Efficient Climate Policies under Technology and Climate Uncertainty

Hermann Held[*,†], Elmar Kriegler[*,‡], Kai Lessmann[*] and Ottmar Edenhofer[*,§]

1st December 2008

Abstract

This article explores efficient climate policies in terms of investment streams into fossil and renewable energy technologies. The investment decisions maximise social welfare while observing a probabilistic guardrail for global mean temperature rise under uncertain technology and climate parameters. Such a guardrail constitutes a chance constraint, and the resulting optimisation problem is an instance of chance constrained programming, not stochastic programming as often employed. Our analysis of a model of economic growth and endogenous technological change, MIND, suggests that stringent mitigation strategies cannot guarantee a very high probability of limiting warming to 2°C since preindustrial time under current uncertainty about climate sensitivity and climate response time scale. Achieving the 2°C temperature target with a probability P^* of 75% requires drastic carbon dioxide emission cuts. This holds true even though we have assumed an aggressive mitigation policy on other greenhouse gases from, e.g., the agricultural sector. The emission cuts are deeper than estimated from a deterministic calculation with climate sensitivity fixed at the P^* quantile of its marginal probability distribution (3.6°C). We show that earlier and cumulatively larger investments into the renewable sector are triggered by including uncertainty in the technology and climate response time scale parameters. This comes at an additional GWP loss of 0.3%, resulting in a total loss of 0.8% GWP for observing the chance constraint.

We obtained those results with a new numerical scheme to implement constrained welfare optimisation under uncertainty as a chance constrained programming problem in standard optimisation software such as GAMS. The scheme is able to incorporate multivariate non-factorial probability measures such as given by the joint distribution of climate sensitivity and response time. We demonstrate the scheme for the case of a four-dimensional parameter space capturing uncertainty about climate and technology.

[*] Potsdam Institute for Climate Impact Research, PO Box 60 12 03, D-14412 Potsdam, Germany

[†] held@pik-potsdam.de

[‡] Department of Engineering and Public Policy, Carnegie Mellon University, Pittsburgh, PA 15213, USA

[§] Technical University of Berlin, Straße des 17. Juni 135, 10623 Berlin, Germany

1 Introduction

The analysis presented in this article investigates efficient climate change mitigation options under uncertainty about energy technologies and the climate system. We follow the dominant decision-analytic approach in assuming that uncertainty can be represented by probability distributions. We employ a model of economic growth and endogenous technological change, MIND (Edenhofer et al., 2005), to identify intertemporally optimal investment streams into competing energy technologies under such uncertainty. In this model, present-day fossil fuel based energy technologies emitting greenhouse gases compete with carbon-free energy generation from fossil fuels using carbon capturing and sequestration (CCS), with renewable energy technologies, and with investments in energy efficiency. MIND includes a simple climate module that maps greenhouse gas emissions onto global mean temperature rise. This allows MIND to derive the welfare maximising investment path under the boundary condition that global mean temperature rise (as against pre-industrial values) shall not transgress 2°C[1]. Following Gerlagh and van der Zwaan (2004) we call the so derived model futures *2DC* scenarios.

The discussion of constrained welfare optimisations has often been accompanied by information about the cost-effectiveness of long-term climate change mitigation policies (e.g. Manne and Richels, 1997). In this context, the costs of efficient (in the sense of welfare maximising) climate policies are evaluated by comparing production and/or consumption streams from the constrained welfare maximisation with an unconstrained *business as usual* (BAU) case, which reflects the optimal behaviour of the economy in the absence of climate change. As a cost metric for the policy vs. BAU case we explore the traditionally used net present value (NPV) of production. We will show (in Fig. 3 displayed below) that in our application relative NPV losses roughly equal relative losses in balanced growth equivalent (a cost metric proposed by Mirrlees and Stern 1972 and employed in Stern (2007)) up to a factor of 1.5.

Cost-effectiveness and welfare analyses of climate policies have been performed for more than a decade, but gained renewed attention recently. The inclusion of endogenous technological change led to the derivation of unprecedented low mitigation costs for reaching ambitious climate protection targets at or below an atmospheric CO_2 concentration of 450 ppm (approximately 0.5% net present value GWP loss 2000-2100 with GWP discounted at 5% per year). Costs of such magnitude were particularly affirmed in the *Innovation Modelling Comparison Project* (IMCP) which included ten economic models (Edenhofer et al., 2006).

