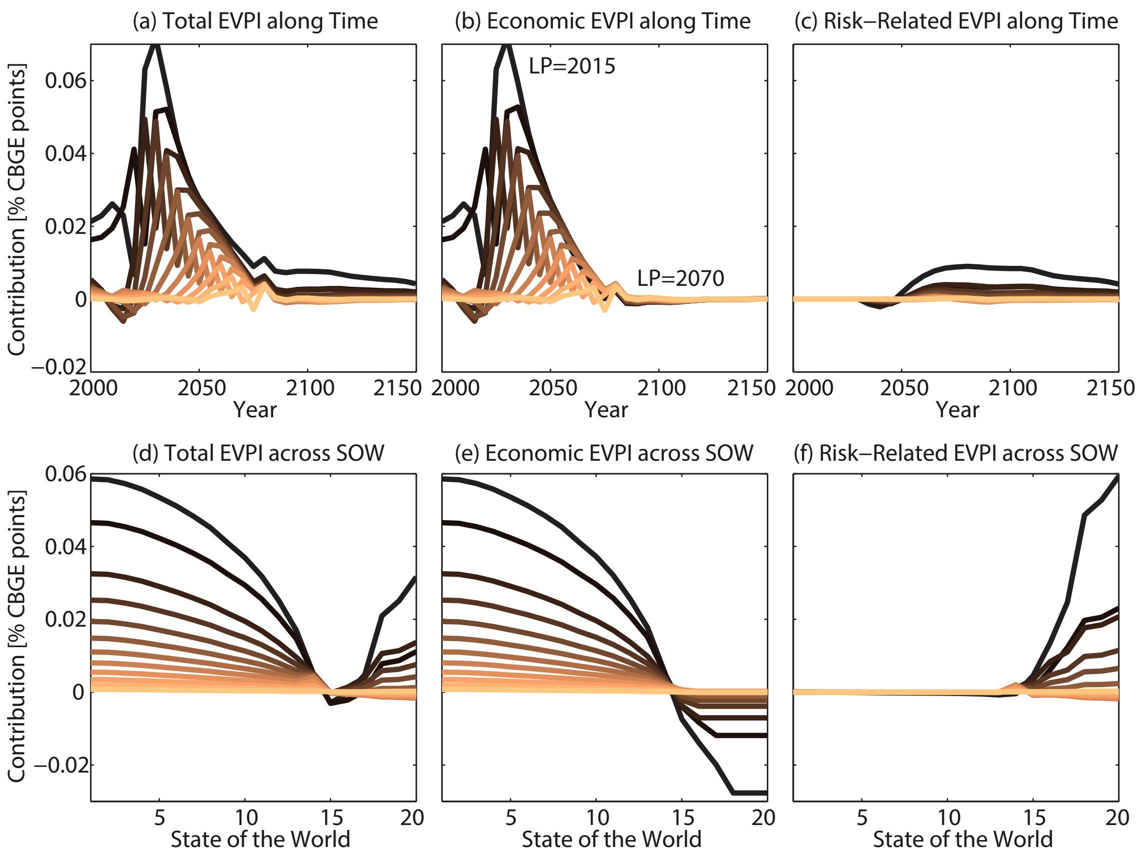

Figure 3.3. – EVPI(LP,2075) where LP is indicated by the color (yellow: 2070, black: 2015) and split up in type, time and states of the world (SOW). (a) shows a split along time, (b) split along time for the economic part, (c) split along time for the risk-related part. (d) shows the split across SOW in total and (e,f) show the economic and risk-related parts. As learning happens earlier the effects are larger. See Appendix A.2 for climate sensitivity values of each SOW.

Because EVPI reduces the ECCP (i.e. learning reduces the cost and risk), it is interesting to look at ECCP to see what the underlying costs and risks are and how they are reduced by learning. If costs or risks are high, then there is high potential for reduction, therefore we would assume that the patterns of ECCP are similar to EVPI.

Figure 3.4 shows the economic and risk-related ECCP split along time and SOW. A similar pattern as the EVPI can be seen. The time shift between the two types of contribution is clearly evident. The figure is constructed from the LEARN(2075) and shows that the economic mitigation costs (yellow) originate from all SOWs equally as mitigation has to be done before learning. The risk originates only from the upper seven SOWs which violate the guard rail[2].

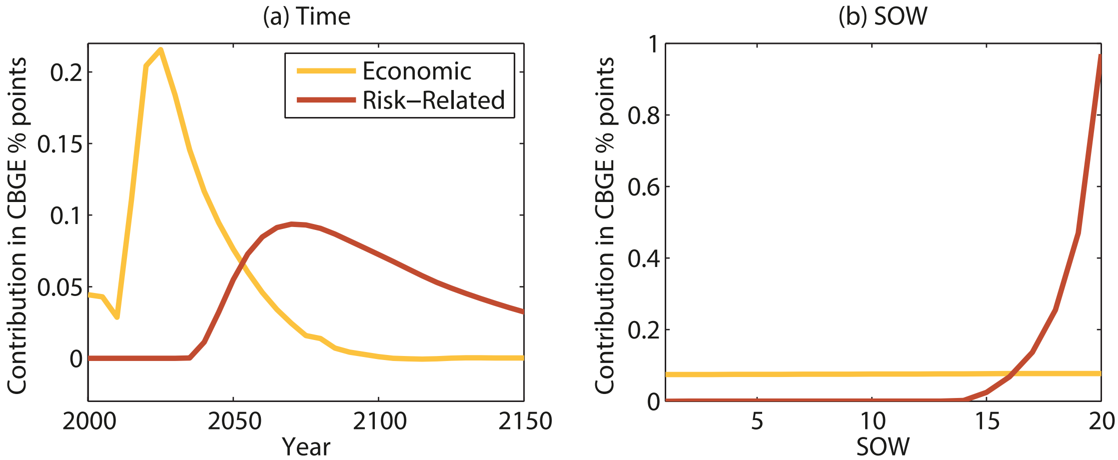

Figure 3.4. – ECCP(2075) divided across (a) time and (b) SOW. The time shift between risk-related and economic costs can be seen in (a) and the fact that all SOWs contribute to the economic mitigation costs although only the upper five SOWs create the risk is shown in (b).

3.2.3. Classifying States of the World

Another tool that can help to assess a CRA simulation is classifying the SOWs into different cases with different binding optimal emission constraints. This classification only makes sense in the deterministic case in which the optimal emissions depend on the true value of the climate sensitivity, i.e. after the learning event. A full derivation of these regimes for the static problem can be found in Chapter 5. We consider two constraints: a maximum and a minimum possible emission strat-

[2]Its worth mentioning here that even though the risk seems to diverge for high climate sensitivities, the area under the curve converges. In other words, the expected risk is finite even if infinite climate sensitivity is considered. Therefore, the dismal proposition discussed in Weitzman (2009) does not apply here. Appendix A.1 argues this is more detail.

egy, representing emissions all available fossil fuels are burnt and maximum viable mitigation respectively. Four regimes can be identified:

1. Optimal emissions would lie above $E_{\max}$ so the optimal course of action is the same as in BAU.
2. It is cheap enough to stay below the guard rail so that trespassing is not viable. Emissions are chosen such that the maximum temperature lies exactly on the guard rail.
3. Optimal emissions are such that temperatures pass the guard rail and a trade-off between cost and risk is made.
4. Optimal emissions would lie below $E_{\min}$ so the optimal course of action is maximum possible mitigation.

These regimes, or cases, can be found in MIND-L by analyzing temperature and emission paths. For LEARN(2015) in the reference case the division of probability of each of the four cases is as follows, in the same order as above: 5% – 65% – 16% – 14%. The biggest share represents Case 2 in which the temperature rises to the guard rail but does not cross, as can also be seen in Figure 2.6. An implication of this analysis is that Cases 1 and 2 allow for economic value of learning whereas Cases 3 and 4 allow for risk-related value.

3.3. Varying Key Parameters

3.3.1. Varying the Calibration

A decision with far reaching implications is the choice of target to calibrate to. Previously we calibrated to a widely discussed target: keeping the probability of staying below 2°C temperature anomaly at 66% (safety). There are of course approximately equivalent targets that reach the same goal by adjusting guard rail and safety along the black line in Figure 2.6 of the previous chapter. In this section we do not investigate these changes but rather show how changing the stringency of a target affects the solution as well as the EVPI. We distinguish increasing and decreasing stringency for the guard rail and safety, where increasing stringency implies decreasing guard rails and increasing safeties.

For the NOLEARN case there is a duality between these possibilities of variation shown graphically in Figure 3.5. For every increase in guard rail there is an equivalent decrease in safety target that produces the same emission strategy and temperature paths. For the learning case this does not hold, therefore we perform both variations with dual values for comparison. The safety is varied between 40% and 70% where the upper bound is given by feasibility limits of the model. In a second simulation the guard rail is varied between the equivalent temperatures 2.58°C and 1.93°C. To

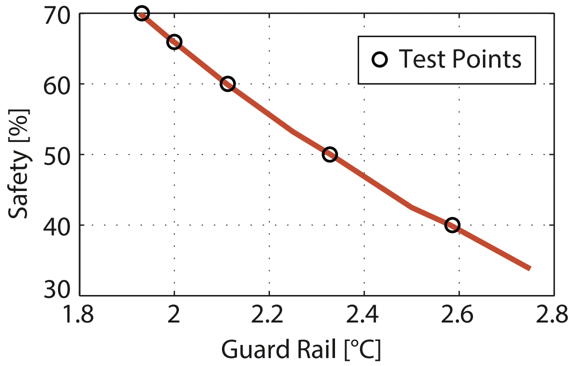

Figure 3.5. – The equality between changing the guard rail and keeping the safety at 66% or changing the safety and keeping the guard rail at 2°C for the NOLEARN case. The test points analyzed are marked by circles.

assess the difference of these variations and the corresponding effect we plot the EVPI for different learning points in Figure 3.6.

For very early learning (e.g. LEARN(2015)) the EVPI decreases as the target becomes less stringent. It decreases more strongly when adjusting the safety instead of the guard rail. For learning after 2030, the EVPI increases in Figure 3.6a and decreases in 3.6b as the target becomes less stringent. In general it is expected that as climate change becomes more of an issue, i.e. targets become stricter, the value of information increases, however, in the case of guard rails and learning in 2040, the opposite is the case.

If given enough time, temperatures gravitate towards the guard rail after a perfect learning event. This makes it very clear that adjusting the guard rail reflects different preferences than adjusting the safety target, even though there is an equivalence in the NOLEARN case .

The question as to how targets should be formulated might arise at this point. Is it better to adjust targets by changing the guard rail or by adjusting the safety target? This question only becomes important if learning is considered and then leads to the conclusion that decreasing safety would decrease the incentive for research into the climate response whereas increasing the guard rail would increase the incentive (for learning after 2030). This increase comes at the cost of higher temperatures in general after learning as can be seen in Figure 3.7. Policy makers have to be aware of the implications when adjusting targets.

The origin of the opposing EVPI(2040,2075) trend for adjustments of guard rail and safety can be found by examining the probability shares of the four regimes described in Section 3.2.3 as well as the origin of the EVPI. In Figure 3.8 we show the division into the four regimes as well as the divisions by type of EVPI for both variations.

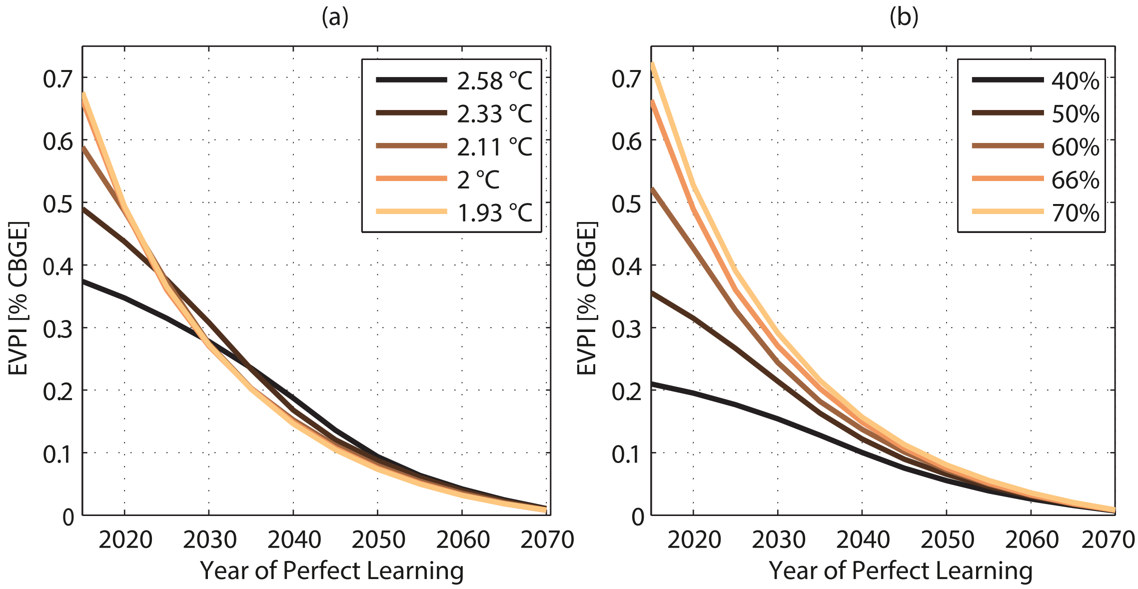

Figure 3.6. – EVPI for different years of perfect learning and different calibration points. Panel (a) keeps the safety at 66% and varies the guard rail whereas panel (b) keeps the guard rail constant at 2°C and varies the safety. A difference can be seen especially for mid-range learning where the EVPI decreases with increasing stringency of the guard rail.

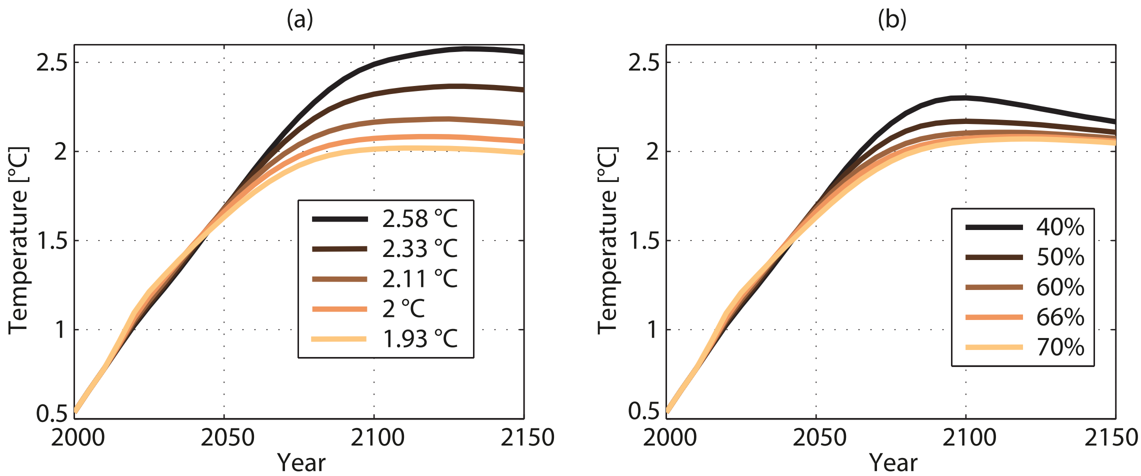

Figure 3.7. – Expected temperature for LEARN(2015) over time for varying (a) guard rail and (b) safety target. Relaxing the guard rail results in slightly higher temperatures than relaxing the safety target.

We make the following observations:

- Increasing the guard rail only changes the shares of the cases slightly. The biggest change is that Case 1 gains in share which means that in more and more SOWs it is possible not to mitigate, without crossing the target. (see Figure 3.8a)

- Increasing the guard rail hardly affects the economic contribution to EVPI, clearly evident from the red solid line in (c) but also implied by the relatively constant sum of the shares of Case 1 and 2 (which are responsible for positive economic value).
- Decreasing the safety in (b) changes the shares of the cases considerably, mainly by reducing the share of Case 2 and increasing shares of Case 3 and 4. The more SOWs enter the regime of a trade-off and pass the guard rail, the more the economic value is reduced. Note that the probability of passing the guard rail in the LEARN case does not have to be the same as the Safety, i.e. the boundary between Case 2 and 3 is not equal to the dotted line in (a) and (b) (see analytical derivation in Chapter 5).
- In (c) the switch in trend for the total EVPI arises from the fact that the economic EVPI for a safety change decreases and for a guard rail change stays relatively stable. This is mainly because increasing the guard rail allows more SOWs to do no mitigation (Case 1 (a)).

One would assume that as the target for calibration is relaxed, the cost of climate protection (ECCP) would decrease and with it also the value of information. But the analysis shows, this is only true for a relaxation of the safety target.

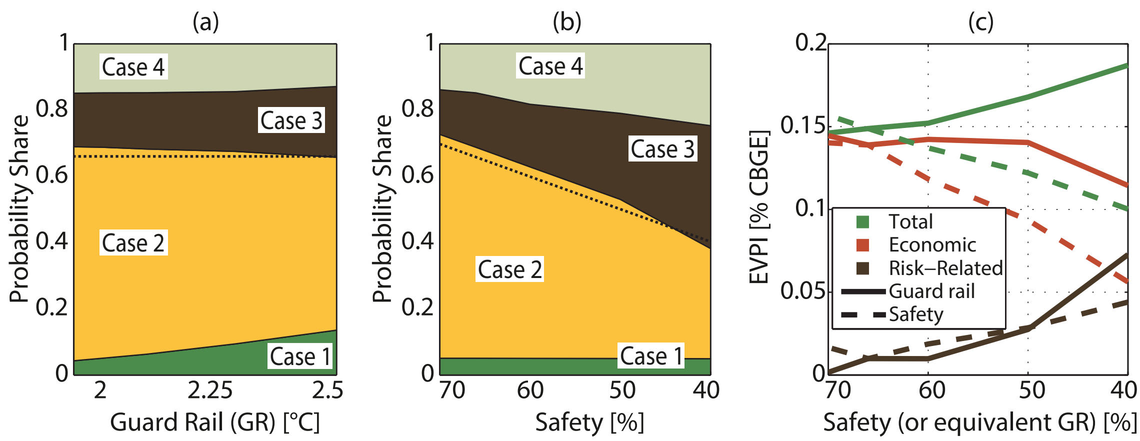

Figure 3.8. – The probability share of the four cases for SOW described in section 3.2.3 for LEARN(2040) and different calibration targets: (a) varying the guard rail and (b) varying the safety. (c) shows the EVPI(2040,2075) split into economic and risk-related value for variations of guard rail and safety plotted on the same axis. Equal at 66% and 2°C. A systematic difference is visible between the two methods to adjust targets, especially considering "Case 2". The dotted line in (a) and (b) show the safeties for the NOLEARN case.

3.3.2. Varying the Risk Metric

A seemingly important choice in setting up the CRA is the risk metric that converts temperatures into risk. The specific function that is implemented per time slice is called the "penalty function". However, we show here that the risk metric plays only a minor role, at least for the value of information. The underlying goal is to stay close to the concept of CEA, having no penalty for temperatures below the guard rail. Only the functional form of the penalty function above the guard rail is varied. In previous sections, we use a linear function, foremost because it resonates well with publications on the concept of degree years by Schneider & Mastrandrea (2005) and because it is the limiting case of a convex penalty function.

To assess the effect of varying curvatures of the penalty function, we use four additional forms, ranging from a logarithmic form to a fourth order polynomial are shown in Table 3.2. Note that we do not investigate more concave functions than the logarithmic form as they produce the sacrificing phenomenon described in Section 2.3. The logarithmic function is analyzed to show that for strongly convex cost functions, as in MIND-L, even partly concave functions can prevent sacrificing. Due to the kink at the threshold at 2°C the function is partly convex and partly concave.

Shorthand	Penalty Function
1. log	$\Theta\left(T-T_{\mathrm{g}}\right)\log\left(T/T_{\mathrm{g}}\right)$
2. lin	$\Theta\left(T-T_{\mathrm{g}}\right)\left(T-T_{\mathrm{g}}\right)$
3. T^2	$\Theta\left(T-T_{\mathrm{g}}\right)\left(T^2-T_{\mathrm{g}}^2\right)$
4. T^3	$\Theta\left(T-T_{\mathrm{g}}\right)\left(T^3-T_{\mathrm{g}}^3\right)$
5. T^4	$\Theta\left(T-T_{\mathrm{g}}\right)\left(T^4-T_{\mathrm{g}}^4\right)$

Table 3.2. – The five different functional forms of the penalty function to find the effect of the risk metric on EVPI. $\Theta(x)$ denotes the Heaviside function of x, T is the temperature of the time slice and SOW and T_{g} is the specified guard rail.

After every change to the risk metric, a recalibration is necessary to ensure that the target is met in the NOLEARN case, as discussed in Section 2.3. The process of recalibration can be formulated in terms of the static model by equation 2.6, reprinted here for convenience:

$$\beta_{\mathrm{cal}}=-\frac{C'(E_{\mathrm{g}})}{\frac{d}{dE}\mathrm{E}\left[R(\gamma E_{\mathrm{g}})\right]}. \tag{3.9}$$

Considering a different penalty function R_2 a recalibration is necessary, if the same E_{g} is to be optimal:

$$\beta_{2,\mathrm{cal}}=\beta_{\mathrm{cal}}\frac{\frac{d}{dE}\mathrm{E}\left[R(\gamma E_{\mathrm{g}})\right]}{\frac{d}{dE}\mathrm{E}\left[R_2(\gamma E_{\mathrm{g}})\right]}. \tag{3.10}$$

In summary, it is an adjustment of the marginal benefit from climate protection so that the marginal benefit at the target optimal strategy is identical for both penalty functions. This also implies that adding a constant to the penalty function does not change the calibration.

In the dynamic setting the situation is more complicated. The recalibrated penalty functions are shown in Figure 3.9b. Note that they all have similar derivatives around a temperature of 2.5°C as shown in Figure 3.9c.

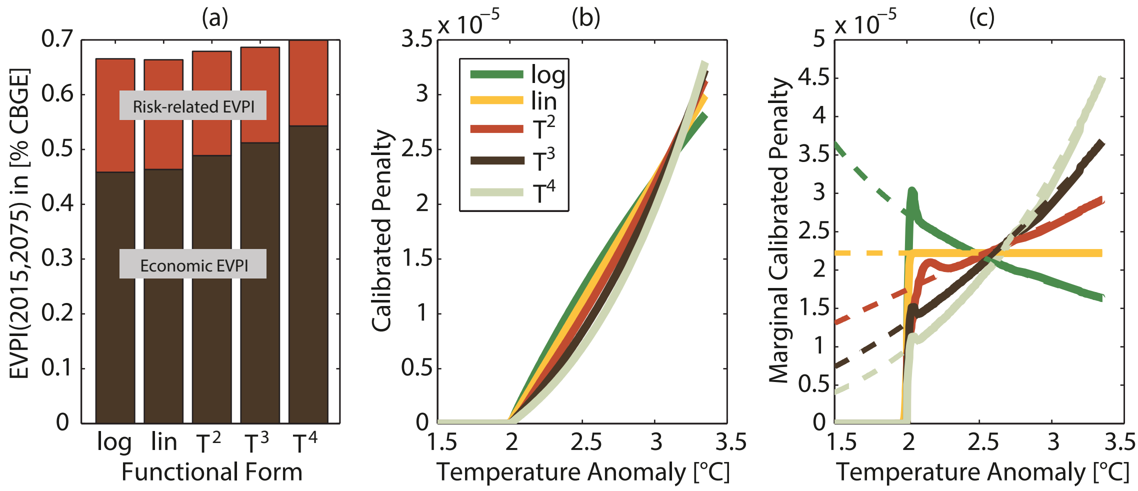

Figure 3.9. – Pane (a) shows the EVPI split into economic and risk-related values for different functional forms of the penalty function. Pane (b) shows the different penalty functions after being recalibrated and (c) show the marginal penalties. The dashed lines indicate the analytical solution without the Heaviside function and the solid lines are created by finite differences from the simulation data. Any deviations come from necessary smoothing of the discontinuity at 2°C. The calibration has the effect that the marginal penalty between 2.5°C and 2.6°C is roughly the same for all functional forms.

The effect of changing risk metrics is plotted in the bar chart in Figure 3.9a. The EVPI is only changed marginally, which is at first surprising as one would expect the EVPI to increase if the power of the penalty function is increased as temperature increases become more dangerous (i.e. catastrophic). The reason for the rather stable EVPI is the compensation of economic and risk-related EVPI changes. The economic part of the EVPI increases whereas the risk-related part decreases.

The origin of this compensating behavior can be found by looking at the ECCP and EVPI split across SOW shown in Figure 3.10. The economic ECCP is constant across SOW and the same for all risk metrics as shown in (a). This is expected because re-calibration ensures that the same amount of mitigation is optimal so that the same target can be reached across all risk metrics. In Figure 3.10b the risk-related ECCP is plotted showing that the risk is generally lower for higher convexity. The reason for this is that the absolute risk is lower for higher convexity between 2°C and 3°C as shown in Figure 3.9b.

Dividing the EVPI(2015,2075) across SOW reveals that risk-related EVPI follows the behavior of ECCP and decreases with increasing convexity, shown in 3.10d. However, this is compensated by an increase in economic EVPI especially for SOW 15 (see Figure 3.10c). Because of the decreasing marginal penalty below 2.5°C (as convexity increases, see Figure 3.9c), temperatures are allowed to grow a little larger if the convexity is higher (counter-intuitively), especially for SOWs falling into Case 3, i.e. in the regime where the temperature just crosses the guard rail and the optimal trade-off between cost and risk is made. This slight increase in temperatures allows for more economic activity and therefore more economic EVPI (see Figure 3.10c).

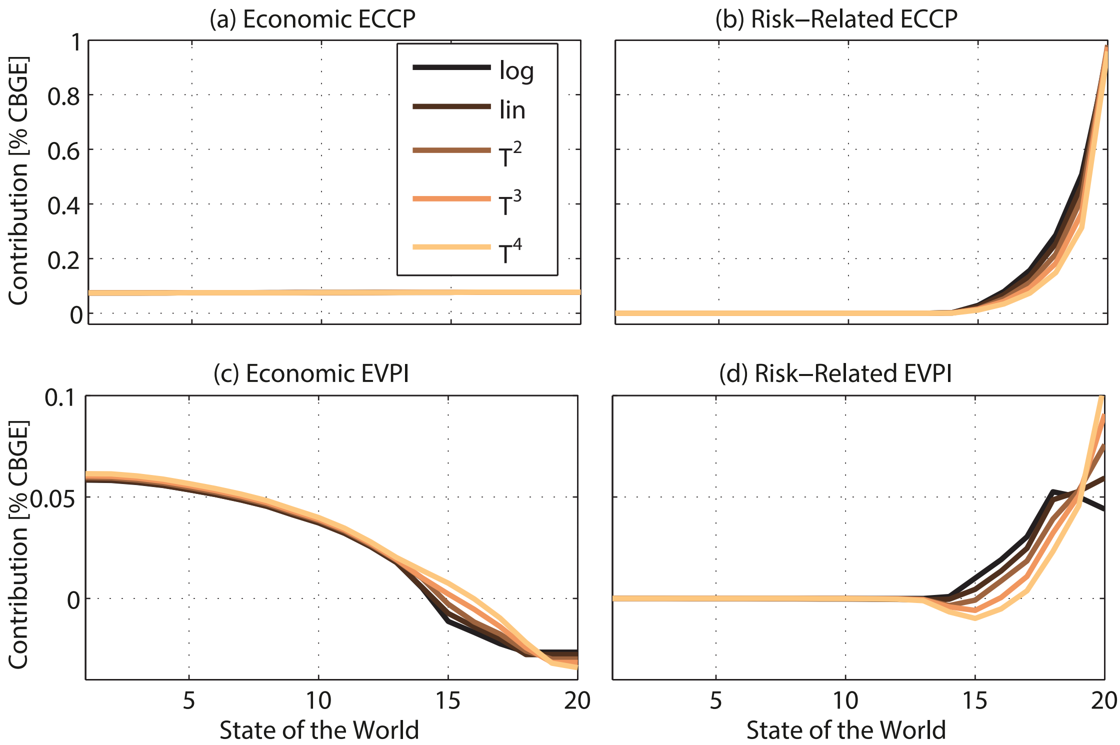

Figure 3.10. – The ECCP(2075) and EVPI(2015,2075) split along type and SOW dimensions for different types of penalty function. Lighter color indicates higher convexity.

3.3.3. Effect of Anticipation

By optimizing a dynamic model the optimal setting for all control variables in all time steps is found simultaneously. This is called a "one shot" optimization and implies that the decision maker knows how the information structure changes in the future, i.e. the decision maker anticipates that learning will happen. It is clear that, under certain conditions, anticipation can change the decisions in "preparation" before the learning event, increasing welfare. This increase, brought on by knowing when learning will happen, is called the value of anticipation.

In a previous publication on MIND-L (Lorenz *et al.*, 2012b) anticipation was investigated with respect to threshold damages and a window of only a few years is found where learning comes with a significant value of anticipation. As CRA has a different structure and no threshold damages but rather a penalty to utility, a discussion of the effect of assuming anticipation in this setting is necessary.

The value of anticipation is calculated by constraining the model to the optimal solution from the NOLEARN scenario up to the learning point. After the learning point, the decision maker is free to maximize welfare. The CBGE change between the constrained simulation and the standard LEARN scenario represents the value of anticipation. This value is always positive, because, by issuing the constraint, the solution is forced to be sub-optimal and therefore the non-anticipated welfare will necessarily be smaller than the welfare from the standard approach with anticipation.

In the analysis presented in this thesis, the effect of anticipation is an order of magnitude smaller than the EVPI as shown in Figure 3.11 by the green lines. For a target of 66% and 2°C the value of anticipation gradually drops from 0.02% to zero. This is very small compared to 0.66% EVPI. For a relaxed target of 50% and 2°C the value of anticipation is below 0.01% (Figure 3.11b).

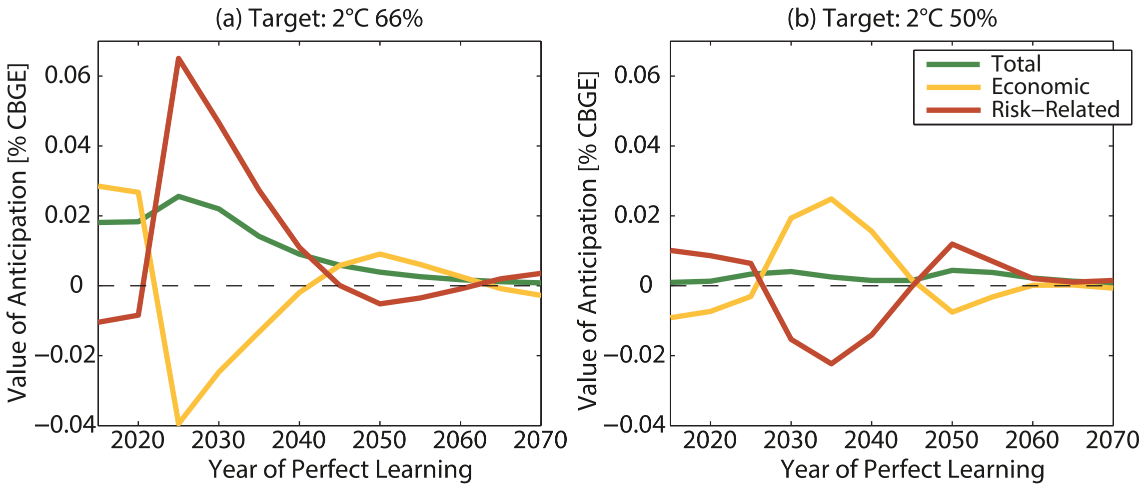

Figure 3.11. – The value of anticipation for two different targets (66% and 50% safety at 2°C) split into economic and risk-related value and plotted for different learning points.

Again there is a strong compensating behavior of economic and risk-related anticipation value which oscillates. The risk-related value in (a) has a strong peak for learning around 2025 which is is inverted in (b). Plotting the value of anticipation across the time and type dimension as shown in Figure 3.12 allows deeper insight.

The upper peaks in Figure 3.12a are at the time points which represent the last time step in which full information is not available (marked by squares). In the case of anticipation, more is consumed just before learning to balance heavy investments as soon as learning happens. Without anticipation, the decision maker is oblivious to the learning event and is "surprised" by it and so cannot cut back in consumption as much. It is interesting to see that after the learning point anticipation reduces welfare, this is balanced by increase in consumption before the learning point.

In Figure 3.12b the risk-related value does not follow a clear pattern. For early learning (dark colors) both a positive and a negative value is possible. The anticipation effect can be summarized as follows:

- For very early learning, anticipation leads to an overall increase in temperature (deduced from the negative risk-related value), hence allowing for an increase in consumption.
- Anticipating learning in 2025 up to 2040 leads to a loss of consumption to allow for a considerable reduction in temperature and risk.

Only an anticipation of learning between 2025 and 2040 leads to a reduction in risk. This seems to be similar to the "anticipation window" that was found by Lorenz *et al.* (2012b).

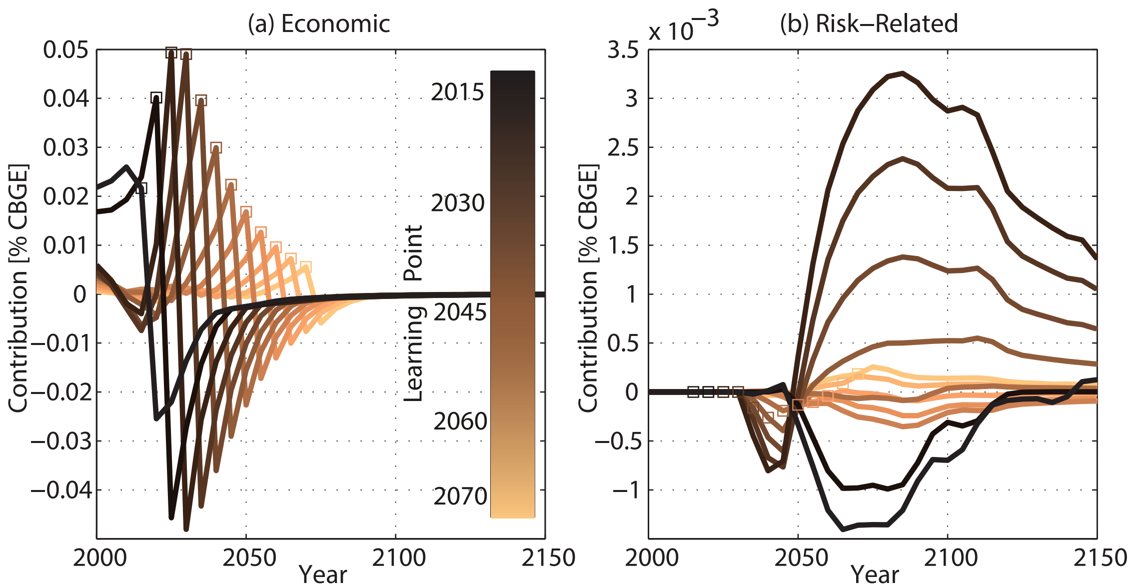

Figure 3.12. – Expected value of anticipation plotted for different learning points (indicated by the color) and divided across time and type. Values for a target of 2°C and 66% safety. Squares indicate the last time point without information.

3.4. Conclusion

This chapter is a collection of methods for analysis of CRA and applications of the same. We investigate the source of EVPI and how it changes if the target or risk metric is changed. We also look at the effect of the implicit assumption made in this thesis that the learning event is anticipated.

Due to the additive structure of CRA it is possible to approximately divide a specific CBGE change, representing the difference between any two simulations, into its constituting components. For the first time, it is possible to disentangle CBGE results in three dimensions: type (economic and risk-related), time and SOW. The division into economic and risk-related value is especially useful because the two parts have very different mechanisms and consequences. It allows us to determine if a value is mainly created by increasing economic activity and less mitigation or by reducing risk and increasing mitigation.

We find that most of the EVPI (around 2/3) is of economic origin especially if learning happens after 2020. If the target is relaxed this no longer is true. EVPI is mostly economic for strict targets because the optimal solution is close to the feasibility limit, which implies that any information received will allow the mitigation effort to be reduced in SOWs with low climate sensitivity. In SOWs with high climate sensitivity more mitigation is not possible and so the risk stays roughly the same.

We also find that adjusting guard rails is not equivalent to adjusting probabilities if learning is considered. Relaxing the guard rail generally leads to higher temperatures after learning than relaxing the safety target. Increasing the stringency of the target increases the EVPI for early learning, no matter if the change in the target is produced by a change in safety or a change of the guard rail. For learning in 2040 an increase in stringency, through a decrease of the guard rail, actually decreases the EVPI.

The investigation of the choice of risk metric shows that it plays only a minor role for the EVPI because this degree of freedom is mostly swallowed by the recalibration. This is a reassuring discovery and somewhat strengthens the reasoning behind the use of CRA, as it is robust w.r.t. such choices. The recalibration ensures that on average the marginal utility from reducing temperatures is equal no matter which risk metric is used. The EVPI only varies by a few percent because there is a compensating effect of an increase in economic value and a decrease in risk-related value.

Anticipation of the learning event leads to a slightly higher (1-5% change) EVPI. This is so small that the differentiation between anticipation and no anticipation is negligible. The interesting discovery was made that the decision maker, in anticipation of the learning event, consumes more than if oblivious of the learning event. This effect is similar to the effect of doing less mitigation in anticipation of learning which Lange & Treich (2008) have discussed. They argue that the change in emissions is not general and may or may not be positive, assuming reversible climate impacts.

4. Discounting and Risk Aversion

4.1. Introduction

Choosing the discount rate and level of risk aversion for a simulation can have a large impact on the outcome and is frequently the subject of intense debate. In MIND-L, in the previous chapters, we use a Constant Relative Risk Aversion (CRRA) of 2 and a Pure Rate of Time Preference (PRTP) of 1%/yr. By the Ramsey equation and an assumed growth rate of around 2%/yr this results in a social discount rate (SDR), or interest rate, of around 5%/yr.

A CRRA greater than zero implies decreasing marginal utility from consumption. This in general means, that if two distributions of consumption have the same mean, the one with the smaller variance is preferred. In the NOLEARN case the consumption path is deterministic and therefore, risk aversion is ineffective regarding uncertainty in this scenario.

A second effect of CRRA is an implied discounting. If the economy grows, a consumption change in the future is of less value than a consumption change earlier. Refer to Dasgupta (2008) and Lontzek & Narita (2011) for more details on risk aversion. For an interesting discussion on the disentangling of the two effects see Traeger (2014). We acknowledge the need for a discussion about such disentanglement but do not consider these issues here.

The PRTP captures the effect that receiving something earlier is preferred. The discount rate is assumed to be constant over time (exponential discounting) in most models to ensure time consistent decisions. A different method is hyperbolic discounting (decreasing discount rates, discussed in Groom *et al.* (2005)) which is the attempt to combine moderate discount rates in the immediate future with low discount rates for the far future, thereby, still attributing some relevance to the long term effects. However, such discounting methods give rise to the problem of time inconsistency. Heal & Millner (2013) discuss the aggregation of preferences of agents with different preferred discounting into an overall single agent with one discount rate. They show that this produces hyperbolic discounting. This thesis only considers exponential discounting to reduce complexity and to ensure time consistent decisions.

This chapter consist of a thorough investigation of the choice of SDR and how it influences the result of the simulation. Throughout the chapter, we refer to the two parameters CRRA and PRTP as “normative parameters” because we acknowledge that the choice of such parameters has a normative component. Discussions between Stern (2007), Nordhaus (2008) and Weitzman (2007a) highlight the different approaches to the choice of SDR. While Stern (2007) bases his assessment mostly on low values and even chooses a PRTP of 0.1%/yr due to the responsibility towards future generations, Nordhaus and Weitzman argue with descriptive arguments, that such low values do not reflect the behavior of society.

The effect of a change in SDR on the outcome of a Cost-Benefit Analysis (CBA) can be drastic, as pointed out by Weitzman (2007a). He further makes the point that the strong mitigation action that Stern (2007) suggests, is a direct result from his choice of low SDR. We confirm the result that SDR has a great impact on the monetary values in a simulation but, in contrast to CBA, has little effect on the optimal policy in Cost-Risk Analysis (CRA). This is not surprising as the optimal policy is primarily defined by the climate target and not by the discount rate.