However the basic results of IMCP are based on average values of key uncertain economic and climate system parameters. Here we ask how these results would change if the investment portfolio was optimised under uncertainty about key parameters in the economic and climate system (for an overview see Kann and Weyant, 2000; Peterson, 2006).

A series of economic studies have undertaken sensitivity studies with respect to such parameters. Nordhaus (1994) performed extensive sensitivity analyses in his seminal work on unconstrained welfare optimisation with the DICE model (weighing costs and avoided damages of mitigation policies Nordhaus,

[1] Employing constraints in an economic optimisation may be interpreted as a high degree of discontinuous risk-aversion.

1994; Nordhaus and Boyer, 2000). Gerlagh and van der Zwaan (2004); van der Zwaan and Gerlagh (2006) demonstrate that general findings from deterministic climate policy optimisations under a temperature constraint are not robust against variations in key uncertain parameters (see also Bürgenmeier et al., 2006). Edenhofer et al. (2006) investigate the sensitivity of optimal portfolios of mitigation options for achieving greenhouse gas concentration targets, and find that changing fossil resource base or learning parameters by a few percent, results in relative changes in the optimal energy mix of the same order of magnitude. (Below we will argue that those parameters are in fact even more uncertain.)

Mastrandrea and Schneider (2004) provide an example of how to transcend mere sensitivity study by propagating probabilistic information on uncertain parameters (climate sensitivity and damages) onto model output (dangerous anthropogenic interference). Yohe et al. (2006) investigate the outcome of (carbon tax) policies on the probability of disrupting the Atlantic thermohaline circulation. Hereby the disruption probability is calculated for prescribed carbon policies based on probability distributions of four climate system parameters. Thus Yohe et al. present a sensitivity analysis over policies.

Once a set of influential uncertain parameters is identified by some sort of sensitivity analysis, it would be of highest importance for decision advice to obtain control paths that are optimal *under* the given uncertainty. In contrast to sensitivity analysis which investigates the dependence of optimal policies on the variation of fixed parameter values, or alternatively the dependence of outcome probabilities on the variation of fixed policies, decision analysis includes the uncertainty about policy outcomes in the objective function. This is the focus of the analysis presented here. Decision criteria under uncertainty, such as the expected utility criterion, have been intensively explored by decision theory. Decision frameworks usually extend to the important case of learning about the uncertain state of the world, and subsequently of adapting the optimal policy to improved information at latter points in time.

Contributions to be found in the literature may be classified in terms of constrained or unconstrained optimisation, how uncertainty is absorbed in the optimisation functional, whether learning is anticipated, and whether the analysis is along a highly simplified discretised or, on the contrary, quasi-continuous control, parameter or output space. For highly stylised conceptual investigations which (other than this article) consider optimising expected utility under learning see e.g. Nordhaus (1994) and Valverde et al. (1999).

Richels et al. (2004) provide the probability P^* of observing the 2°C temperature target as function of a carbon tax. They then invert this information in order to obtain the tax that would induce a desirable P^* (Table 7), given uncertainty in climate sensitivity, climate response time scale and GWP growth rate (with $3 \times 3 \times 5$ *states of the world*). Richels et al. can utilise their sensitivity programming to assess decision making under uncertainty, but the setup crucially depends on P^* being a monotonous function of the control parameter "tax".)

Yohe et al. (2004) analyse optimisation under an uncertain temperature target, in combination with uncertain climate sensitivity, the value of both to be revealed under a single *learning step* in the year 2035. They optimise over a binary control space (tax or no tax) and obtain that a modest initial tax is the dominant strategy. According to their result, it would be rational to mitigate now instead of waiting until uncertainty has been reduced. This is in line with the findings of Nordhaus (1994) when including uncertainty in the optimisation. Quite to the contrary, Keller et al. (2004) find decreasing near-term abatements, when considering the effects of uncertainty stemming from highly nonlinear damage functions. (Hereby the authors optimise over a quasi-continuous intertemporal control space, involving a sophisticated non-local optimiser.) This underlines that including new dynamical effects may result in unexpected implications of uncertainty.

In contrast to unconstrained optimisation where costs and benefits are fully reflected in the objective function, Syri et al. (2007) optimise expected utility under a constraint on temperature increase which safeguards the climate. Within their approaches that assume just a few possible states of the world (involving among other uncertain parameters climate sensitivity), they manage to obtain emissions paths that observe the temperature target in *each* state of the world. Ambrosi et al. (2003) solve the same functional for an intertemporal abatement control, using the intertemporal optimisation software GAMS (General Algebraic Modeling System; Brooke et al., 1992). They show that inclusion of uncertainty results in earlier abatement.

However, expected utility optimisation subject to a constraint generally leads to a chance-constrained programming problem (Charnes and Cooper, 1959). If full compliance with the constraint cannot be guaranteed under uncertainty (for any policy), framing the decision problem as stochastic programming (i.e., optimisation of expected utility while observing the constraint for *all* sampled combinations of parameters) will lead to conceptual problems, since the feasibility of a solution becomes a function of the uncertainty sampling. In particular, full (numerical) compliance with the constraint may be possible only under severe undersampling of the probability measure. A fat tail of a probability density distribution will generally prohibit the existence of a policy that can observe a temperature guardrail with (almost) certainty. For that reason, by employing a stochastic programming approach, the above mentioned contributions could not address the fat-tail of the probability distributions assembled for climate sensitivity IPCC (2007b) adequately.

In order to address this fact, one needs to employ *chance-constrained programming*: optimisation while prescribing a minimum probability P^* of observing a certain (temperature) constraint. McInerney and Keller (2008) operationalise this by adding a penalty to the optimised expected utility ("reliability constraints" in their terminology). They find that reducing the odds of a North Atlantic meridional overturning circulation collapse to 1/10, would require an almost complete shutdown of emissions within a few decades.

In this article we present a *different* numerical scheme for optimising nonlinear welfare under a chance constraint that does not rely on penalty functions. The scheme employs the optimisation software GAMS and is suitable for complex models of economic growth that would not allow for an analytic treatment (in MIND dozens of paths of prognostic variables are coupled in a nonlinear way). We constrain global mean temperature rather than greenhouse gas concentration as temperature is closer to the politically relevant impacts of global warming, hereby following the arguments by Commission of European Communities (2007); Gerlagh and van der Zwaan (2004); Richels et al. (2004).

While we acknowledge that learning and adaptive strategies are of great importance for climate mitigation policies, we do not include this case in our analysis here. This is owed to the complexity of the analysis which involves – in addition to MIND's complexity mentioned above – a four-dimensional uncertainty space, and a high dimensional policy space. Given this overall complexity, we begin with the natural first step to investigate efficient one-shot policies, and defer the incorporation of learning to future research.

This article is organised as follows. Section 2 introduces our decision criteria for optimisation under uncertainty, including our definition of a probabilistic guardrail. We also summarise the model MIND and elaborate on those key parameters that are considered uncertain in this study. In Section 3 we define our strategy of sampling from probability distributions on the space of uncertain parameters and our numerical implementation of optimisation under uncertainty within GAMS. In Section 4 we operate the deterministic version of MIND in sensitivity mode to outline the relative influence of the uncertain parameters. In our main Section 5 we present results from probabilistic optimisation. Finally in Section 6 we summarise the main findings of this article.

2 Methodology and model

For our analysis, we are using MIND, the **M**odel of **In**vestment and Technological **D**evelopment. The original model version 1.0 was presented in Edenhofer et al. (2005). Later on, an endogenous carbon capturing and sequestration (CCS) module was added (version 1.1, Bauer (2005)), and a more elaborate carbon cycle and atmospheric chemistry module included (version 1.2IMCP Edenhofer et al., 2006). In this paper, we use MIND in its previous version 1.1 with only a few changes[2].