Therefore, it can be argued that the choice of trade-off parameter in CRA is also normative (or the calibration point that is used). This would not be the case in CBA with a damage function which is the result of careful impact assessment. This chapter shows that a great strength of CRA lies in the reduction of the importance of SDR by attributing the implied value judgment to the choice of a climate target.

Weitzman (2007a) considers the choice of SDR “the biggest uncertainty of all” when constructing an integrated assessment model of climate change. To incorporate such uncertainty in the model, Pizer (1999) treats the normative parameters like any other uncertain parameter and a probability distribution is assumed. We follow a different approach here and implement a sensitivity study to not mix uncertainties in physical processes with normative discussions. We further analyze the forgone welfare, if decisions are made under “wrong” assumptions of SDR.

The chapter is structured as follows. First the normative parameters are introduced and the process of recalibrating CRA is explained. Section 4.3 shows the impact of SDR on the expected costs of climate protection (ECCP), on the expected value of perfect information (EVPI) as well as on the decisions. The final part analyzes the loss due to a misjudgment of SDR. The results support the choice of low values and confirm that the choice is not as important in CRA.

4.2. Normative Parameters in MIND

In the model MIND-L a CRRA function is used to calculate utility. This utility is discounted by the PRTP which is an exponential discount rate. The utility loss ("risk") from climate change is discounted by the same rate. For a deterministic discrete consumption path, and neglecting population growth, the welfare would be calculated according to:

$$W = \frac{1}{1-\eta} \sum_t C(t)^{1-\eta} e^{-t\delta} - \beta R(t) e^{-t\delta} \tag{4.1}$$

The parameters η (CRRA) and δ (PRTP) depend on interpretation and ethical considerations and may vary across people, societies and situations. In this analysis we offer a sensitivity study on these parameters. The interested reader is referred to Dasgupta (2008) and Arrow *et al.* (2013) and for further discussion of such parameters. In other publications that use CRRA and PRTP parameters (Cline, 1992; Pizer, 1999; Weitzman, 2007b; Stern, 2007; Nordhaus, 2008; Garnaut, 2008) the CRRA range from 1.5 to 2 and the PRTP from 0.05% to 3%/yr.

4.2.1. Social Discount Rate

Both normative parameters appear in the exponent of the economic part of the welfare equation 4.1 and so, for exponential growth, they both have a discounting effect. Assuming a growth rate α we get:

$$\frac{C_0^{1-\eta}}{1-\eta} \sum_t e^{\alpha t} e^{-t(\alpha\eta+\delta)} \tag{4.2}$$

The effective discount rate is SDR $= \alpha\eta + \delta$ which is called the social discount rate, or interest rate. This formula is also known as the Ramsey equation. The SDR is the rate with which future consumption changes have to be discounted to calculate the net present value. Intuitively, the reason that CRRA influences the effective discount rate is the fact that if there are decreasing marginal returns on consumption, the future is valued less because it is very rich. The growth rate α will depend on the normative parameters in some small way as well, but these second order effects are neglected here.

In the simulations carried out in this thesis, we test 24 combinations of CRRA and PRTP, shown in Table 4.1 together with the SDR. The growth rate is calculated by a least square fit of an exponential function to the consumption path and shown later in Figure 4.1c. The range of CRRA (0.75-5) and PRTP (0.5%-4%) is broader than the choices in the literature, although we lack extremely low values. Stern (2007) arrives at an implied SDR as low as 1.4%. The lowest we consider is 2.5% and the reference case has an SDR of around 5.4%. This is a common value, implying a drop to 1% of value after 90 years.

		PRTP (%/yr)			
		0.5	**1**	**2**	**4**
CRRA	**0.75**	2.5	2.8	3.7	5.6
	1.5	3.9	4.4	5.3	7.2
	2	4.9	**5.4**	6.3	8.2
	3	6.9	7.4	8.3	10.2
	4	8.8	9.3	10.2	12.1
	5	10.7	11.1	12.2	14.1

Table 4.1. – The 24 combinations of CRRA and PRTP and the equivalent SDR in %/yr. The value for the reference case from the previous chapter is in bold face.

Before we investigate the effects in MIND-L we briefly formulate the expected effects as the SDR increases:

1. The future becomes less important, therefore the savings rate is reduced which in turn reduces the growth rate.
2. The impacts of climate change mostly occur far in the future, so we expect the risk-related ECCP to be smaller.
3. As the EVPI results from economic adjustments in the future, we would expect an increasing SDR to decrease the EVPI.

An adjustment of the SDR causes a different strategy to be optimal. However, calibration dictates that in the NOLEARN case a specific target has to be exactly fulfilled. Therefore, a recalibration is necessary to ensure this goal. The next section discusses the effect on the calibration parameter.

4.2.2. Estimating the Calibration Parameter

In CRA, the trade-off parameter is calibrated so that the optimal solution fulfills a predefined target. If the normative parameters are changed, the balance is shifted due to changing marginal costs and a recalibration of the trade-off parameter is needed so that the target is met. In what way the trade-off parameter is affected by a change in normative parameters is explored in this section. We show that the recalibration can be internalized by formulating the trade-off parameter in terms of the normative parameters.

The range for the trade-off parameter β spans many order of magnitudes and depends strongly on CRRA (η) and PRTP (δ) as well as the cost structure of MIND-L. An empirical analysis of β from equation 2.15 and its changes with η and δ led to the conclusion that for MIND-L there is an independent and exponential dependency of β on η and δ . The parameter β increases exponentially in δ and decreases exponentially in η. This can be supported by intuitive arguments:

- When δ increases it increases the devaluation of the future which leads to the necessity of increasing the penalty for future temperature rises to reach the same goal as climate change mainly happens in the future.
- When η increases, the marginal benefit from consumption declines. To adjust the marginal benefit from risk reduction such that the same decisions are again optimal, β also has to fall.

The empirical dependence that was found is of exponential form:

$$\beta = a\, e^{-b\eta + c\delta} \tag{4.3}$$

where a, b and c are greater than zero and for the reference CRA (2°C, 66%) take on the values $a = 0.0209$, $b = 3.894$ and $c = 89.98$. To find the reason for this particular form, we simplify the problem starting from the original formulation:

$$W = \max_X \sum_{t=0}^{t_{\text{end}}} \sum_{s=1}^{S} p_s \left\{ \underbrace{U(X,t)}_{\text{economic}} - \underbrace{\beta R(T(X,t,s))}_{\text{risk-related}} \right\} e^{-\delta t}. \tag{4.4}$$

As we only look at the NOLEARN case, the economic part is not probabilistic. We also make use of the fact that risk happens later than economic welfare and attribute a representative time step t_C and t_R accordingly. In this way the problem is converted to a static problem and can be used to find supporting arguments for the exponential dependency:

$$W = \max_X \left\{ U(X)e^{-\delta t_C} - \beta \sum_{s=1}^{S} p_s R(T(X,s)) e^{-\delta t_R} \right\}. \tag{4.5}$$

Using the cumulative emissions E as a condensed control variable and absorbing the probabilistic parts into $R(E)$ leaves:

$$\max_E \left\{ U(E)e^{-\delta t_C} - \beta R(E) e^{-\delta t_R} \right\}, \tag{4.6}$$

where U is the utility function that maps cumulative emissions E to an associated economic utility incorporating the economic costs of mitigation, and R is the risk as a convex function in cumulative emissions. The time steps t_C and t_R are the representative time steps for the economy and risk. The economic utility is influenced mostly by decisions during the first decades due to the early investments needed to transform the energy sector to fulfill the strict climate target. Risk comes into play later in time as temperatures rise slowly. Therefore the value of t_R is larger than t_C.

In the calibration process we find the value for β which ensures that the climate target is met. This goal is reached with the same emissions, E_{cal}, independent of the normative parameters as it depends on physical processes. Put differently, β is tuned in order to make E_{cal} the optimal solution. In the optimal point, the derivative of the functional is zero. Assuming a CRRA utility function $U(E)$ we have:

$$\beta = \frac{U'(E_{\text{cal}})}{R'(E_{\text{cal}})} e^{\delta(t_R - t_C)}.$$

This equation shows the exponential dependency of β on δ. We chose t_R and t_C to be the representative time points and referring back to the empirical finding $c = 89.98$ gives an estimate for the time lag between costs and risk of about 90 years.

We further show that the empirically found dependency on η can also be approximately found in the static case describe above if we assume a utility function $U(E) = g(E)^{1-\eta}/(1-\eta)$ with an economic equation $g(E)$:

$$\beta = g(E_{\text{cal}})^{-\eta} \frac{g'(E_{\text{cal}})}{R'(E_{\text{cal}})} e^{\delta(t_R - t_C)} \tag{4.7}$$

$$\beta = \frac{g'(E_{\text{cal}})}{R'(E_{\text{cal}})} \exp\left[-\ln(g(E_{\text{cal}}))\eta + \delta(t_R - t_C)\right]. \tag{4.8}$$

A comparison with the empirical equation for β yields: $a = g'(E_{\text{cal}})/R'(E_{\text{cal}})$, $b = \ln(g(E_{\text{cal}}))$ and $c = (t_R - t_C)$. This explains the functional form of the calibration equation. If the exponential form of β is included in the dynamic welfare equation of CRA we have:

$$W = \max_X \sum_{t=0}^{t_{\text{end}}} \sum_{s=1}^{S} p_s \left\{ \underbrace{U(X,t)}_{\text{economic}} - a\, e^{-b\eta + c\delta} \underbrace{R(T(X,t,s))}_{\text{risk-related}} \right\} e^{-\delta t} \tag{4.9}$$

and reformulated:

$$W = \max_X \sum_{t=0}^{t_{\text{end}}} \sum_{s=1}^{S} p_s U(X,t) e^{-\delta t} - \sum_{t=0}^{t_{\text{end}}} \sum_{s=1}^{S} p_s R(T(X,t,s))\, a\, e^{-b\eta} e^{-\delta(t-c)}. \tag{4.10}$$

This implies that β effectively compensates for the large time shift in economic and temperature effects by bringing the risk-related impacts into the same time frame as the economic impacts by shifting them 90 years. Any change in δ now has equal net effects on the marginal benefits of risk and economy. The risk aversion parameter η changes the marginal benefits in the economy. The change is balanced by a similar shift in the risk-related part caused by the calibration of β. The coefficients a and b are tied to consumption levels and model parameters of MIND-L.

The next part of the chapter answers the question what the effect of the change in normative parameters together with the subsequent recalibration is.

4.3. Varying SDR

In this section we vary both CRRA and PRTP in the manner presented in Table 4.1. Sometimes the results give a smooth result with respect to SDR, which hints at the fact that the CRRA's role as risk aversion is negligible. Other times, the results form a point cloud, i.e. SDR does not explain all effects induced by CRRA and PRTP. To clarify the situation, lines of equal PRTP are included in most plots. When interpreting the results it is useful to keep in mind that the reference case analyzed in previous chapters has a SDR of 5.4%.

4.3.1. Cost of Mitigation

The EVPI is greatly dependent on ECCP (cost of mitigation plus welfare loss from risk). The smaller ECCP is, the less room for improvement there is, and EVPI is generally smaller. To assess the effect of SDR, we first look at the ECCP. Figure 4.1a shows the total ECCP and its economic and risk-related components. A visible effect is that decreasing the discount rate from 5.4%/yr decreases the ECCP strongly, economic and risk-related equally. Increasing the SDR also decreases the ECCP, mostly through the economic part.

The economic value (red curve in Figure 4.1a) follows a functional relation to SDR that resembles a log-normal distribution. The fact that all points lie on one line indicates that CRRA and PRTP are interchangeable. We use this fact to find the reason for such a behavior of the economic ECCP.

We make the assumption that the bulk of the economic costs are realized between t_1 and t_2 (for example 2015 and 2065, estimated from Figure 3.4). We further assume that the costs do not change w.r.t. variations in SDR as roughly the same mitigation has to be done to reach the climate target. This time frame is weighted differently as the discount rate changes. We model this by aggregating two constant consumption streams which differ between t_1 and t_2 by a constant a . The constant a represents the added cost of mitigation. A discounting of δ is applied which represents the total SDR in this case. To calculate the CBGE between the consumption streams a linear utility function is assumed for simplicity:

$$\mathrm{CBGE} = \frac{\int_0^\infty C_0 e^{-\delta t}\mathrm{d}t + \int_{t_1}^{t_2} a e^{-\delta t}\mathrm{d}t}{\int_0^\infty C_0 e^{-\delta t}\mathrm{d}t} - 1. \tag{4.11}$$

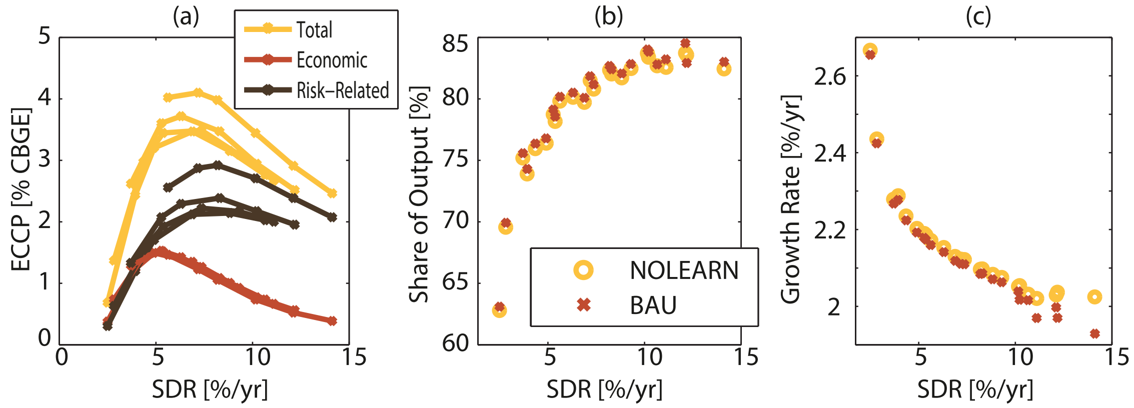

Figure 4.1. – (a) ECCP divided into risk-related (brown) and economic (red) parts for different social discount rates (SDR). The lines indicate equal PRTP. There are four lines, one for each PRTP that was used in the simulation (0.5%, 1%, 2%, 4%). The further the line is located to the right, the higher the PRTP. Panels (b) and (c) show indications of endogenous growth which is closely related to the consumption share of output. Values are plotted for NOLEARN and BAU cases. LEARN cases would lie indistinguishably between these.

This simplifies to

$$\mathrm{CBGE} = \frac{a}{C_0}\left(e^{-\delta t_1} - e^{-\delta t_2}\right), \tag{4.12}$$

with an extremum at

$$\delta_{\max} = \frac{\ln\left(t_2/t_1\right)}{t_2 - t_1}. \tag{4.13}$$

As seen in Figure 3.4, the majority of costs are realized between 2015 and 2065. Inserting these as years from the starting point of the model (1995) into the equation above yields a value of $\delta_{\max} = 2.8\%$. In the limits of δ, i.e. for very strong and no discounting, the CBGE goes to zero. This is the reason for the log-normal type form of the economic ECCP. It is a resonance between the time scale of the costs, i.e. transforming the energy system, and the discounting. We see this behavior again for the EVPI for the same reasons.

The anticipated effect of readjusting decisions, due to a change in preferences, is played down as a result of the recalibration. There is, however a change in the economy which can be seen in Figure 4.1b and c, which shows the consumption share of output as well as the growth rate of consumption for the NOLEARN and BAU case. The LEARN cases are not plotted as they are largely indistinguishable from NOLEARN. As the SDR increases, the consumption share increases and the growth rate of the economy decreases in-line with general economic wisdom and equals what we anticipated earlier.

4.3.2. Value of Information

Figure 4.2a shows the EVPI for different learning points and different SDRs. As expected, the value is smaller for all choices of SDR if learning arrives later. The effect of SDR can be described by considering SDR as valuing the future: an increase in the SDR devalues the future and therefore also decreases the EVPI stronger, the later information arrives. For low discount rates (which are more relevant for policy makers) the EVPI is smaller for immediate learning but decreases much slower and therefore results in higher values for learning in 2040.

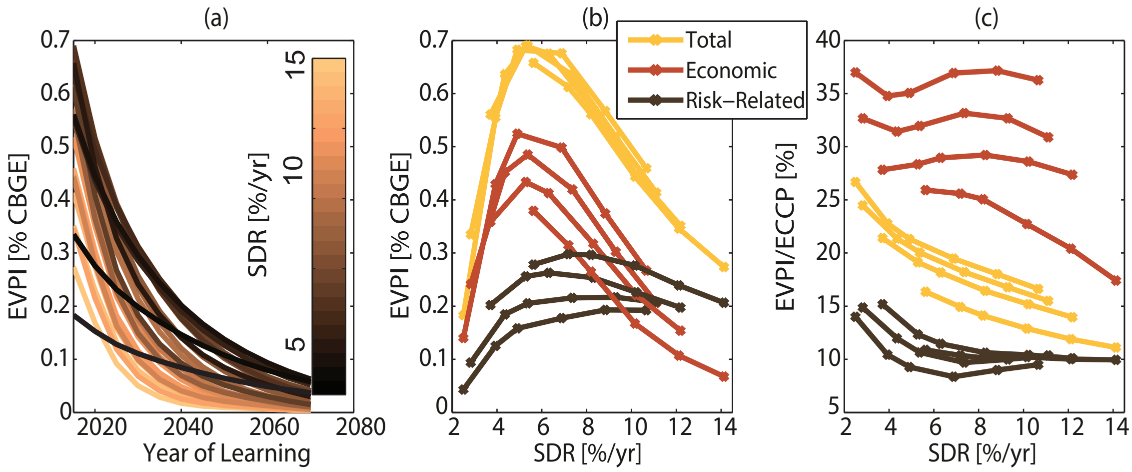

Figure 4.2. – EVPI depending on SDR. (a) EVPI(L,2075) where L is varied from 2015 to 2070. The color gradient indicates the SDR. For lower SDR, the EVPI drops more slowly as the year of learning is moved into the future. (b) and (c) shows the EVPI(2015,2075) and ECCP(2075) respectively split into economic (red) and risk-related (brown) parts. Lines represent constant PRTP (0.5%, 1%, 2%, 4%). If a line ends (or starts) further to the right it has a higher PRTP. From (c) it is evident that the relative value of information mostly depends on PRTP and not on CRRA.

In Figure 4.2b a similar concave behavior can be seen for total EVPI as for economic ECCP in the previous section. The lines in Figure 4.2b and c represent lines of constant PRTP. If a line ends (or starts) further to the right it has a higher PRTP. If the lines lie on a shared path, PRTP and CRRA have the same effect on the analyzed value. For the total EVPI the reason the points lie clustered on a line is a compensating effect of economic and risk-related value for a change in PRTP. As PRTP increases, the risk-related EVPI increases (as brown lines move to the right, points move upward) but the economic EVPI decreases (red lines further to the right are lower).

Figure 4.2c shows the relative reduction of the ECCP by learning in 2015. Changing CRRA does not seem to affect the economic reduction of ECCP much (point to point on one line) as the lines are relatively flat, however, a change in PRTP (line to line) can affect the relative value greatly. This implies that PRTP defines the relative economic benefit from perfect information whereas CRRA plays a minor role. Overall, the relative reduction in cost of mitigation (economic ECCP) is about three times larger (25%-37%) than the relative reduction of risk (10%-15%).

4.3.3. Effect on Decisions

What remains to be shown is the absolute effect of SDR on temperature and on decisions. For sake of simplicity only the NOLEARN case is considered to find the effect of normative parameters on the mitigation strategy. Due to the nature of CRA, and the recalibration that is done for each combination of normative parameters, the effect on the mitigation strategy is small. This is because the same target has to be reached in all cases.

In Figure 4.3a, the annual emissions are plotted. All are rapidly decreased in order to reach the target. A bundle of trajectories decrease to zero emissions about a decade later. These are the simulations with a PRTP of 4%. One would assume that due to these extra emissions the resulting temperatures have to be higher. This is only the case on average, the maximum temperature is actually the same. The reason for this is that an investment into fossil fuels causes short term emissions of aerosols creating a short term reduction in climate forcing. This shifts the peaking of the temperature a few years and allows for a slight increase in emissions. After the peak, the temperatures always stays about 0.05°C higher which incurs a greater nominal risk, however, due to the larger discounting, this is accepted.

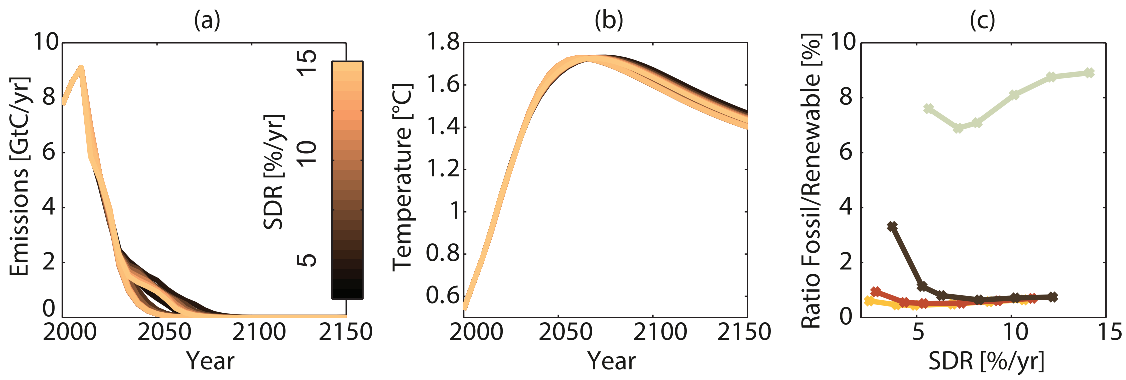

Figure 4.3. – (a) Emissions, (b) expected temperature and (c) fossil to renewable energy ratio in 2050 for the NOLEARN scenario and 24 combinations of CRRA and PRTP plotted as SDR (shades of brown). (c) shows lines of equal PRTP (yellow 0.5%, red 1%, black 2% and cyan 4%). (a) shows a difference in emissions in the year 2050 which is again reflected in (c) (cyan line). For 4% PRTP more is emitted but due to cooling effect of aerosols, temperatures still reach the target at the expense of higher temperatures later. (b) is evidence for little effect of SDR on the temperature.

Figure 4.3b shows expected temperature and it is easy to see that, no matter what the choice of normative parameters is, the solution has the same maximum. Overall the effect on the actual temperature is very small. This highlights a key benefit of CRA: the decisions are only influenced marginally by the choice of SDR.

The increased fossil fuel activity of the 4% PRTP case can be seen by plotting the ratio of fossil to renewable energy for the year 2050, as shown in Figure 4.3c. The gray curve shows around 8% whereas all other curves lie mostly below 2%. This case is not of much relevance because 4% PRTP is widely viewed as unsupportable but it shows nicely that different preferences can lead to different solutions even though the same climate target is reached in an optimal fashion.

4.4. Misjudged Discounting

4.4.1. Correct Normative Parameters

For the following analysis, we make an exploratory assumption that there is a choice of normative parameters that best reflects the behavior or ethical preferences of the global community. The parameters that are used for the policy advising scenario might deviate from this best choice. This section finds the impact of such a "wrong" choice in terms of forgone welfare.

Before going into the analysis a few points to be considered are presented. Firstly, misjudging the SDR will over or under estimate the impact of the climate in relation to the economic decisions in general. CRA relies on a recalibration to a target, which is itself a normative decision. This recalibration balances the risk-related values to any changes in normative parameters as shown in Section 4.2.2. This effect greatly reduces the impact of a wrong selection of parameters on the climate.

Secondly, the economy will develop according to the savings rate, which is derived from the interest rate, which in turn is represented by the SDR. The Ramsey type growth model used follows the paradigm that the economy will develop in an optimal sense. Lets assume that the interest rate is internal to the economy and determined by the behavior of agents, which will then drive the economy along an optimal path. This implies that it is not possible to force the economy along a non-optimal path, provided the market is free to adjust. This implies that a climate policy might lead to suboptimal decisions in the energy sector but the rest of the economy is free to adjust to the situation in an optimal way.

The counter assumption would be that the policy maker, steers the economy (or places incentives accordingly) to follow a calculated optimal path even though the SDR that was used to calculate this optimal path might not reflect the preferences of society. We refer to this situation as a fixed economy. Both assumptions allow for the calculation of the loss of welfare, if optimization is done under the assumption of a "wrong" SDR. The two scenarios are:

1. Fixed economy: all investment decisions, consumption paths, and risk is found using a particular set of parameters p_0. The loss is calculated by evaluating these decisions under parameter set p_1.
2. Free economy: only investments in the energy sector are fixed, i.e. the climate policy is fixed to the strategy found in a simulation with parameters p_0. The loss is calculated if these decisions were evaluated under parameter set p_1 while other decisions, such as consumption and investment into the common good, can adjust to maximize utility.

4.4.2. Loss due to Misjudgment

To find the welfare loss described above we compare optimal solutions with suboptimal solutions. We divide the control variables into two camps: the ones that control energy policy x and others y. In the "free economy" case, the controls y are chosen by the market mechanism to be optimal, whereas x is chosen by the energy policy defined by the policy maker. Furthermore, we consider two possible control paths x_0 and x_1 representing the optimal decisions for a set of normative parameters p_0 and p_1 respectively. We denote the welfare as $W(x_i, y_j; p_j)$ as it depends on controls and normative parameters. If $i = j$ the solution is optimal, however, if a policy x_0 is chosen on the basis of p_0 but the policy is evaluated in a world where p_1 is true then the welfare is necessarily suboptimal.

With a fixed economy, we compare the welfare in a world where the controls y can not be adjusted by the market mechanism and are set by the optimization under the assumption of normative parameter set p_0. If it turns out that the true parameters are p_1 the CBGE loss can be calculated by:

$$\Delta_{\text{fix}} = \left[\frac{W(x_0, y_0;\, p_1)}{W(x_1, y_1;\, p_1)}\right]^{\frac{1}{1-\eta}} - 1. \tag{4.14}$$

On the other hand, if the economy can adjust freely to the true parameters we have a smaller loss calculated by:

$$\Delta_{\text{free}} = \left[\frac{W(x_0, y_1;\, p_1)}{W(x_1, y_1;\, p_1)}\right]^{\frac{1}{1-\eta}} - 1. \tag{4.15}$$

Note that the welfare values, that are compared, always have to be based on the same normative parameters, otherwise the calculation of CBGE is not valid. Figure 4.4a shows a contour plot for the "fixed economy" case. The SDR is calculated by the Ramsey equation from the normative parameters CRRA and PRTP as discussed in Section 4.2. The hypothesis (the parameters used for policy finding) represents p_0 whereas the reality is p_1. The loss of welfare can reach up to 40% CBGE which is an astonishingly large value. It can be interpreted as the value that would be lost if society were forced along a path devised by a dictator with completely different discounting preferences than the aggregate society or an ill-posed incentive scheme. This large loss only appears if the true SDR differs by more than a factor two from the hypothesis.

Figure 4.4b shows the case for the "free economy". Here the economy can adjust to the "true" normative parameters which reduces the welfare loss greatly. The leftover welfare loss originates only from suboptimal investments into the energy sector. This loss can be interpreted as the loss of welfare if an energy policy is fixed that is based on wrong assumptions about the SDR of society. It seems that it is worse to assume a high discount rate if it is low in reality (bottom, right), than vice versa. In general, it is soothing to know that an error by about 2% points SDR only creates a loss of below 0.2% as long as the market is free to adjust.

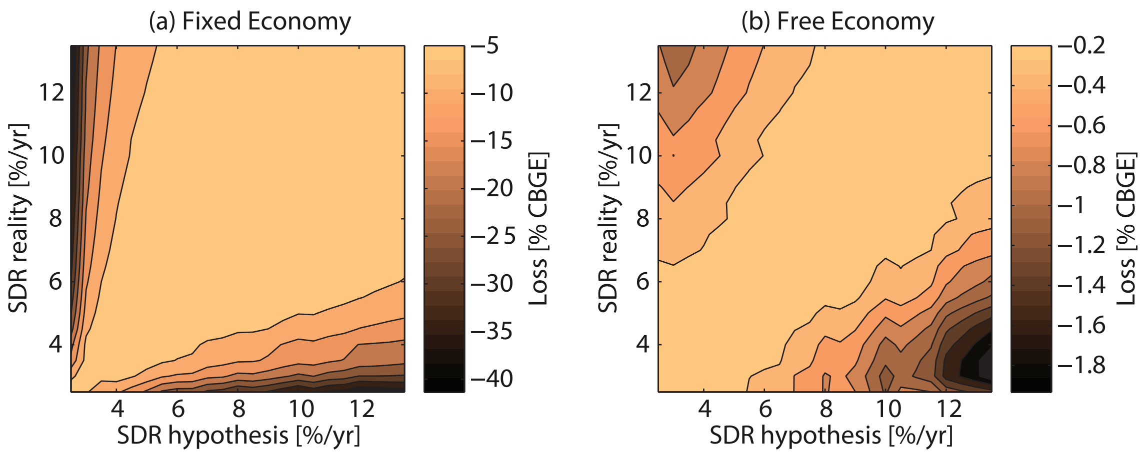

Figure 4.4. – Loss of welfare in CBGE percentage if policies are optimized under different normative discount rates, than the model is subject to. (a) shows the result including the whole economy in the fixed decisions and (b) shows the situation if only the energy sector is decided bases on the hypothesized parameters.

This small loss is partly due to the recalibration which is at the heart of CRA and constitutes one of its strengths. If it is assumed, that the economy will grow along an optimal path, and it is up to the policy maker to decide the energy policy (investments), using a CRA with a target greatly reduces the impact of the choice of normative parameters. The underlying reason being that CRA acts similar to a CEA by reducing the climate policy question to "How to fulfill the target at the lowest cost possible?". A real trade-off between costs and benefits, as is done in CBA, is only a second order effect. The normative decisions are exported to the choice of climate target.

It is unclear what this result means for the policy maker and we offer some ideas here. Policy makers should rather use lower discount rates to find optimal climate policy, as the error made, if the choice of SDR turns out to be incorrect, is smaller than vice versa. This argument is also supported by Weitzman (2007a). An argument to further support a choice of lower than normal SDR is that society has a lower discount rate for climate problems than for the actions in the economy. Similarly, people tend to have much lower discount rates when considering the future of their child, than when making purely economic decisions.

It is generally easier for policy makers to discuss targets rather than normative parameters, therefore the recommendation would be to focus on setting a widely accepted consistently posed target, rather then figuring out which the correct normative parameters are.

4.5. Discussion

This chapter shows the, possibly extreme, effect a change in normative parameters can have on a wide variety of aspects. But what implications does this have and how can this information be used productively?

It is still an open question how to map society's preference into a model of the climate problem. To some extent, a climate target takes on the role of determining how important a stable climate is to the decision maker. In a CBA this is done through the damage function in conjunction with the discount rate. The role of the discount rate is greatly reduced in CRA due to the recalibration, hence, SDR ceases to be "the biggest uncertainty of all" (Weitzman, 2007a). Consequently, for finding climate policy, improved normative parameters are not as important when using CRA because the preference order is captured by the target used for calibration.

However, if EVPI or cost of mitigation is sought to be the major output of the analysis, then the normative parameters have a large effect and more research is needed into which normative parameters are supported by society. Model output can vary an order of magnitude for these monetary results of the analysis. The ECCP varies from 0.5% CBGE to a maximum of 4% CBGE. The maximum is reached at an SDR between 5% and 6%/yr. The economic ECCP, i.e. the cost of mitigation, peaks with a value of 1.5% CBGE.

The EVPI for learning in 2015 varies in a similar fashion between 0.2% and 0.7% CBGE. Formulated as a ratio, the EVPI makes up between 15% and 25% of ECCP. This fraction increases, if just the economic parts are considered, to a range of 25% to 37%, i.e. roughly 1/3 of the cost of mitigation can be saved if information arrives in 2015.

The reason for the strong dependence on discounting is that the SDR is responsible for the relative importance of time steps. Suppose a stream of utility changes is strongest in a specific time interval, as for example the value of information is between the compared learning time points. Then the SDR will define the importance of this time interval and therefore change the value greatly, even though decisions have not changed at all. Discounting may even be so strong, that the time interval considered is discounted into irrelevance, or discounting is so small that the time interval looses importance next to all the subsequent time steps. This discussion clarifies that there is an SDR for which values such as EVPI and ECCP reach their maximum, i.e. when the time constant of the discount rate resonate with the time horizon of the effect that is analyzed.

In MIND-L the SDR which resonates with the cost of mitigation as well as the value of information lies around 5%/yr which is a good average of the popular choices. This is equivalent to a time constant of 20 years.

To find out if 5%/yr would be a good choice, we find the loss that we would endure if, in truth, society endorsed a different SDR than the one used to find the optimal climate policy. For discount rates in the span of 2.5%/yr to 12%/yr the loss of welfare would lie below 0.5% CBGE which is in the same order of magnitude as the EVPI.

We find that it is generally better to underestimate the SDR than to overestimate it. This is beneficial because ethical discussions tend to drive the discount rate to lower values and it is in line with Weitzman (2007a), who explains that in problems with such long periods it is best to use interest rates at the lower end of possible values. Therefore, we recommend using an SDR below 5%/yr.

Our analysis is also related to the work by Dietz & Matei (2013) who devise a theory of time-stochastic dominance to find sets of normative parameters in which one solution is dominant. Our approach relaxes this theory to a more moderate question of finding an optimal solution which performs well in a wide range of possible SDRs. Another related theory is the minimax regret theory by Savage (1951) which aims to find the choice which brings the least regret in the worst case. This can also be applied here but is heavily dependent on the range of SDR that is analyzed and therefore the usefulness is questionable in this situation.

5. Analytical CRA and further Approximations

5.1. Introduction

The overall goal of this chapter is to harvest benefits from using simplified models to understand fundamental relations and interactions. Throughout the analysis presented so far, simple relationships have been found between parameters and model results such as the exponential dependence of the trade-off parameter on discount rate and risk aversion or the dependence of value of information on the social discount rate. These are indicators that MIND-L's evaluation with a Cost-Risk Analysis (CRA) and a climate target could be represented by simpler equations and still yield important insights.

The aim of the first part of this chapter is to find analytical expressions for optimal emissions and to understand the mechanisms behind the optimization. Giving an analytical expression for value of information is difficult and beyond the scope of this thesis, therefore we restrict ourselves to calculating optimal emissions which still allows for some insight. The first sections introduce the simplifications that are needed to formulate the static model. Next, the no-information case is calculated and the calibration to a target is conducted. Lastly, we calculate the optimal emissions in the perfect information case, distinguishing four cases. An example plot is provided to show the situation if the climate target is very strict.

The second half of the chapter looks into the shape of the optimum of a standard CRA simulation with MIND-L, without learning, i.e. the same scenario presented in Chapter 2. The immediate vicinity of the optimal solution is probed by adding small variations to the main control variables. The main control variables are chosen to be the investments into renewable and fossil based energy until 2150, yielding 52 variables.

By studying the welfare loss that is inflicted if the solution is forced away from the optimum by randomly drawn offsets, a second order function can be found that describes the shape of the welfare equation at the optimal solution. The coefficients of the function allow for an assessment of which decisions, or variables, are most important.

5.2. Analytical Static CRA

In Chapter 2 CRA was introduced by using a static model. A lot of insight can be gained by constructing a static model, as it is much simpler to grasp and often is a good first approximation for a dynamic problem. In discussion about the trade-off parameter and recalibration in Chapter 4 it was demonstrated that the climate's response time scale also expresses itself in CRA. This leads to the key assumption that the climate problem can be conceptually split into two time steps. In the first time step, the economic decisions are made and in the second step, the climate impacts are realized depending on the economic decisions. This allows for the construction of a static model and can be backed up by considering the plot of expected costs of climate protection (ECCP) divided across time as shown in Figure 3.4 and in Held *et al.*, 2009.

Therefore, CRA is reformulated as a stylized tradeoff between a decision to mitigate now and the effect of climate later. In essence this is a two time step model but without mitigation decisions in the second step as well as no climate effect in the first step. The reason being that a target such as 2°C requires strong early mitigation and temperatures have not risen far enough to matter during the first time step. A further simplification is reached by only considering no information or perfect information situations. A more complex framework might include act-learn-act dynamics but to understand what CRA does this is not necessary.

This section formulates CRA and calculates the optimal emissions analytically to enable future analysis on this level. We show that there are four cases as already described in Section 3.2.3.

The dynamic problem for the case without learning is reprinted here for convenience:

$$W = \max_X \sum_{t=0}^{t_{\text{end}}} \sum_{s=1}^{S} p_s \left\{ \underbrace{U(X,t)}_{\text{economic}} - \underbrace{\beta R(T(X,t,s))}_{\text{risk-related}} \right\} e^{-\delta t} \tag{5.1}$$

We conduct the following approximations and simplifications to arrive at a static model:

1. The control variable is assumed to be cumulative emissions E. This choice has enjoyed wide popularity due to its simplicity and good capacity to capture future temperature rise.

2. Temperature is calculated by $T = E\theta\rho_1$ where θ is the climate sensitivity which is uncertain and ρ_1 is a conversion parameter. The climate sensitivity is distributed according to $f(\theta)$ which is assumed to be log-normal throughout this thesis. The parameter γ that was used in Chapter 2 is equivalent to $\theta\rho_1$ but the added distinction is needed here to be able to interpret θ as the climate sensitivity.
3. The risk from a temperature increase is calculated by the exceedance amplitude above the guard rail: $\Theta(T - T_{\mathrm{g}})(T - T_{\mathrm{g}})$ where Θ is the Heaviside function and T_{g} is the guard rail.
4. The economic utility is replaced by a deviation from the exponential growth path, i.e. as negative costs of mitigation. We choose $C(E) = \rho_2 E^{-n}$ where ρ_2 is again a conversion parameter. It is a reasonable choice because it bears close resemblance to the CRRA function and features infinite marginal costs at $E = 0$.
5. Discounting is done by introducing representative time steps, as has been done in Section 4.2.2.

These changes result in the following simplified form:

$$W = \max_E \int_0^\infty \left\{-\rho_2 E^{-n} e^{-\delta t_C} - \beta\Theta\left(E\theta\rho_1 - T_{\mathrm{g}}\right)\left(E\theta\rho_1 - T_{\mathrm{g}}\right)e^{-\delta t_R}\right\} f(\theta)\mathrm{d}\theta \quad (5.2)$$

If no learning is modeled, the cumulative emission E (in the following only referred to as "emissions" for simplicity) is a scalar, i.e. for all states of the world the same decisions are made. If learning is modeled, the emissions are a function of the value θ that is "learned". The first step in a CRA is to calibrate the no-information scenario to the desired climate target. Then the optimal emissions for the perfect information case can be determined.

5.2.1. The No Information Case

First we discuss the case of no learning and conduct the calibration, similar to Section 4.2.2. A calibration target is given by the temperature guard rail T_{g} and "safety" p_{g} (probability of staying below T_{g}). This target is just reached if E is equal to E_{g} which follows from the temperature equation: $T_{\mathrm{g}} = E_{\mathrm{g}}\theta_{\mathrm{g}}\rho_1$ where θ_{g} is the p_{g} quantile of the climate sensitivity, i.e.: $F(\theta_{\mathrm{g}}) = p_{\mathrm{g}}$. To simplify the expression we pull parameters together into $\hat{\beta}$ and omit any constants as they do not affect the maximization problem ($\hat{W} = W/\left(\rho_2 e^{-\delta t_C}\right)$). The welfare reads:

$$\hat{W} = -E_{\mathrm{g}}^{-n} - \hat{\beta}\int_{\theta_{\mathrm{g}}}^\infty \left(E_{\mathrm{g}}\theta\rho_1 - T_{\mathrm{g}}\right) f(\theta)\mathrm{d}\theta \quad (5.3)$$

with

$$\hat{\beta} = \beta \frac{e^{-\delta t_R}}{\rho_2 e^{-\delta t_C}} = \frac{\beta}{\rho_2} e^{-\delta(t_R - t_C)}. \tag{5.4}$$

Further, we define a general function $H(E)$, to replace the integral in equation 5.3 (for $E = E_{\mathrm{g}}$). We also define $\bar{\theta}$ as the expected value of θ. The welfare function is reduced to $\hat{W} = -E^{-n} - \hat{\beta} H(E)$ with:

$$H(E) := E\bar{\theta}\rho_1 - T_{\mathrm{g}} + \rho_1 E \int_0^{\frac{T_{\mathrm{g}}}{E\rho_1}} F(\theta) d\theta, \tag{5.5}$$

It is helpful to calculate the derivative of $H(E)$ w.r.t. E:

$$\hat{h}(E) := H'(E) = \bar{\theta}\rho_1 - \frac{T_{\mathrm{g}}}{E} F\left(\frac{T_{\mathrm{g}}}{\rho_1 E}\right) + \rho_1 \int_0^{\frac{T_{\mathrm{g}}}{\rho_1 E}} F(\theta) d\theta. \tag{5.6}$$

The expressions $T_{\mathrm{g}}/\rho_1 E$ is repeated often and represents the climate sensitivity θ^* that is necessary for the emissions E to reach a temperature T_{g}. By transforming the coordinates and carrying out the integral in equation 5.6 a simplified representation of $\hat{h}$ can be found[1]:

$$h(\theta^*) = \bar{\theta}\rho_1 \left[1 - F\left(\theta^* e^{-\sigma^2}\right)\right]. \tag{5.7}$$

For the further calculation we define:

$$\tilde{h}(\theta^*) := 1 - F\left(\theta^* e^{-\sigma^2}\right) \tag{5.8}$$

which is only dependent on the distribution of θ, shown in Figure 5.1. For a log-normal distribution $f(\theta) = \mathcal{LN}(0.973, 0.4748)$ and climate target of 2°C with a safety of 66%, the approximate value of $\tilde{h}(\theta_{\mathrm{g}}) \approx 0.5$.