In particular, we use the carbon cycle-climate model employed in Edenhofer et al. (2005) (see Appendix C). This carbon-cycle climate module (introduced in Petschel-Held et al. (1999) and described in detail in Kriegler and Bruckner (2004)) uses carbon dioxide and sulphur dioxide emissions as inputs, and converts them into concentrations and subsequent radiative forcing of the earth system. The sum total of the forcing is calculated by adding an exogenous scenario for the aggregate radiative forcing of other greenhouse gases (OGHGs) and aerosols (see Fig. 1). In MIND version 1.1, this scenario has been assumed to follow the SRES B2 scenario (IPCC, 2001). Since constraining global mean temperature change to 2°C warming since preindustrial time is extremely difficult with CO_2 abatement alone (Hare and Meinshausen, 2006; IPCC, 2007a), we have assumed an OGHGs forcing scenario modelled after a proposal of Hansen and Sato (2001) incorporating aggressive mitigation measures. This scenario follows the lowest available SRES scenarios (B1 for halocarbons, A1T for nitrous oxide), and goes even further for methane and tropospheric ozone, whose forcings are assumed to be reduced by 60% until 2100. After 2100, the declining trend of most radiative agents is linearly extrapolated. The resulting aggregated forcing for the other greenhouse gases and aerosols (excluding CO_2, sulphate and carbonaceous aerosols from fossil fuel burning) is shown in Fig. 1.

[2] The population scenario was updated from SRES B2 Nakićenović and Swart (2000) to the common POLES/IMAGE baseline (CPI, van Vuuren et al. (2003)). Instead of the SRES B2 scenario for CO_2 emissions from land-use change, we use the SRES A1T scenario. In Edenhofer et al. (2005), the 2°C climate policy also limited temperature increase per decade to 0.2 K, but this guardrail is not included here.

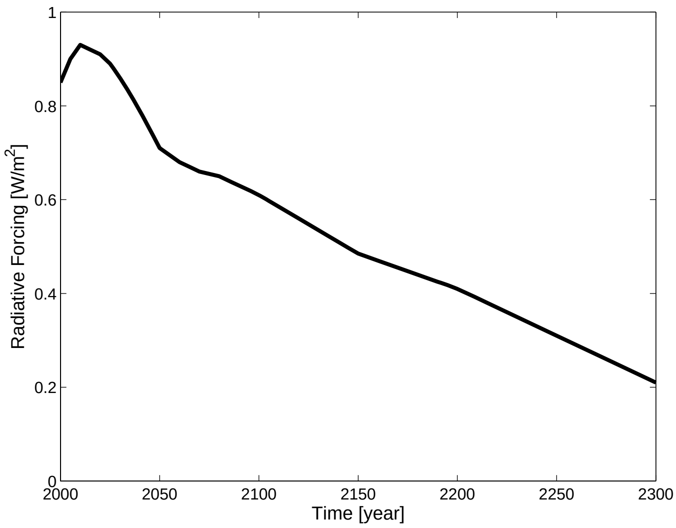

Figure 1: Total radiative forcing from other greenhouse gases and aerosols (excluding CO_2, sulphate and carbonaceous aerosols from fossil fuel burning) reflecting an aggressive exogenous mitigation policy (Hansen and Sato, 2001) for those radiative agents.

Similar to Nordhaus' DICE model, MIND is based on an optimal growth framework to describe the macroeconomy (Nordhaus, 1994; Nordhaus and Boyer, 2000). However, it differs from DICE in mainly two respects. Firstly, technological change is endogenous to MIND. This is important to us, because technological change induced by environmental policy may substantially differ from business-as-usual projections, especially over the course of a century or more. Secondly, the energy sector of the economy includes the technological detail reminiscent of energy system models. This enables us to discuss mitigation options such as energy from renewable sources and carbon capture and storage explicitly. In this sense, MIND is a hybrid energy system-macroeconomic growth model.

The following Subsections introduce those portions of MIND in greater detail that are relevant to the paper. This includes the model equations defining the parameters under investigation, i.e., the development of renewable energy technologies, in particular their expected learning by doing effects, the availability of fossil fuels, and the temperature response to anthropogenic forcing. It also includes the specific setup of the welfare maximisation in MIND, both in its deterministic version, and its extension to the probabilistic optimisation problem investigated here. More complete discussions of the model equations may be found in Edenhofer et al. (2005); Bauer (2005); Edenhofer et al. (2006).

2.1 Deterministic welfare analysis

The objective of economic activity in MIND is maximisation of social welfare. As a proxy for welfare, we use the present value of utility which itself is given by the discounted logarithm of per capita consumption under an exogenously given population scenario $L(\cdot)$. Compared to the DICE model where economic output is distributed over consumption, a generic emissions control, and investment in production capital (Nordhaus, 1994), MIND provides greater flexibility in allocating output Y. It can be consumed (C) or invested (I) in (i) industrial capital, (ii) R&D sectors generating labour and energy augmenting technological change, (iii) generic fossil and renewable energy technologies, and (iv) carbon capture and sequestration. Hence, the control variables of MIND are constituted by the investment paths $I(\cdot)$ in various economic and energy sectors (to be discussed in greater detail below), with the residual output allocated to consumption according to the budget constraint

$$C(t) = Y(t) - \sum_i I_i(t)\,, \qquad C(t) \geq 0\,. \tag{1}$$

The investment paths are chosen according to

$$\max_{I(\cdot)} \qquad \int_{1995}^{2300} \mathrm{d}t\ L(t)\ \ln\left(\frac{C(t;I(\cdot))}{L(t)}\right)\ \mathrm{e}^{-\rho t} \tag{2}$$

$$\text{subject to} \qquad \forall_{t\in[1995,2400]} \quad T(t;I(\cdot)) \leq 2°\mathrm{C}.$$

Hereby, ρ is the pure rate of social time preference (PRTP), and set to 1% per year in this study. The temperature constraint, which we may refer to as "guardrail" in this study, signifies the climate policy requirement to be reached in an efficient manner. The time horizon of the guardrail exceeds the horizon of the economy to account for inertia in the climate system. For these years we assume continuous resource extraction and land use change to the extent of the year 2300 and continuous leakage from CCS sites to be sources of CO_2 emissions. In the BAU case, the objective function is maximised without such constraint. Obviously, nonlinear optimisation problem (2) is deterministic in nature, since utility and temperature are assumed to be known with certainty for given investment decisions $I(\cdot)$.

2.2 Extension to probabilistic welfare analysis

In case several model parameters – summarised by vector a – are regarded as uncertain, the outcome of a climate policy in terms of temperature and intertemporal utility will be a function of a. Therefore, above optimisation problem needs to be generalised. For the purpose of this study, we (i) assume that uncertainty on $a \in A$ can be modelled by means of a probability distribution $F : A \to [0, 1]$, and (ii) choose to maximise the expected utility of an investment decision $I(\cdot)$. Although this reflects the conventional approach to decision making under uncertainty, we note that both assumptions (of probability and of maximising the expectation) may have to be dropped in sophisticated applications which seek robust policies under deep uncertainty about the state of the world.[3] Alternative uncertainty representations and decision criteria have been explored, e.g.,

[3] In fact Knight (1921) highlighted as early as almost a century ago that decisions under "risk" (based on known probability distributions) should be distinguished from other sorts of decision situations involving lack of information or incomplete knowledge (so called deep or "Knightian" uncertainty). In this article, we disregard "Knightian" uncertainty in the sense that we assume that the available information on uncertain parameters is sufficient to justify the use of probability distributions. On the other hand one may argue that we address part of

by Lempert (2002) (and references therein) using an evaluation of regret, and Kriegler et al. (2007) employing imprecise probability theory. However, in line with our incremental approach to the analysis of efficient climate policies under uncertainty, we make the natural first step (in the framework of complex economic models) in this paper, and will investigate the consequence of relaxing the conventional assumptions for our application of MIND in future work. The key step for generalisation, however, will be the inclusion of sequential learning under uncertainty.[4]

In order to implement expected utility maximisation with a chance constraint, we also need to generalise the temperature "guardrail" imposed in optimisation problem 2. The future temperature trajectory $T(\cdot; I(\cdot), a)$ now depends on a as well, and therefore it may be impossible to exclude the violation of the guardrail with certainty. In Kleinen (2005), it was suggested to generalise the temperature guardrail by including a lower limit for the probability P of observing it, i.e.,

$$\forall_t \; \{P(T(t; I(\cdot), a) \leq 2°\mathrm{C}) \geq P^*\} \\ \Leftrightarrow \quad \forall_t \; \{P(T(t; I(\cdot), a) > 2°\mathrm{C}) \leq 1 - P^*\} \tag{3}$$