With these helper functions we can reduce the maximization problem to:

$$\max_E \left\{-E^{-n} - \hat{\beta} H(E)\right\}. \tag{5.9}$$

[1]For this the following equation is necessary: $\int_0^x F(\theta)\mathrm{d}\theta = xF(x) - \bar{\theta}F\left(xe^{-\sigma^2}\right)$ assuming that $F(\theta)$ is the cumulative distribution function of a log-normal distribution with parameters μ and σ.

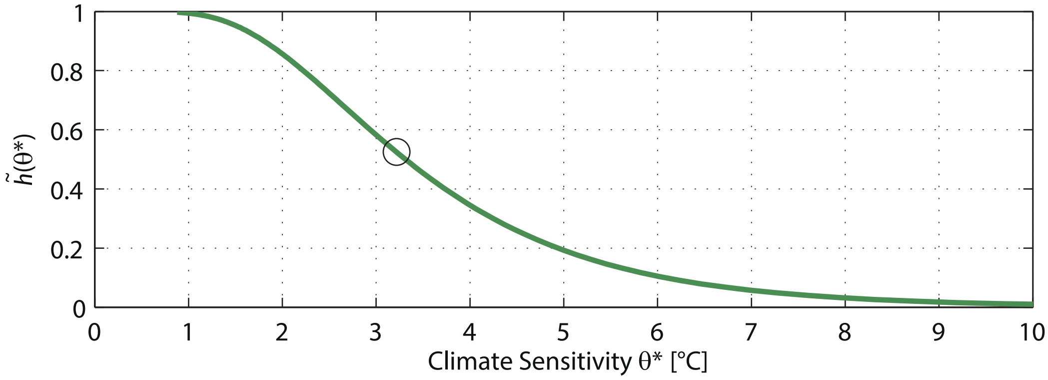

Figure 5.1. – A plot of $\tilde{h}(\theta^*)$ where θ^* is the climate sensitivity which makes temperatures rise to T_g if emissions are E, i.e. $\theta^* = T_\mathrm{g}/(\rho_1 E)$. The circle marks the value of $\tilde{h}(\theta_g)$ which is roughly around 0.5.

The optimality condition, that the first derivative with respect to E is equal to zero, leads to the following expression for the trade-off parameter, ensuring that the derivative is zero exactly at $E = E_\mathrm{g}$:

$$\hat{\beta} = \frac{n E_\mathrm{g}^{-n-1}}{h\left(\theta_\mathrm{g}\right)}. \tag{5.10}$$

Reinserted into the original welfare equation we have an expression for the optimal welfare:

$$\hat{W}_\mathrm{NL} = -E_\mathrm{g}^{-n} - \frac{n E_\mathrm{g}^{-n-1}}{h\left(\theta_\mathrm{g}\right)} H\left(E_\mathrm{g}\right). \tag{5.11}$$

This completes the calibration process and describes the case with uncertainty and no learning. The optimal emission strategy is E_g by design and equal for all states of the world θ. For the business as usual (BAU) case, we assume that the maximum feasible emissions are E_max which would lead to a (scaled) welfare $-E_\mathrm{max}^{-n-1}$ because there are no effects from the climate in the BAU case. The scaled welfare is multiplied by $\rho_2 e^{-\delta t_C}$ to be scaled back into the original welfare units. Assuming $n > 0$ and a large value for E_max results in BAU welfare to be close to zero. The cost of mitigation is represented by the change of welfare from the BAU welfare level to the economic part of W_NL given by $-\rho_2 e^{-\delta t_C} E_\mathrm{g}^{-n}$.

5.2.2. The Full Information Case

In the case of full information, the optimal emissions $E(\theta)$ depend on the value of θ that was learned. To reach the overall welfare the expected value across θ is calculated. The complete and calibrated welfare equation reads:

$$\hat{W} = \int_0^\infty \max_{E(\theta)} \left\{ -E(\theta)^{-n} - \frac{nE_{\mathrm{g}}^{-n-1}}{h\left(E_{\mathrm{g}}\right)} \Theta\left(E(\theta)\theta\rho_1 - T_{\mathrm{g}}\right)\left(E(\theta)\theta\rho_1 - T_{\mathrm{g}}\right) \right\} f(\theta)\mathrm{d}\theta. \tag{5.12}$$

For clarity, we first look at the maximization for a specific value of θ and later include the expected value operation. If upper and lower feasibility limits on cumulative emissions are assumed, the optimization has four possible regimes in θ that are analogous to those identified in Section 3.2.3. The regimes are distinct in the way their optimal emissions are calculated and can be separated by three limiting values of climate sensitivity: θ_{low}, θ_{sw} and θ_{high}.

1. If $\theta < \theta_{\mathrm{low}}$ even the BAU case emissions do not produce temperatures above the guard rail and therefore no mitigation has to be done. This leaves the optimal cumulative emissions to increase to:

$$E_1^* = E_{\max}. \tag{5.13}$$

2. If $\theta_{\mathrm{low}} < \theta < \theta_{\mathrm{sw}}$ the optimal solution is to stay below the guard rail, i.e. just reach it, and pay the costs for doing so, because the marginal risk from crossing the guard rail surpasses the marginal cost saving. The optimal emissions are given by:

$$E_2^*(\theta) = \frac{T_{\mathrm{g}}}{\rho_1\theta}. \tag{5.14}$$

3. If $\theta_{\mathrm{sw}} < \theta < \theta_{\mathrm{high}}$ the full trade-off between risk and cost is made. The first order condition derived from equation 5.12 reads:

$$nE(\theta)^{-n-1} = \frac{nE_{\mathrm{g}}^{-n-1}}{\bar{\theta}\tilde{h}\left(\theta_{\mathrm{g}}\right)}\theta. \tag{5.15}$$

Therefore, optimal emissions are given by:

$$E_3^*(\theta) = E_{\mathrm{g}} \left(\frac{\theta}{\bar{\theta}\tilde{h}\left(\theta_{\mathrm{g}}\right)} \right)^{-\frac{1}{n+1}}. \tag{5.16}$$

4. The last regime covers the case that $\theta > \theta_{\mathrm{high}}$ with resulting optimal emissions below the feasibility limit forcing the solution to maximum mitigation:

$$E_4^* = E_{\min}. \tag{5.17}$$

The lower limit θ_{low} is found by equating $E_{\max}$ and $E_2^*(\theta_{\text{low}})$ to find the value of climate sensitivity for which even maximum emissions would not breach the guard rail:

$$\theta_{\text{low}} = \frac{T_{\text{g}}}{\rho_1 E_{\max}}. \tag{5.18}$$

The value of climate sensitivity (θ_{sw}) for which the regime switch from 2 to 3 takes place when the marginal cost of staying below the guard rail is equal to the marginal risk of crossing it. This is the same as finding the point at which the optimal emissions for case 2 and 3 are identical, i.e. $E_2^*(\theta_{\text{sw}}) \stackrel{!}{=} E_3^*(\theta_{\text{sw}})$. For convenience and future reference the optimal emissions at this point are referred to by E_{sw}:

$$E_{\text{sw}} = E_{\text{g}} \left(\frac{\theta_{\text{sw}}}{\overline{\theta}\tilde{h}\left(\theta_{\text{g}}\right)} \right)^{\frac{1}{-n-1}} = \frac{T_{\text{g}}}{\rho_1 \theta_{\text{sw}}}, \tag{5.19}$$

$$\theta_{\text{sw}} = \theta_{\text{g}} \left(\frac{\theta_{\text{g}}}{\overline{\theta}\tilde{h}\left(\theta_{\text{g}}\right)} \right)^{\frac{1}{n}}. \tag{5.20}$$

The value of θ_{high} is more complicated as it depends on the previous values. It lies at the point where the otherwise optimal emissions begin to be less than $E_{\min}$. A differentiation has to be done by the relation between $E_{\min}$ and E_{sw}. This defines which of the cases (2 or 3) is active when $E_{\min}$ is reached. Once this is known, θ_{high} is calculated by equating the optimal emissions accordingly:

$$E_{\min} > E_{\text{sw}}: \quad E_2^*(\theta_{\text{high}}) \stackrel{!}{=} E_4^*(\theta_{\text{high}}), \tag{5.21}$$

$$E_{\min} < E_{\text{sw}}: \quad E_3^*(\theta_{\text{high}}) \stackrel{!}{=} E_4^*(\theta_{\text{high}}). \tag{5.22}$$

From this it follows:

$$\theta_{\text{high}} = \begin{cases} \frac{T_{\text{g}}}{\rho_1 E_{\min}} & E_{\min} > E_{\text{sw}} \\ \overline{\theta}\tilde{h}\left(\theta_{\text{g}}\right) \left(\frac{E_{\text{g}}}{E_{\min}}\right)^{n+1} & E_{\min} < E_{\text{sw}} \end{cases}. \tag{5.23}$$

Note that if $E_{\min} > E_{\text{sw}}$, case 3 is never realized, as the optimal emissions reach the bound before a trade-off with the risk would be optimal.

We have now calculated all four possible optimal emission expressions and the limiting values for θ. Hence, we can now formulate the expected welfare for the perfect information case. The completely expanded expression is given in Appendix A.4, we show here the simplified end result assuming that $\theta_{\text{high}} > \theta_{\text{sw}}$ and using the previously determined optimal emissions. We also divide the total welfare into economic welfare $\hat{W}_C$ and risk-related welfare $\hat{W}_R$ components:

$$\begin{aligned}\hat{W}_C = & - E_{\text{max}}^{-n} F(\theta_{\text{low}}) \\ & - \int_{\theta_{\text{low}}}^{\theta_{\text{sw}}} E_2^*(\theta)^{-n} f(\theta) d\theta \\ & - \int_{\theta_{\text{sw}}}^{\theta_{\text{high}}} E_3^*(\theta)^{-n} f(\theta) d\theta \\ & - E_{\text{min}}^{-n} \left(1 - F\left(\theta_{\text{high}}\right)\right),\end{aligned} \tag{5.24}$$

$$\begin{aligned}\hat{W}_R = & \hat{\beta}\Big(\int_{\theta_{\text{sw}}}^{\theta_{\text{high}}} \left(E_3^*(\theta)\theta\rho_1\right) f(\theta) \mathrm{d}\theta \\ & + E_{\text{min}}\rho_1 \int_{\theta_{\text{high}}}^{\infty} \theta f(\theta) \mathrm{d}\theta \\ & - T_{\text{g}} \left(1 - F\left(\theta_{\text{sw}}\right)\right) \Big).\end{aligned} \tag{5.25}$$

5.2.3. Special Case of Strict Targets

The tedious analytical derivation of welfare now bears fruit by giving insights into the situation if the target is strict, i.e. the target is such that it can only be met in the no-information case if emissions are close to E_{min}. We highlighted characteristics thus far in the thesis that lead us to the claim that the 2°C and 66% target is quite strict and produces optimal emissions that are close to the feasibility limit. We make use of this as a simplification by stating that the optimal emissions in the no-learn case, E_{g}, are very close to the minimum emissions, E_{min}. To study the implications we first calculate an expression for E_{sw} by inserting equation 5.20 into equation 5.19:

$$E_{\text{sw}} = E_{\text{g}} \left(\frac{\overline{\theta} \tilde{h}\left(\theta_{\text{g}}\right)}{\theta_{\text{g}}} \right)^{\frac{1}{n}}. \tag{5.26}$$

The mean climate sensitivity lies at 2.96 whereas θ_{g} is the 66% quantile (for a 66% safety target) and lies at 3.22. We already stated above that $\tilde{h}\left(\theta_{\text{g}}\right) \approx 0.5$ for a 2°C guard rail and 66% safety target. Together with the knowledge that $n > 1$ this allows the conclusion that $E_{\text{sw}} < E_{\text{g}} \approx E_{\text{min}}$. Therefore, case 3 is never realized and the following holds for the upper threshold of the climate sensitivity:

$$\theta_{\text{high}} = \frac{T_{\text{g}}}{\rho_1 E_{\text{min}}}. \tag{5.27}$$

This finding can explain the fact that temperatures gravitate towards a 2°C anomaly after learning (see Figure 2.7a). In other words, the guard rail acts as an attractor because it is the only optimum which is not at the emission bounds.

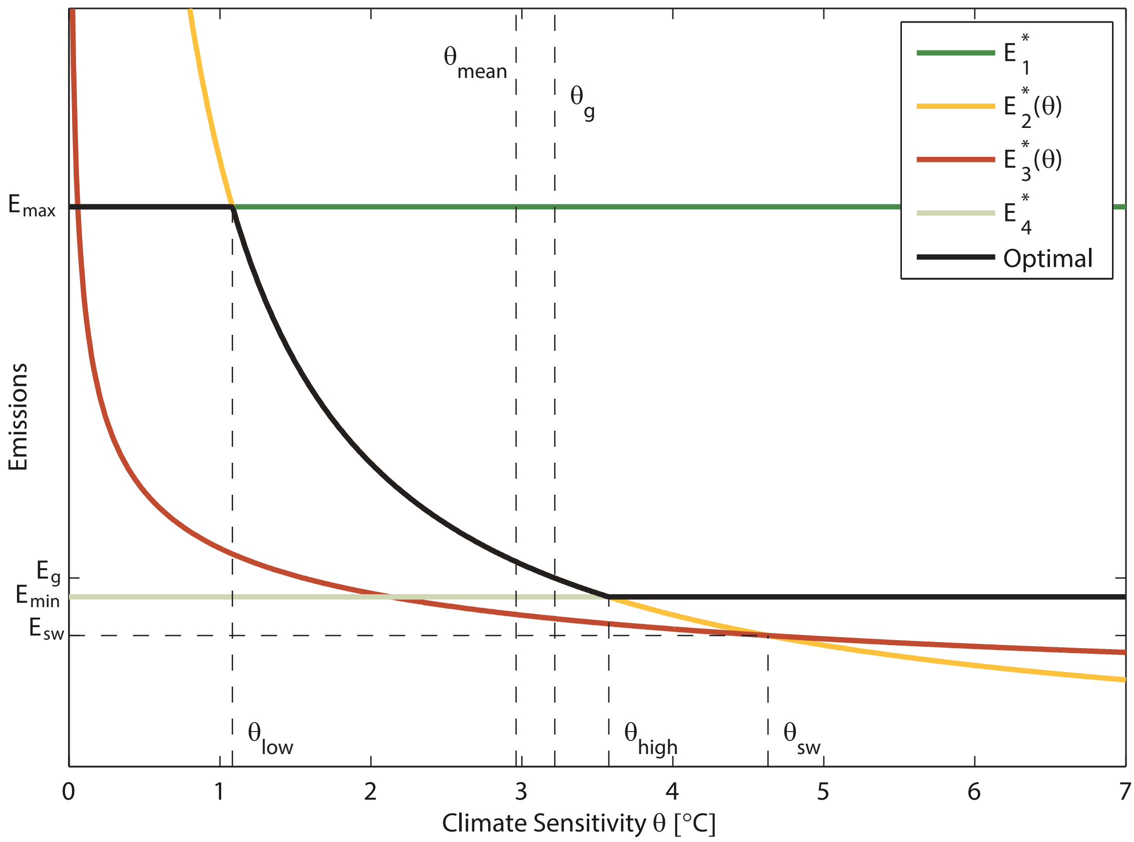

Figure 5.2. – Optimal emissions (black) for perfect information about the climate sensitivity θ. The optimal emissions for each of the four cases are also plotted in thick colored lines. Various values of θ and emissions of interest are marked. The assumptions made are: a target of 2°C and 66%; a minimum possible emissions level at 90% of the optimal emissions of the no information case; for the cost function: $n = 2$.

We close the analytic investigation with an example plot of optimal emissions shown in Figure 5.2. Additional assumptions needed to create this plot are boundaries (minimum and maximum) of possible emissions. The optimal emissions of the third case (red curve) are never globally optimal. During the entire Case 2, i.e. between

θ_{low} and θ_{high}, the temperature exactly meets the guard rail which allows for the assumption that the average temperature will be close to the guard rail. Another interesting aspect is that the emissions in the no information case (E_{g}) are at the lower end and so the value of incoming information is created by the possibility of increasing the emissions if climate sensitivity is learned to be smaller than θ_{g}.

Risk-related value of information originates from the reduction in temperature and emissions if climate sensitivity is learned to be above θ_{g} which only has a probability of 34% by construction.

5.3. Approximating the Welfare Function

An IAM in an optimization framework finds the optimal decisions, i.e. those maximizing welfare. The result is a point in the control space and there is little information about the surrounding points of such an optimal point. For real world applications it is relevant to know if the area around the optimum is flat or not. A flat optimum has the benefit that even if decisions deviate from the optimal solution by small amounts, the welfare loss is tolerable.

For MIND-L without learning we investigate the optimal solution to explore the shape of the welfare function in the optimum and to develop an understanding about what variations in decisions are acceptable. We will assume that 0.1% CBGE (certainty and balanced growth equivalent) is an acceptable change in welfare as it is in the range of year to year variability of gross world product. We hope to answer two questions:

1. What are the most important decisions? By how much can decisions vary without affecting the outcome by more than 0.1% CBGE.
2. Is there a time horizon for the decisions after which they become ineffective. A simulation of the decisions further into the future than such a horizon would not yield any benefit. This could optimize the horizon of the simulation for the economy in IAMs and improve the model speed.

To this end, we vary the important control variables, namely investment into fossil and renewable energy sectors over time. In total 52 variables are analyzed by taking 20000 draws from normal distributions centered around the respective optimal value and constrained by feasibility limits (refer to Section 5.3.1). The reference case is the NOLEARN scenario from Chapter 2. To estimate the welfare function we fit a second order function to the welfare losses with the extremum set to equal the optimal solution.

5.3.1. Control Variables

The control variables that are considered for study are investments into renewables and investments into fossil fuels between 2015 and 2140. This results in a vector of 52 variables (MIND-L works with 5 year steps) which we denote by $x = \begin{pmatrix} x_{\text{fossil}} & x_{\text{renew}} \end{pmatrix}$ and defined as:

$$x = I - I_{\text{opt}} \tag{5.28}$$

where I and I_{opt} are the test point and the optimal values of the investments respectively. The vector x represents deviations from the optimum control values and is determined by drawing random samples from a normal distribution with a zero mean for each variable. The normal distribution is the same for each variable and chosen with a variance that produces welfare losses of around 0.1 % CBGE, the acceptable level decided upon above. The standard deviation is equal to 0.25 trillion USD (which is about 1% of the initial production output of the economy in MIND-L, see Figure 5.3 for clarification). If I is less than the lower limit of investments for any of the 52 variables, that particular variable is bounded by the lower limit. Hence, the distribution of x_{fossil} is one-sided because the optimal investments are zero and can not be decreased (see 5.3a and c).

Figure 5.3a and b show the distributions of x_{fossil} and x_{renew} respectively whereby it can be seen that the distribution is equal over time. These deviations are added to the optimal investment paths and the resulting envelopes of investment paths in the sample is shown in red in Figure 5.3c and d. Note that the distributions are independent and identically distributed (i.i.d.), therefore, it is extremely unlikely that an investment path follows one of the borders of the envelope exactly.