This implies a time-point-wise observation of the guardrail with a certain probability P^* that needs to be decided on in addition to the temperature value of the guardrail. However, since this condition ignores the path dependency of T on a we assert that for most global warming impacts it would be desirable to observe the stricter condition

$$P\left(\forall_t \; \{T(t; I(\cdot), a) \leq 2°\mathrm{C}\}\right) \geq P^* \\ \Leftrightarrow \quad P\left(\exists_t \; \{T(t; I(\cdot), a) > 2°\mathrm{C}\}\right) \leq 1 - P^* \tag{4}$$

Condition (4) implies Condition (3). The stricter condition recognises that for many impacted systems it matters whether the guardrail was violated at all, but not the time of violation.

Then the probabilistic optimisation reads

$$\max_{I(\cdot)} \quad \int \mathrm{d}F(a) \int_{1995}^{2300} \mathrm{d}t \; L(t) \; \ln\left(\frac{C(t; I(\cdot), a)}{L(t)}\right) \; \mathrm{e}^{-\rho t} \tag{5}$$

$$\text{subject to} \quad P\left(\forall_{t \in [1995, 2400]} \; \{T(t; I(\cdot), a) \leq 2°\mathrm{C}\}\right) \geq P^*. \tag{6}$$

For a given emission path (or a distribution of paths), the left hand side of the inequality constraint can be derived from the probability measure on the parameters of the carbon carbon cycle-climate model, and the model's mapping from emissions onto temperature. The measure that we use is given in

Knightian uncertainty as we are willing to subscribe to the concept of subjective probabilities as representation of epistemic uncertainty (i.e., uncertainty that comes from lack of knowledge rather than real-world stochastic processes).

[4]In the context of exploring alternatives to expected utility within GAMS, we are currently investigating a utility quantile criterion under probabilistic uncertainty that is a promising candidate to provide important insight with respect to "robustness" of policies. The criterion is very flexible as it allows to continuously interpolate between a maximin (a pessimist's) and a maximax (an optimist's) criterion by a quantile parameter Q. See Appendix A for details.

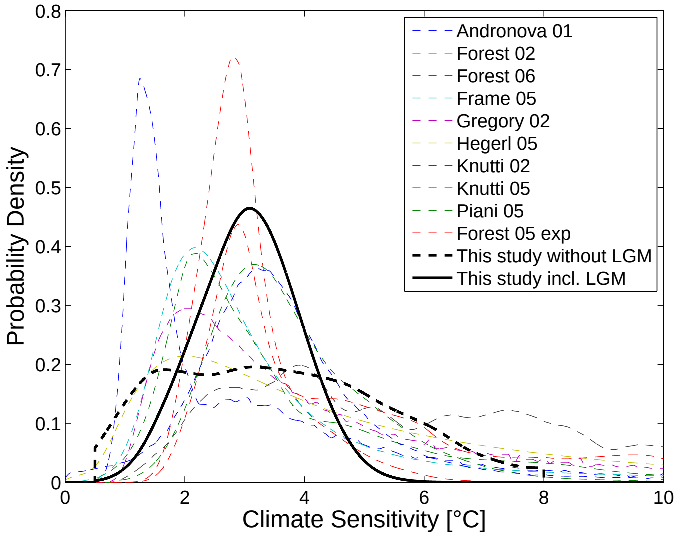

Figure 2: Probability density functions (pdfs) for climate sensitivity (CS) according to the latest IPCC report (thin lines). When marginalising our joint pdf on CS, a rather flat pdf results (bold dashed line). After learning from additional data from the last glacial maximum (Schneider von Deimling et al., 2006), the pdf sharpens significantly (solid line). We assume that usage of such data will become standard in the future and proceed with the solid line in our optimisation.