After all 52 variables have been drawn, these values are used to fix the investments while all other decision variables are allowed to adjust during the optimization process. The decision maker is not allowed to “waste” investment, i.e. any investment into fossil fuels has to be converted into emissions eventually. The calculated welfare is translated into a CBGE percentage loss compared to the optimum.

The CBGE loss that this non-optimality induces is shown in Figure 5.4a and is on average around 0.1% CBGE by construction because we specifically chose the standard deviation of x to give welfare losses on this order of magnitude (by numerical approximation). This allows an estimate to be made on how much the investment decisions can vary without causing a significant loss of welfare, i.e. without diverging too far from the optimum and creating more than 0.1% CBGE loss.

Therefore, the envelopes shown in Figure 5.3c and d can be viewed as the allowable variation in investments, i.e. a variation with a variance of 0.25 trillion USD. What the analysis is lacking is the inclusion of time. The goal is to find the possible variances for each time step, that produce the same welfare effect. This can be achieved by performing a regression to find the effect of a change of investment in each time step.

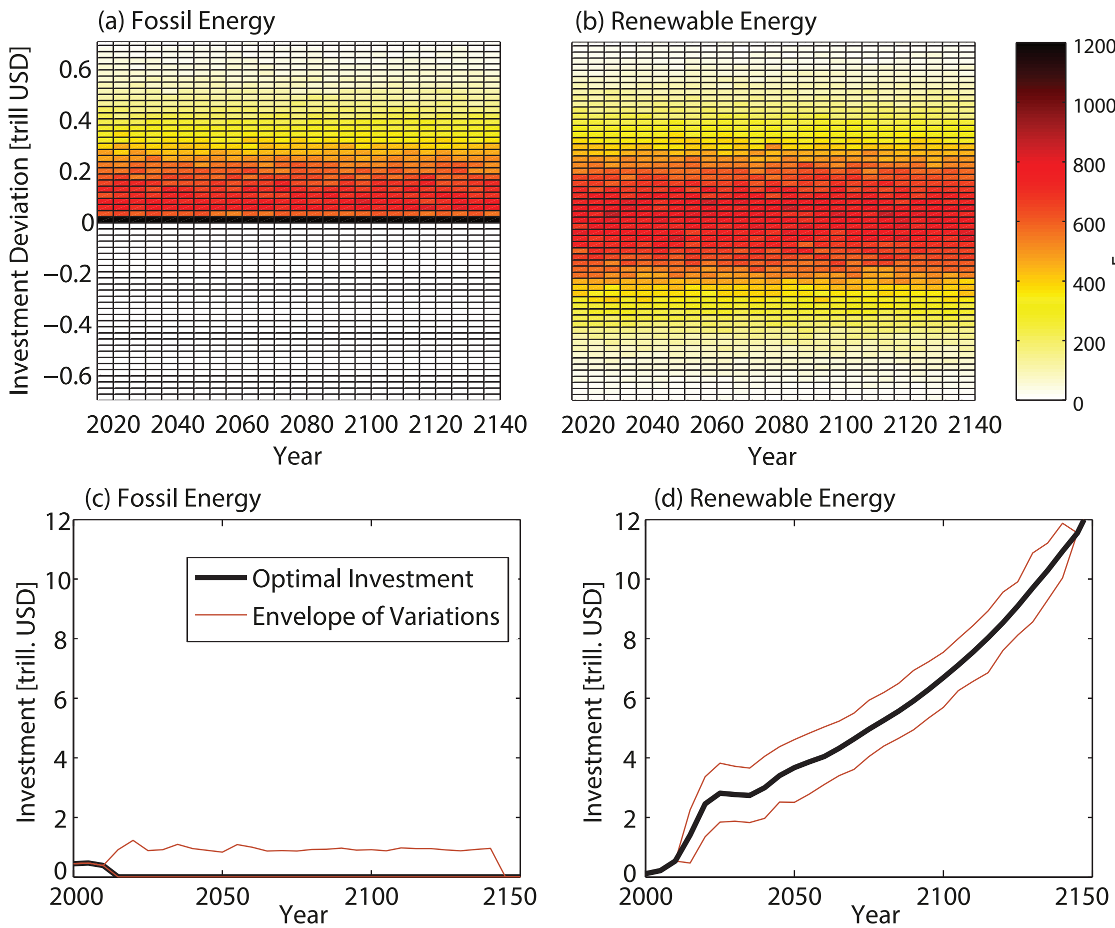

Figure 5.3. – Distribution of x, i.e. deviations of the variables from the optimum for (a) fossil energy and (b) renewable energy. The absolute investment paths are shown in (c) and (d). Fossil energy investment is close to zero in the optimum so negative deviations are not possible.

5.3.2. Approximating Function

The functional form of the fit for the regression has to be chosen in a way that forces the maximum to be at $x = 0$. This follows trivially from the fact that the point denotes the optimal solution in the model run, therefore, any deviation has to result in a loss of welfare. This is supported empirically by observing that there are no positive values in Figure 5.4a. For there to be a maximum the function has to have at least second order terms. Furthermore, the gradient at $x = 0$ has to be zero forcing all linear and constant terms to be zero. We are left with only the second order terms (1378 coefficients). The condition for $x = 0$ to be a maximum is that the Hessian has all negative eigenvalues. This condition is difficult to implement as a constraint on the regression and is therefore only checked in retrospect. The first 48 eigenvalues of the Hessian are negative whereas the last 4 are positive.

The root mean square error of the regression is equal to 0.0087% CBGE which is

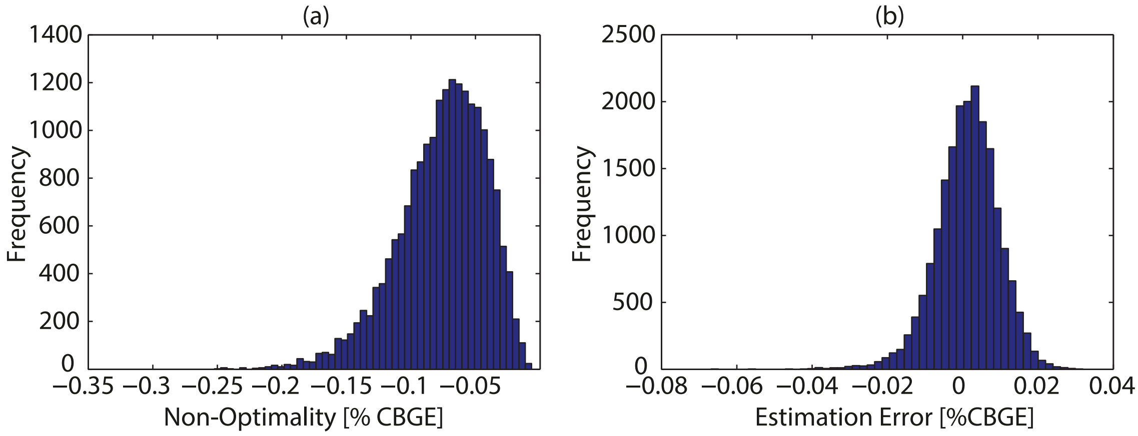

Figure 5.4. – (a) Distribution of the CBGE loss given in % with respect to the optimal solution for 20000 random draws of control vector x. (b) The difference between the second order fit estimation and the true deviations for the 20000 draws.

acceptable considering the order of magnitude of the non-optimality itself. The distribution of the error is plotted in Figure 5.4b. In summary, the approximation is well suited to represent MIND-L with CRA in the regime of 0.1% CBGE losses.

5.3.3. Coefficients

The resulting square coefficients are plotted in Figure 5.5. All coefficients that are not significant to a 5% level according to the student-t test p-value, or are zero, are white. Positive values are colored red and negative values are blue to black. There are only 464 significant coefficients which are non-zero of a total of 1378. The diagonal terms are the largest. The 95% confidence intervals, calculated from the regression, are plotted in Figure 5.5d for the diagonal terms. The mixed energy terms are the least significant as most values are small compared to the others. The units of the coefficients are utils per square US dollar and very specific to the model scaling and formulation, therefore, the absolute values are of little interest.

The first observation that can be made is that the welfare is more sensitive to changes in investments into fossil energy than investments into renewable energy, for changes of the same absolute magnitude. The reason for this is that fossil investments are on the whole much smaller than investments into renewables (as seen in Figure 5.3c and d). Especially early investments into fossil fuels have a strong negative effect on welfare.

After 2050, both types of investments have roughly the same influence which is much lower than in early years. The reason for this results from the fact that any investments after 2050 do not affect the climate risk as much anymore and therefore the welfare loss is merely due to the consumption loss which is the same for both types of investments.

Lastly, the effects of a change in one type of investment is largely independent from a change in the other type as can be seen by the random pattern in Figure 5.5c. The only notable mixed terms are for fossil investments at different time instances in (a). This is due to the budget effect for emissions. If a sum was invested into fossil fuels in one time instance then any subsequent investment will be worse because the overall cumulative emissions are raised.

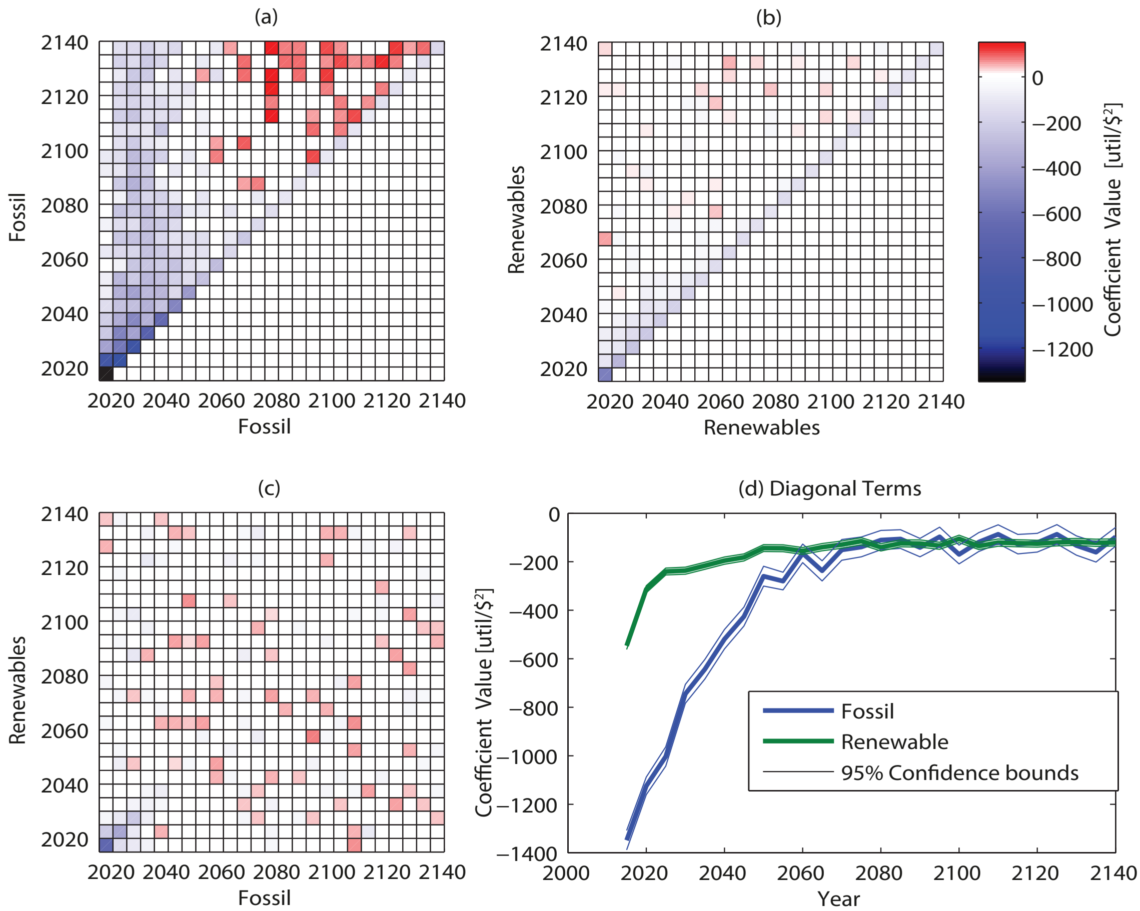

Figure 5.5. – Coefficients of second order polynomial fit to estimate welfare of MIND-L around the optimum of a NOLEARN scenario. The coefficients are divided into groups where (a) both variables are investments into fossil energy, (b) both are investments into renewable energy and (c) mixed terms. In Pane (d) the diagonal terms are shown on their own with 95% confidence bounds. Only significant coefficients (5% level) are shown, all others are colored white. The main finding is that the investments into fossil fuels in the first decades are the most important.

5.4. Discussion

The analytical model of CRA for the climate problem allows us to find analytical expressions for the optimal emissions. We find four possible regimes for the optimal emissions in a perfect information scenario depending on the climate sensitivity that is learned. The first case represents the decision to do no mitigation because climate sensitivity is very low. In the second case, emissions are chosen to exactly hit the guard rail whereas the third case consists of the actual trade-off and temperatures cross the guard rail. The fourth case represents maximum mitigation.

Due to the fact that climate change is already quite advanced, the widely discussed 2°C target moves closer and closer to the feasibility limit (as modeled in MIND-L). We show that this implies that an actual trade-off between risk and cost (Case 3) is not made, instead the decision is made to either stay below the guard rail or pass it as little as possible (by doing maximum mitigation). In the full dynamical model this is not completely true but the probability share of cases that go into a trade-off is around 16% of which most only violate the target for a few decades.

The second half of the chapter concludes that a variability of investments of up to 0.25 trillion USD (2012) creates about 0.1% CBGE welfare loss which is usually treated as acceptable. The approximation of the welfare equation is intended to be an exploratory study to gain insight into shape of the optimum. We found the optimum to be flat with respect to investments after 2050 which leaves the possible deduction that it suffices to represent the economy until around 2060. Another possible inference is that learning after 2050 is of little value.

In contrast to a CEA, which produces infeasibilities if a small change in the investment decisions leads to the violation of the target, CRA allows for such an analysis although a climate target is included in a fashion.

6. Conclusion

Cost-Risk Analysis (CRA) is, so far, the only tool that enables a community that supports climate targets to assess climate policies while including the possibility of resolving uncertainty at some point in the future. CRA was first proposed for the climate problem by Schmidt *et al.* (2011) and is a trade-off between expected utility from consumption and expected utility loss from increasing temperatures, defined by the risk metric.

This thesis applies the method and discusses the adjustments necessary to yield a productive analysis. For the analysis, we borrow the concept of degree years as a risk measure for climate change from Schneider & Mastrandrea (2005) and further assume that a community exists which supports a climate target of staying below 2°C temperature rise with 66% probability (called "safety"). This assumption is partly based on discussions of the UNFCCC (2011).

In the following, the main findings and contributions of this thesis are summarized and we conclude with an outlook onto possible future applications and evolutions of CRA.

6.1. Summary

The contributions and findings are divided into three sections: (i) the framework, including all aspects that are necessary to calculate and understand the value of information, (ii) robustness of CRA, showing how CRA reacts to changes of key parameters and finally, (iii) other aspects of a learning event besides the value of information.

(i) A Working CRA Framework

The originally proposed CRA (Schmidt *et al.*, 2011) focused on a risk metric that was based on the probability of violating a guard rail as it reflects Cost-Effectiveness

Analysis (CEA) as close as possible. We showed that this leads to “sacrificing” if the climate sensitivity is learned to be very high, i.e. mitigation is abandoned. To correct this behavior, which is not in line with the assumed preferences of the community endorsing climate targets, the risk metric is based on expected discounted degree years (expected and discounted area below the temperature path and above the guard rail). This prevents the sacrificing behavior as there is always an incentive to reduce temperatures, even for small reductions.

The remaining degree of freedom for the risk metric is the trade-off parameter connecting risk to utility. The trade-off parameter sets the marginal value of reducing temperatures and its value was not discussed specifically in previous publications. We use a calibration technique, previously unexplored in climate science, to set the trade-off parameter to the lowest value that still produces an optimal solution that fulfills the climate target. The set up of the model for the calibration should reflect the situation in which the target was devised so that the preferences of the community can be properly reflected. The calibration in this thesis assumes a scenario without future learning.

The CRA solution was compared to CEA with chance constrained programming, with the same climate target. We find that optimal decisions are similar until 2050 and then diverge, where CRA mandates more mitigation and therefore leads to lower temperatures than CEA. In total the cost of mitigation increases from 1.30% (in CEA) to 1.52% CBGE[1]. We infer that CEA simulation results remain useful for short term decisions if an uncertainty is considered that is not resolved.

By the novel approach of decomposing a change in welfare along three dimensions we are able to determine the origin of such a change. The three dimensions are time, states of the world and the differentiation between economic and risk-related sources. The decomposition is done by linearizing and rescaling the CBGE equation. Applied to the expected value of information (EVPI) and expected cost of climate protection (ECCP, mitigation costs together with monetized risk) the following conclusions can be summarized:

1. ECCP is divided roughly equally across economic and risk-related sources (see Table 6.1) but is strongly skewed in the time dimension. The risk-related ECCP mainly originates after 2050 and only from states of the world with a climate sensitivity above 3.5°C, whereas the economic ECCP originates mainly before 2050 and equally across all states of the world. This confirms the often assumed conceptual time scale division between economic action and climate impacts.
2. EVPI for learning in 2015 compared to learning in 2075 results in a total value of 0.66% CBGE. About one third is due to risk-related changes, i.e. a reduction in temperature, and the rest is due to consumption increases. Most

[1]CBGE is the certainty and balanced growth equivalent and can be interpreted as a change in consumption each year. See Anthoff & Tol (2009) for details. Applied to 2012 gross world product and assuming a 75% consumption share, 1% CBGE is equal to about 500 billion USD.

of the economic value originates from the time between the compared learning events (2015 – 2075) as this is the interval in which an information advantage exists. The states of the world with a climate sensitivity below 3.5°C generate most of the EVPI. The reason being that learning is most valuable if it allows to increase emissions.

3. Economic EVPI makes up around 1/3 of economic ECCP (i.e. the cost of mitigation). This is a substantial reduction and is comparable to the added value of the carbon capture and storage technology (Luderer *et al.*, 2011). The reason for such a large value lies in the fact that, if information is not available, the decision maker is forced to do close to maximum mitigation in order to reach the target. Whereas, if available, around 66% of the states of the world can emit more than before.

A summary of the EVPI and ECCP is given in Table 6.1.

% CBGE	Economic	Risk-Related	Total
EVPI(2015,2075)	0.46	0.20	0.66
ECCP(2075)	1.52	1.93	3.45

Table 6.1. – EVPI(2015,2075) split into economic and risk-related terms and compared to their ECCP counterparts. The economic ECCP is interpreted as the cost of mitigation. Values are given in percentage CBGE which gives the certainty equivalent change in initial consumption that would produce the same difference in welfare for the same utility function and identical growth rates.

(ii) Robustness of CRA

CRA exhibits a robust behavior with respect to many aspects of the decision problem. These are summarized in the following:

1. If the risk metric is changed, for example by changing the linear penalty above the guard rail to a quadratic function, the recalibration of the trade-off parameter to the given target compensates the change in marginal risk. Therefore, the impact of the form of penalty function on the analysis is greatly reduced. The EVPI, as well as the decisions, stay stable for a variety of penalty functions above the guard rail. We tested functions ranging from a logarithmic function up to a fourth order polynomial penalty function.
2. If the social discount rate (SDR) is varied[2], the resulting optimal climate policy is only marginally affected (optimal emissions only deviate around the year 2040 and only for extreme values of 4%/yr as a pure rate of time preference). The calibration cancels out most of the effect of the SDR on the valuation of

[2]The social discount rate, or interest rate, is chosen implicitly by setting the pure rate of time preference and the constant relative risk aversion.

future climate risks. This is one of the strengths of CRA: it allows for the partial substitution of the choice of SDR by the choice of climate target as a valuation of climate effects. In contrast, the optimal decisions of a standard cost-benefit approach with a damage function, are more sensitive to the choice of SDR.

3. Although the effect of SDR on the optimal climate policy is small (see point 2), a misjudgment of the SDR can create a loss of welfare due to the resulting suboptimal climate policy. This welfare loss is found to be moderate, around 0.2% CBGE for a misjudgment of 2%/yr. An underestimation of SDR has less impact than an overestimation. We therefore recommend using values at the lower end of the commonly used SDR, in line with Weitzman (2007a). This also agrees well with ethical discussions about how the climate problem should be discounted.

4. The safety for a case with learning is similar (ca. 70%) to the calibrated safety target of 66%. This implies that our choice to do a calibration with the no-learn scenario is not an important choice. If a learning scenario were to be chosen, similar results would be obtained, underlining the robustness of CRA. The underlying mechanism for this robustness is that 66% safety is close to the maximum feasible safety in MIND-L of 74%.

5. An aspect of robustness is the sensitivity of the welfare to small variations in the decisions. CRA with a climate target and without learning is found to give a flat optimum relative to decisions after 2050. The most important policy decisions are the investments into fossil fuels before 2050. The allowed variability of investment decisions that produces welfare losses of only 0.1% CBGE has a standard deviation of 0.25 trillion USD. This is substantial compared to a value of around 4 trillion USD for investments into renewables in 2050. We conclude that it is not necessary to follow the optimal decisions exactly and a certain economic variability is allowed without threatening the welfare substantially. In the presence of thresholds such as a climate target, this is generally not the case because a small variation can lead to the violation of the target. The CRA approach to targets is robust w.r.t. this aspect.

(iii) Effects of Perfect Learning in CRA

Due to the relatively unexplored field of learning under climate targets, CRA has enabled some conclusions relating to learning:

1. Changing the target by adjusting the guard rail or adjusting the probability target has different effects only if learning is considered. Generally it is expected that relaxing a target yields a lower value of information, as the cost of mitigation is also lower. However, relaxing the target by increasing the guard rail increases the EVPI for information after 2020 (due to more states of the

world doing zero mitigation) although EVPI is decreased if the target is relaxed by decreasing the necessary safety. This strengthens the need for clear communications and formulations of a climate target and specifically how it is adjusted, if need be.

2. We distinguish four possible regimes ("cases") in which the optimal decisions can lie depending on what value of the climate sensitivity is learned and we calculate the optimal emissions for a static problem. The first regime (Case 1) plays a minor role (5% of state of the worlds) and includes the business as usual decision if climate sensitivity is very low and the guard rail is not reached. Due to the rather stringent target of 2°C and 66% safety, the range of climate sensitivity that leads to a trade-off between cost and risk (Case 3) is very small (16%) and even zero for a static case. Most of the the solutions (65%) drive temperatures to just under the guard rail, therefore preventing any risk (Case 2). A long term violation of the target (Case 4) is due to learning of a climate sensitivity which is so large that maximum mitigation is reached (about 14%). Over long periods, the temperature will gravitate towards the guard rail, which is a consequence of the structure of the penalty function, also expressed by the small share of Case 3.

3. The EVPI is influenced greatly by the choice of SDR although optimal decisions are not effected (as stated in the 2. point of the previous section). Both, EVPI and ECCP, peak at an SDR of about 5%/yr and for extreme SDR (2%/yr-14%/yr) a reduction by an order of magnitude is possible. This is due to a matching of time scales of SDR and the climate problem itself. High SDRs will devalue most of the relevant period between 2015 and 2100 and low SDR will also reduce its weight because of increasing weight given to the time after 2100. This means that for extremely low SDR values, counter-intuitively, the value of information is strongly reduced, as shown in Figure 4.2b, and with it the incentive for research. The effect that giving the future more value would reduce incentive for research seems odd and hints at the possibility that the climate problem and the economy should be discounted at different rates.

Considering the thesis as a whole, there are two mechanisms that mainly drive the results: (i) the robustness of the calibration technique and (ii) the fact that a 2°C target with a likely chance is increasingly difficult to reach. The calibration technique is a powerful tool to include preferences into an analysis. Instead of formalizing all possible effects of climate change into one damage function under great uncertainty, we make use of the already aggregated preferences present in the formulation of a climate target. In this way we take some pressure off the quest to evaluate all possible impacts and provide a tool to examine climate policy presuming that society is in favor of a specific goal until uncertainty around impacts are reduced sufficiently.

In general, the thesis supports increased investment into the reduction of the uncertainty around the climate sensitivity as well as the transient climate response. The expected value of resolving the uncertainty in the next decade is on the order of 100 billion USD (in 2012) per year. We hope that this thesis motivates scientists and policy makers to consider the implications when formulating climate targets under uncertainty and be clear about their preferences in the light of future learning.

6.2. Outlook

The calculation of the value of information, as was done in this thesis, allows for the estimation of the value of projects that reveal such information. This thesis concentrates on the uncertainty present in the climate response, hence the value of information found can indicate the value of research and projects that are able to reduce the uncertainty on the climate response. As Stevens & Bony (2013) point out, research about the physicochemical properties of water in the atmosphere could reduce this uncertainty. Increased measurement accuracy of the temperature anomaly (also studied in Cooke *et al.* (2013)) and the ocean heat content would also decrease the uncertainty. In general, measuring campaigns such as the Argo system[3] (3560 robotic probes deployed across all oceans) or remote sensing satellites contribute to the reduction in uncertainty and all these investments work towards harnessing part of the EVPI found in this thesis.

The structure of CRA that was used here, allows for the choice of a pure rate of time preference and a constant relative risk aversion to describe the economic behavior. To ensure time consistency, the risk-related utility loss was discounted with the same pure rate of time preference as was used to discount the utility stream from consumption. However, the fact that society might have more complex preferences raises the question how these can be uncoupled and how the emerging time inconsistency could be handled. To some extent this is done by calibration to a climate target, however, possible future studies might separate these “normative” parameters into economic and climate related parts. Furthermore, the constant relative risk aversion also brings with it an elasticity of inter-temporal substitution. Traeger (2014) has suggested to decompose the two effects to better map the preferences of society.

There are other possible applications of CRA to problems that CEA struggles with. Bosetti *et al.* (2009) and Luderer *et al.* (2011) study the effect of late participation with CEA but are limited to studying time shifts that still allow compliance with the target. Luderer *et al.* (2011) find that a delayed participation until the year 2030 does not allow a 450 ppm CO_2 target in 2100. If participation is delayed until 2020 the authors allow for an overshoot before 2100 to establish feasibility. If large uncertainty and learning is included in the analysis, it becomes increasingly difficult to produce a feasible CEA (with chance constrained programming) simulation. The underlying reason is that CEA is not an expected utility framework and therefore a violation of the climate target constraint cannot be weighed against the danger

[3]http://www.argo.ucsd.edu/

of doing so. CRA allows for overshooting the target indefinitely and therefore stays operational even after elongated times of inaction which opens the possibility to do thorough analysis of delayed participation.

CRA is designed to bridge the gap until impact models have matured. But when is the point reached where cost-benefit analysis with damage functions should be preferred over CRA? A hybrid decision framework is also possible by including all the specific impact functions that are known and augmenting the analysis with a calibrated risk function to cover impacts that were left out. Such hybrid models promise to be an interesting field of future research.

A. Appendix

A.1. Numerical Considerations

Approximating the Heaviside Function

To ensure that the solver CONOPT can cope with the non-differentiable functions (Heaviside function) they are approximated by error functions. This somewhat blurs the guard rail targets on the order of 0.1°C. This is acceptable because natural variability is also on this order of magnitude.

Finite Horizon Effect

The model runs until year 2200, i.e. a finite horizon. Effects that can arise include:

1. As the end approaches, the decision maker will cease to invest into mitigation because the effects will only manifest beyond the horizon. This results in heavy investment into fossil fuels in the last time steps of the model run for a climate protection scenario.

2. The discount factor is not at zero in 2200, so that theoretically the years that come beyond the horizon should still have an influence on the welfare. Without them the importance is skewed towards the beginning, i.e. the future has less importance than would be the case with an infinite horizon.

To reduce the impact of the finite horizon effect, the very last time step is discounted by a value that includes the geometric sum of the discount factor continued into infinity. In effect, this pretends that the economy is continued with the constant consumption level of the last time step into infinity. We consider only the results up to 2150 for the plots.

Probability of Violation

Including a calculation of the safety (1 - probability of violating the guard rail) is necessary in a CEA in the form of chance constrained programming. As CRA is a trade-off analysis this is not necessary anymore. However, during the calibration a measure of probability has to be introduced. The probability can either be calculated per time slice (soft) or per temperature path (hard). As the calibration is done in the NOLEARN scenario, in which the temperature paths of different SOWs cannot cross each other, soft and hard crossing probabilities are the same. This strengthens CRA as it does not depend on the definition of the probability measure. For reference the equations for the different variants are given here.

The soft variant calculates the probability per time slice and then looks for the maximum:

$$P_{\text{soft}} = \max_t \left\{ \sum_{s=0}^{S} p_s \Theta \left(T(t,s) - T_{\text{g}} \right) \right\}. \tag{A.1}$$

Here Θ represents the Heaviside function. The hard variant has a stronger condition:

$$P_{\text{hard}} = \sum_{s=0}^{S} p_s \Theta \left(\max_t \{T(t,s)\} - T_{\text{g}} \right) \tag{A.2}$$

No matter when the SOW crosses the guard rail, it is counted into the total probability of violation.

When calculating the probability for the LEARN cases in retrospect, we adopt the hard interpretation by extracting the maximum temperature of each realization of climate sensitivity. The probability is found by sorting the maximum temperatures and giving each value a 5% probability quantile through which a cumulative distribution is constructed (see Figure 2.6).

Optimization procedure

To make sure that the solution of the optimization is optimal, a strategy was followed that tries to work against the effect of multiple optima. The concept relies on the fact that if multiple optima exist they will be characterized by one solution with higher temperatures on the whole and one solution with lower temperatures. Any GAMS optimization can be started with an initial guess. The strategy consisted of starting the model once with high temperatures (BAU case) and once with low temperatures (NOLEARN case). The solutions used for the initial guess originate from previous runs.

The next step in the strategy is to restart the optimization four times to guarantee that the solver found an optimum. The solutions for both strategies tended to converge. We conclude that no relevant local optima exist that could complicate the optimization.

Convergence for Infinite Climate Sensitivity

Weitzman (2009) opened a discussion by showing that expected damages can turn out to be infinite. For the approach in this thesis the equivalent case would be if expected risk were to be infinite. This is not the case due to the log-normal distribution and the fact that the risk is linear in climate sensitivity as shown in Figure A.1.

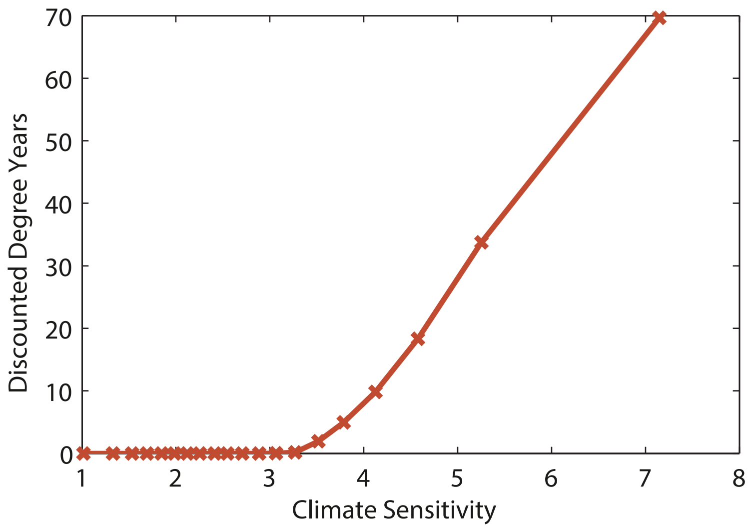

Figure A.1. – Discounted degree years (“risk”) for different values of climate sensitivity and for the same emissions.

Assuming the linear dependency of the risk on the climate sensitivity θ, the convergence can be tested by calculating the integral over the domain of the climate sensitivity by assuming a log-normal distribution $f(\theta)$ with parameters μ and σ. The integral evaluates to less than infinity as it is the definition of the expected value of the distribution:

$$\int_0^\infty \theta f(\theta) \mathrm{d}\theta = e^{\mu + \frac{\sigma^2}{2}} < \infty. \qquad \text{(A.3)}$$

Therefore, by using a discounted degree years as the risk metric, we avoid the dismal proposition so long as the probability distribution has a finite mean value.

We tested the convergence numerically by choosing a sampling technique which includes the 99% quantile of the distribution and we found no substantial change in model behavior.

A.2. Descriptive Sampling

To be able to sample the distribution representatively, this thesis uses the same strategy as was used by Lorenz *et al.* (2012b) and is similar to the ideas in Saliby (1990) but instead of finding equiprobable samples numerically we do it analytically. The derivation follows two steps: (i) division of the distribution into equally probable intervals and (ii) representation of the intervals by mean value. The calculations are shown here.

The mean value $\bar{\theta}$ of a log-normal distribution $f(\theta)$ is defined as follows:

$$\bar{\theta} = \int_0^\infty \theta f(\theta) d\theta \tag{A.4}$$

To find N samples that have the same probability, the expected value of each of the N quantiles has to be found. The above integral only works because the integral over $f(\theta)$ is unity. If the mean value over only a part of the distribution is to be found, the result has to be normalized. Over an arbitrary interval $[a, b]$ the expected value equals:

$$\theta_{a,b} = \frac{\int_a^b \theta f(\theta) d\theta}{\int_a^b f(\theta) d\theta} \tag{A.5}$$

The intervals for N samples are $[(n-1)/N, n/N]$ for n between 1 and N. The denominator evaluates to $F(b) - F(a)$, i.e. for the samples of interest: $1/N$. Assuming F^{-1} is the inverse cumulative distribution function then the nth sample is given by:

$$\theta_n = N \int_{F^{-1}\left(\frac{n-1}{N}\right)}^{F^{-1}\left(\frac{n}{N}\right)} \theta f(\theta) d\theta \tag{A.6}$$

This can be solved analytically to give:

$$\theta_n = \frac{1}{2}\bar{\theta}N \left(\mathrm{Erf}\left[\frac{\sigma}{\sqrt{2}} + \mathrm{InverseErfc}\left[\frac{2(n-1)}{N}\right]\right] - \mathrm{Erf}\left[\frac{\sigma}{\sqrt{2}} + \mathrm{InverseErfc}\left[\frac{2n}{N}\right]\right]\right), \tag{A.7}$$

where Erf and InverseErfc correspond to the error function and the inverse of the complementary error function. The parameter σ is the second parameter for the log-normal distribution. This is a analytically exact way of calculating the samples. The beauty lies in the fact that

$$\frac{\sum_n \theta_n}{N} = \bar{\theta}. \tag{A.8}$$

We used $N = 20$, visualized in Figure A.2 and the values are given explicitly in Table A.1.

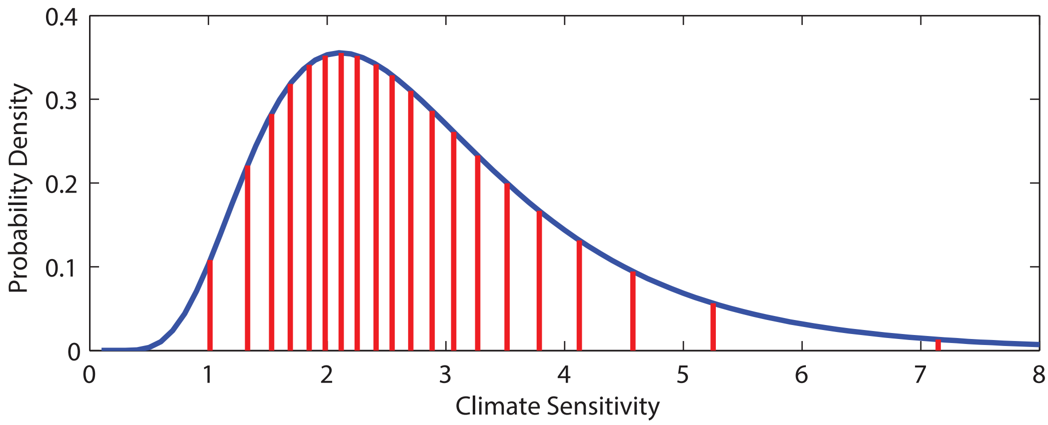

Figure A.2. – Sampling of the climate sensitivity distribution, $\mathcal{LN}(0.973, 0.4748)$, taken from Wigley & Raper (2001).

SOW	CS [°C]	CDF [%]	SOW	CS [°C]	CDF [%]
1	1.01	2.1	11	2.73	52.5
2	1.33	7.4	12	2.90	57.5
3	1.53	12.5	13	3.08	62.5
4	1.70	17.5	14	3.28	67.5
5	1.85	22.5	15	3.52	72.5
6	1.99	27.5	16	3.79	77.5
7	2.13	32.5	17	4.13	82.6
8	2.27	37.5	18	4.58	87.6
9	2.42	42.5	19	5.27	92.6
10	2.57	47.5	20	7.17	98.2

Table A.1. – Climate sensitivity (CS) and the value of the cumulative distribution function (CDF) for each of the 20 SOWs.

A.3. CBGE Linearization Error

In Chapter 3 the change of CBGE between two simulations is calculated as a measure of welfare loss and then split into parts according to the differences in welfare. Here we show the comparison of the actual total CBGE with the sum of the CBGE parts. In the notation used in the chapter the error that is plotted in Figure A.3 is given by:

$$\text{Error} = \left(\frac{\tilde{\Delta}}{\Delta} - 1\right) * 100. \tag{A.9}$$

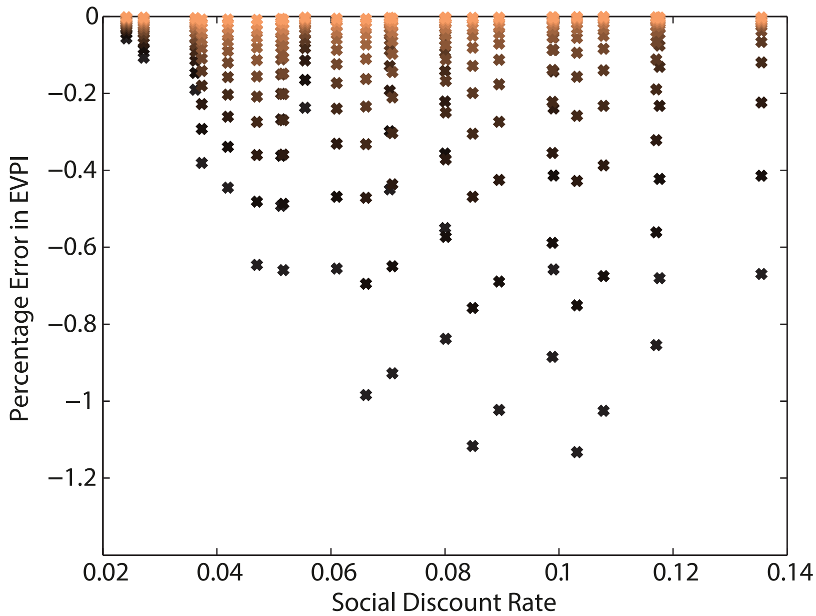

Figure A.3. – The relative percentage error of EVPI that is made when linearizing the CBGE calculation for different SDR. Color gradient indicates the learning point. The lighter the color, the later the learning point. Learning points lie between 2015 and 2075.

A.4. Analytical Welfare Equation

For the interested reader the complete expression for the welfare in a perfect information case and assuming $\theta_{\text{high}} > \theta_{\text{sw}}$ is given below. The probability distribution of the climate sensitivity is given by $f(\theta)$ and the cumulative distribution by $F(\theta)$. For the case where the above inequality does not hold, it suffices to set $\theta_{\text{high}} = \theta_{\text{sw}}$ to attain the correct expression.

$$\hat{W} = \hat{W}_C + \hat{W}_R \tag{A.10}$$

$$
\begin{aligned}
\hat{W}_C = & - E_{\max}^{-n} F(\theta_{\text{low}}) \\
& - \left(\frac{T_{\text{g}}}{\rho_1}\right)^{-n} \int_{\theta_{\text{low}}}^{\theta_{\text{sw}}} \theta^n f(\theta) d\theta \\
& - E_{\text{g}}^{-n} \left(\frac{\rho_1}{h\left(E_{\text{g}}\right)}\right)^{\frac{n}{n+1}} \int_{\theta_{\text{sw}}}^{\theta_{\text{high}}} \theta^{\frac{n}{n+1}} f(\theta) d\theta \\
& - E_{\min}^{-n} \left(1 - F\left(\theta_{\text{high}}\right)\right)
\end{aligned}
\tag{A.11}
$$

$$
\begin{aligned}
\hat{W}_R = & - \frac{n E_{\text{g}}^{-n-1}}{h\left(E_{\text{g}}\right)} \Bigg(E_{\text{g}} \left(\frac{\rho_1}{h\left(E_{\text{g}}\right)}\right)^{\frac{n}{n+1}} \int_{\theta_{\text{sw}}}^{\theta_{\text{high}}} \theta^{\frac{n}{n+1}} f(\theta) \mathrm{d}\theta \\
& + E_{\min} \rho_1 \int_{\theta_{\text{high}}}^{\infty} \theta f(\theta) \mathrm{d}\theta \\
& - T_{\text{g}} \left(1 - F\left(\theta_{\text{sw}}\right)\right) \Bigg)
\end{aligned}
\tag{A.12}
$$

Bibliography

Acemoglu, D., Aghion, P., Bursztyn, L. & Hemous, D. (2012). The environment and directed technical change. *American Economic Review*, 102(1), pp. 131–66.

Ackerman, F. & Munitz, C. (2012). Climate damages in the FUND model: A disaggregated analysis. *Ecological Economics*, 77, pp. 219–224.

Allen, M.R. & Frame, D.J. (2007). Call off the quest. *Science (New York, N.Y.)*, 318(5850), pp. 582–3.

Allen, M.R., Frame, D.J., Huntingford, C., Jones, C.D., Lowe, J.a., Meinshausen, M. & Meinshausen, N. (2009). Warming caused by cumulative carbon emissions towards the trillionth tonne. *Nature*, 458(7242), pp. 1163–6.

Andronova, N., Schlesinger, M., Dessai, S., Hulme, M. & Li, B. (2007). The concept of climate sensitivity: history and developement. In: *Human-induced Climate Change: An interdisciplinary Assessment*, chapter 1, Cambridge university Press.

Anthoff, D. & Tol, R.S.J. (2009). The Impact of Climate Change on the Balanced Growth Equivalent: An Application of FUND. *Environmental and Resource Economics*, 43(3), pp. 351–367.

Arrow, K., Cropper, M., Gollier, C., Groom, B., Heal, G.M., Newell, R.G., Nordhaus, W.D., Pindyck, R.S., Pizer, W.A., Portney, P.R., Sterner, T., Tol, R.S.J. & Weitzman, M.L. (2013). How should benefits and costs be discounted in an intergenerational context? The Views of an Expert Panel (December 19, 2013). *Resources for the Future Discussion Paper No. 12-53*.

Athanassoglou, S. & Xepapadeas, A. (2012). Pollution control with uncertain stock dynamics: When, and how, to be precautious. *Journal of Environmental Economics and Management*, 63(3), pp. 304 – 320.

Azar, C. & Lindgren, K. (2003). Catastrophic events and stochastic cost-benefit analysis of climate change. *Climatic Change*, 56(3), pp. 245–255.

Baker, E. & Shittu, E. (2008). Uncertainty and endogenous technical change in climate policy models. *Energy Economics*, 30(6), pp. 2817–2828.

Barrage, L. (2012). Optimal Dynamic Carbon Taxes in a Climate-Economy Model with Distortionary Fiscal Policy. *Working Paper.*

Blau, R. (1974). Stochastic programming and decision analysis: an apparent dilemma. *Management Science*, 21(3), pp. 271–276.

Bordley, R.F. & Pollock, S.M. (2009). A Decision-Analytic Approach to Reliability-Based Design Optimization. *Operations Research*, 57(5), pp. 1262–1270.

Bosetti, V., Carraro, C., Galeotti, M., Massetti, E. & Tavoni, M. (2006). WITCH - A World Induced Technical Change Hybrid Model. *University Ca' Foscari of Venice Economic Research Paper No. 46/06.*

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

Bruckner, T. & Zickfeld, K. (2008). Inverse Integrated Assessment of Climate Change: the Guard-rail Approach. In: *International Conference on Policy Modeling (EcoMod2008).*

Cline, W.R. (1992). *The economics of global warming.* Peterson Institute.

Cooke, R. (2013). Model uncertainty in economic impacts of climate change: Bernoulli versus Lotka Volterra dynamics. *Integrated Environmental Assessment and Management*, (2006), pp. 1–11.

Cooke, R., Wielicki, B.a., Young, D.F. & Mlynczak, M.G. (2013). Value of information for climate observing systems. *Environment Systems and Decisions.*

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

den Elzen, M.G.J., Meinshausen, M. & van Vuuren, D.P. (2007). Multi-gas emission envelopes to meet greenhouse gas concentration targets: costs versus certainty of limiting temperature increase. *Global Environmental Change*, 17(2), pp. 260–280.

den Elzen, M. & van Vuuren, D.P. (2007). Peaking profiles for achieving long-term temperature targets with more likelihood at lower costs. *Proceedings of the National Academy of Sciences of the United States of America*, 104(46), pp. 17931–17936.

Dietz, S. & Matei, A. (2013). Spaces for agreement: a theory of time-stochastic dominance. *Working Paper*, (153).

Edenhofer, O., Bauer, N. & Kriegler, E. (2005). The impact of technological change on climate protection and welfare: Insights from the model mind. *Ecological Economics*, 54(2-3), pp. 277–292.

Edenhofer, O., Knopf, B. & Barker, T. (2010). The economics of low stabilization: Model comparison of mitigation strategies and costs. *The Energy Journal*, 31, pp. 11–48.

Frame, D.J. (2005). Constraining climate forecasts: The role of prior assumptions. *Geophysical Research Letters*, 32(9), p. L09702.

Fritz, S., Scholes, R., Obersteiner, M., Bouma, J. & Reyers, B. (2008). A conceptual framework for assessing the benefits of a global earth observation system of systems. *Systems Journal, IEEE*, 2(3), pp. 338–348.

Garnaut, R. (2008). *Garnaut climate change review: Interim report to the Commonwealth, state and territory governments of Australia*. Garnaut Climate Change Review.

Geden, O. (2013). Modifying the 2Â°C Target. *German Institute for International and Security Affairs*.

Gollier, C. (2004). *The economics of risk and time*. The MIT Press.

Golub, A., Narita, D. & Schmidt, M. (2013). Uncertainty in integrated assessment models of climate change: Alternative analytical approaches. *Environmental Modeling & Assessment*, pp. 1–11.

Groom, B., Hepburn, C., Koundouri, P. & Pearce, D. (2005). Declining discount rates: The long and the short of it. *Environmental and Resource Economics*, 32(4), pp. 445–493.

Ha-Duong, M., Grubb, M. & Hourcade, J. (1997). Influence of socioeconomic inertia and uncertainty on optimal CO2-emission abatement. *Nature*, 390.

Ha-Duong, M. (1998). Quasi-option value and climate policy choices. *Energy Economics*, 20(5-6), pp. 599–620.

Heal, G. & Millner, A. (2013). *Discounting under Disagreement*. Working Paper 18999, National Bureau of Economic Research.

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

Hope, C. (2011). The Social Cost of CO_2 from the Page09 Model. *Economics Discussion Paper No. 2011-39*.

Howitt, R. (1995). Positive mathematical programming. *American journal of agricultural economics*, 77(2), pp. 329–342.

Iverson, T. & Perrings, C. (2012). Precaution and proportionality in the management of global environmental change. *Global Environmental Change*, 22(1), pp. 161–177.

Jagannathan, R. (1985). Use of sample information in stochastic recourse and chance-constrained programming models. *Management Science*, 31(1), pp. 96–108.

Keller, K., Kim, S.R., Baehr, J., Bradford, D.F. & Oppenheimer, M. (2007). What is the economic value of information about climate thresholds? *Human-Induced Climate Change: an Interdisciplinary Assessment, SchlesingerM, KheshgiH, SmithJ,* *De La ChesnayeF, ReillyJM, WilsonT, KolstadC (eds). Cambridge University Press: Cambridge*, pp. 343–354.

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

Kleinen, T. (2005). *Stochastic information in the assessment of climate change*. Dissertation, University of Potsdam, Germany.

Kriegler, E. & Bruckner, T. (2004). Sensitivity analysis of emissions corridors for the 21st century. *Climatic Change*, 66(3), pp. 345–387.

Kriegler, E., Riahi, K., Bauer, N., Schwanitz, V.J., Petermann, N., Bosetti, V., Marcucci, A., Otto, S., Paroussos, L., Rao, S., CurrÃ¡s, T.A., Ashina, S., Bollen, J., Eom, J., Hamdi-Cherif, M., Longden, T., Kitous, A., MÃ©jean, A., Sano, F., Schaeffer, M., Wada, K., Capros, P., van Vuuren, D.P. & Edenhofer, O. (2014). Making or breaking climate targets: The {AMPERE} study on staged accession scenarios for climate policy. *Technological Forecasting and Social Change*, (0), pp. –.

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

Lavalle, I. (1986). On information-augmented chance-constrained programs. *Operations research letters*, 4(5), pp. 225–230.

Lempert, R.J., Groves, D.G., Popper, S.W. & Bankes, S.C. (2006). A general, analytic method for generating robust strategies and narrative scenarios. *Management Science*, 52(4), pp. pp. 514–528.

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

Lontzek, T.S. & Narita, D. (2011). Risk-averse mitigation decisions in an unpredictable climate system. *The Scandinavian Journal of Economics*, 113(4), pp. 937–958.

Lorenz, A., Kriegler, E., Held, H. & Schmidt, M.G.W. (2012a). How to measure the importance of climate risk for determining optimal global abatement policies? *Clim Change Econ*, 3.

Lorenz, A., Held, H., Bauer, E. & Deimling, T.S. (2009). Constraining ocean diffusivity from the 8.2Â ka event. *Climate Dynamics*, 34(5), pp. 719–734.

Lorenz, A., Schmidt, M.G.W., Kriegler, E. & Held, H. (2012b). Anticipating climate threshold damages. *Environmental Modeling and Assessment*, 17(1-2), pp. 163–175.

Luderer, G., Bosetti, V., Jakob, M., Leimbach, M., Steckel, J.C., Waisman, H. & Edenhofer, O. (2011). The economics of decarbonizing the energy system – results and insights from the RECIPE model intercomparison. *Climatic Change*, 114(1), pp. 9–37.

Luderer, G., DeCian, E., Hourcade, J.C., Leimbach, M., Waisman, H. & Edenhofer, O. (2012). On the regional distribution of mitigation costs in a global cap-and-trade regime. *Climatic Change*, 114(1), pp. 59–78.

Luderer, G., Leimbach, M., Bauer, N., Kriegler, E., Tino, A., Baumstark, L., Bertram, C., Giannousakis, A., Hilaire, J., Klein, D. *et al.* (2013). Description of the remind model (version 1.5).

Manne, A. (2005). MERGE: an integrated assessment model for global climate change. *Energy and Environment*.

Mastrandrea, M., Field, C., Stocker, T., Edenhofer, O., Ebi, K., Frame, D., Held, H., Kriegler, E., Mach, K., Matschoss, P., Plattner, G.K., Yohe, G. & Zwiers, F. (2010). Guidance Note for Lead Authors of the IPCC Fifth Assessment Report on Consistent Treatment of Uncertainties. *Intergovernmental Panel on Climate Change (IPCC).*

Mastrandrea, M.D. & Schneider, S.H. (2004). Probabilistic integrated assessment of "dangerous" climate change. *Science*, 304(5670), pp. 571–5.

Meinshausen, M., Hare, B., Wigley, T.M.L., Van Vuuren, D.P., Den Elzen, M.G.J. & Swart, R. (2006). Multi-gas emissions pathways to meet climate targets. *Climatic Change*, 75(1), pp. 151–194.

Meinshausen, M., Meinshausen, N., Hare, W., Raper, S.C.B., Frieler, K., Knutti, R., Frame, D.J. & Allen, M.R. (2009). Greenhouse-gas emission targets for limiting global warming to 2°C. *Nature*, 458(7242), pp. 1158–1162.

Messner, S. & Strubegger, M. (1995). User's Guide for MESSAGE III. *Rep. WP-95-69, International Institute for Applied Systems Analysis, Laxenburg, Austria*, 14.

Mirrlees, J. & Stern, N. (1972). Fairly Good Plans. *Journal of Economic Theory*, 4(2), pp. 268–288.

Morgan, M.G., Henrion, M. & Small, M. (1992). *Uncertainty.* Cambridge University Press.

Nordhaus, W.D. (2008). A question of balance: economic modeling of global warming. *Yale University, New Haven, Connecticut.*

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

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

O'Neill, B.C., Crutzen, P., GrÃŒbler, A., Duong, M.H., Keller, K., Kolstad, C., Koomey, J., Lange, A., Obersteiner, M., Oppenheimer, M., Pepper, W., Sanderson, W., Schlesinger, M., Treich, N., Ulph, A., Webster, M. & Wilson, C. (2006). Learning and climate change. *Climate Policy*, 6(5), pp. 585–589.

Oppenheimer, M. & Petsonk, A. (2005). Article 2 of the unfccc: Historical origins, recent interpretations. *Climatic Change*, 73(3), pp. 195–226.

Ott, K., Klepper, G., Lingner, S., Schäfer, A., Scheffran, J., Sprinz, D. & Schröder, M. (2004). Reasoning goals of climate protection. specification of article 2 unfccc. *Report for the Federal Environmental Agency, Berlin. Europäische Akademie GmbH, Berlin.*

Otto, A., Todd, B.J., Bowerman, N., Frame, D.J. & Allen, M.R. (2013). Climate system properties determining the social cost of carbon. *Environmental Research Letters*, 8(2), pp. 24–32.

Patt, A.G. (1999). Extreme outcomes: the strategic treatment of low probability events in scientific assessments. *Risk Decision and Policy*, 4(1), pp. 1–15.

Pearlman, J., Craglia, M., Bertrand, F., Nativi, S., Gaigalas, G., Dubois, G., Niemeyer, S. & Fritz, S. (2011). Eurogeoss: an interdisciplinary approach to research and applications for forestry, biodiversity and drought. In: *Proceedings of the 34th international symposium on remote sensing of environment*, pp. 1–4.

Pindyck, R.S. (2013). Climate Change Policy: What Do the Models Tell Us? *Journal of Economic Literature*, 51(3), pp. 860–872.

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

Rogelj, J., Hare, W., Lowe, J., van Vuuren, D.P., Riahi, K., Matthews, B., Hanaoka, T., Jiang, K. & Meinshausen, M. (2011). Emission pathways consistent with a 2Â°C global temperature limit. *Nature Climate Change*, 1(11), pp. 413–418.

Rogelj, J., Meinshausen, M. & Knutti, R. (2012). Global warming under old and new scenarios using IPCC climate sensitivity range estimates. *Nature Climate Change*, 2(4), pp. 248–253.

Rotmans, J. & van Asselt, M. (2001). Uncertainty in integrated assessment modelling: A labyrinthic path. *Integrated Assessment*, 2(2), pp. 43–55.

Saliby, E. (1990). Descriptive sampling: A better approach to monte carlo simulation. *The Journal of the Operational Research Society*, 41(12), pp. 1133–1142.

Savage, L.J. (1951). The theory of statistical decision. *Journal of The American Statistical Association*, 46(253), pp. 55–67.

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

Schneider, S.H. & Mastrandrea, M.D. (2005). Probabilistic assessment of "dangerous" climate change and emissions pathways. *Proceedings of the National Academy of Sciences of the United States of America*, 102(44), pp. 15728–35.

Schneider von Deimling, T., Held, H., Ganopolski, A. & Rahmstorf, S. (2006). Climate sensitivity estimated from ensemble simulations of glacial climate. *Climate Dynamics*, 27(2-3), pp. 149–163.

Stern, N. (2007). *The Economics of Climate Change*. HM Treasury London, ISBN 9780521700801.

Stevens, B. & Bony, S. (2013). Water in the atmosphere. *Physics Today*, 66(6), p. 29.

The World Bank (2013). 4.8 world development indicators: Structure of demand.

Tol, R.S.J. (1997). On the optimal control of carbon dioxide emissions: an application of FUND. *Environmental Modeling and Assessment*, 2(3), pp. 151–163.

Traeger, C. (2014). Why uncertainty matters: discounting under intertemporal risk aversion and ambiguity. *Economic Theory*, pp. 1–38.

UNFCCC (2010). Report of the Conference of the Parties on its sixteenth session, held in Cancun from 29 November to 10 December 2010. Addendum. Part Two: Action taken by the Conference of the Parties at its sixteenth session. Decision 1/CP.16:The Cancun Agreements: Outcome of the work of the Ad Hoc Working Group on Long-term Cooperative Action under the Convention FCCC/CP/2010/7/Add.1.

UNFCCC (2011). Report of the Conference of the Parties on its seventeenth session, held in Durban from 28 November to 11 December 2011. Addendum. Part Two: Action taken by the Conference of the Parties at its seventeenth session. Decision 1/CP.17: Establishment of an Ad Hoc Working Group on the Durban Platform for Enhanced Action. FCCC/CP/2011/9/Add.1.

van Vuuren, D.P., Lowe, J., Stehfest, E., Gohar, L., Hof, A., Hope, C., Warren, R., Meinshausen, M. & Plattner, G.K. (2011). How well do integrated assessment models simulate climate change? *Climatic Change*, 104(2), pp. 255–285.

WBGU (1997). *Targets for Climate Protection, 1997*. Technical report, German Advisory Council on Global Change.

Weaver, C.P., Lempert, R.J., Brown, C., Hall, J.A., Revell, D. & Sarewitz, D. (2013). Improving the contribution of climate model information to decision making: the value and demands of robust decision frameworks. *Wiley Interdisciplinary Reviews: Climate Change*, 4(1), pp. 39–60.

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

Webster, M. (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*, (67).

Weitzman, M.L. (2007a). A Review of the Stern Review on the Economics of Climate Change. *Journal of Economic Literature*, XLV(September), pp. 703–724.

Weitzman, M.L. (2007b). Subjective expectations and asset-return puzzles. *The American Economic Review*, 97(4), pp. 1102–1130.

Weitzman, M.L. (2009). Additive Damages , Fat-Tailed Climate Dynamics , and Uncertain Discounting. *Economics: The Open-Access, Open-Assessment E-Journal*, 3, pp. 0–22.

Wigley, T.M. & Raper, S.C. (2001). Interpretation of high projections for global-mean warming. *Science*, 293(5529), pp. 451–4.

Acknowledgments

First of all, I would like to express my sincere gratitude towards my supervisor, or as you say in German: “PhD-Father”, Prof. Dr. Hermann Held, for all the long hours of supporting, interesting and inspiring conversation over the years. My gratitude extends also to subjects unrelated to the PhD but specifically to the helpful feedback and ideas about the direction in which to keep going. Thank you very much!

I am very grateful for the openness of Uwe Schneider to discuss any GAMS and conceptual problem when programming, especially his open door policy was very appreciated. I also thank Steven Dirkse for support on the GAMS list and for help with programming the calibration technique. I thank all attendants, lecturers and organizers of the Summer School FEEM/EAERE/VIU of 2013, which was a great experience and I can only recommend the event.

My wife, Theresa, has been very patient with me, especially in the last months and I thank her dearly for ongoing support throughout the endeavor. My thanks also extends to my parents who always encouraged me and enabled me to have a shot at university in the first place.

My colleagues, especially Jana Müller, Natalie Trapp, Felix Bunzel, Marius Stankoweit, Benjamin Blanz and Robert Roth should not be left unmentioned as they were the ones that never got tired of discussing topics related to my work and sat through many practice presentation without actually working on the topic themselves. They were always there to provide a welcome distraction when needed. Special thanks goes out to Benjamin and Robert who offered to proof read the thesis and I thank them for helpful comments.

The IMPRS-ESM was a great help to smooth the many obstacles on the PhD road and therefore I thank Bjorn Stevens as my panel chair and Antje Weitz as coordinator of the school for always being there for questions and support. I am grateful for the stipend that was made possible by the IMPRS-ESM and allowed me to fully concentrate on the thesis.

Work would have been much harder if it were not for our central IT services, ZMAW library and Joeys. Thank you all!

Die gesamten Veröffentlichungen in der Publikationsreihe des MPI-M
„Berichte zur Erdsystemforschung“,
„Reports on Earth System Science“,
ISSN 1614-1199

sind über die Internetseiten des Max-Planck-Instituts für Meteorologie erhältlich:

http://www.mpimet.mpg.de/wissenschaft/publikationen.html

ISSN 1614-1199