Appendix D.2. It is one of the few available that capture climate system uncertainty in terms of a joint probability for climate sensitivity (CS) and climate response time scale (dominated by ocean heat uptake). For reasons of comparison with probability estimates in the literature, Fig. 2 displays a 1D- projection of our 2D probability measure w.r.t. climate sensitivity (i.e. the marginal distribution in CS). All probability densities in dashed lines incorporate observational data of the 20th century, IPCC distributions as thin curves (IPCC, 2007b), ours in bold. It can be seen that we consider a rather non-informative distribution in CS.

In Section 5, we will find that using climate sensitivity pdfs incorporating only 20th century observations leads to very low probabilities for achieving the 2° degree target, even for the most stringent emission reductions that remain feasible.

The maximum P^* that can be achieved with such distribution is closer to 1/2 than to 1, yielding a compliance probability that would be difficult to defend as a chance constraint. Only if we absorb additional constraining information from paleo data (Schneider von Deimling et al., 2006) in the climate sensitivity pdf (solid line) we obtain a more confined distribution allowing for $P^* \geq 75\%$. A confined distribution like this was also obtained by Annan and Hargreaves (2006), and a similar interval for plausible values of climate sensitivity was presented by IPCC (2001). For all those reasons we continue with a pdf for climate sensitivity confined by paleo data (black solid line) and the underlying 2D probability measure as working hypothesis for this paper.

2.3 Formulation of the macroeconomy

In MIND, gross world product is produced by a single industrial sector with a technology of constant elasticity of substitution (CES). Inputs to production are industrial capital, labour, and energy. The latter two may be enhanced by research and development (R&D) investments in corresponding knowledge stocks. Substitution possibilities between the three factors are limited by choosing an elasticity of substitution of 0.4. Gross world product is allotted to consumption and investment flows into various sectors in order to maximise the objective function. Sectors comprise the industrial production sector, R&D sectors to enhance labour and energy productivity, respectively, and four stylised sectors describing the energy system (fossil fuel extraction, fossil energy generation, carbon capture and sequestration and renewable energy generation).

The energy input to production is an additive compound differentiating three main energy sources: energy from fossil fuel combustion with or without capture and storage of the resulting emissions, renewable energies, and traditional nonfossil energies (mainly large-scale water power and nuclear power, both via an exogenous scenario).

The energy sector and the carbon cycle-climate dynamics as represented in the MIND version of this article are described in Appendices B and C, respectively. We allow for investments in a fossil sector, renewable sources as well as energy efficiency. The fossil sector emits carbon and sulphur dioxide that both interfere with the radiative balance of the climate system.

3 Numerical implementation of probabilistic optimisation

We consider two uncertain technology parameters, the fossil resource base χ_3 and the economic potential of renewable energy as described by the technology learning rate $\tilde{lr}$ and the technology floor costs c_{floor} (see Appendix B). For reducing the dimensionality of the uncertainty space, we have made the simplifying assumption that learning rate and floor costs of the renewable energy technology are fully correlated. χ_3 assumes values between 2500 and 7000 GtC, $\tilde{lr}$ between 0.08 and 0.24, and c_{floor} between 390 and 630 \$ per kW final energy production capacity (see Appendix D; final energy includes electricity, heat and fuel). In addition we treat climate sensitivity (CS; long-term climate response to carbon dioxide concentration elevation) and the response rate α of the climate system as uncertain (see Appendices D&E and Fig. 2). By sampling these four uncertain parameters within the given ranges, we cover a large fraction of uncertainty relevant for the decision problem of this article.

We assume that the uncertainty about the fossil resource base χ_3, the overall learning rate of renewable technologies determined by the duplet $(c_{\text{floor}}, \tilde{lr})$, and the temperature response parameter duplet (CS,α) are independent of each other, as these parameters emerge in very different domains. The assumption of independence permits us to combine the corresponding samples for χ_3 (indexed by $k \in \{1, ..., K = 10\}$), $(c_{\text{floor}}, \tilde{lr})$ (indexed by $l \in \{1, ..., L = 10\}$), and (CS,α) (indexed by $m \in \{1, ..., M = 301\}$) by full factorial design. Although a suboptimal choice from a sampling point of view, full factorial design allows to restrict the associated sample of model realisations to the parameters that actually affect the individual model variables, which in most cases will be smaller than the size of the overall sample with 30100 ($10 \times 10 \times 301$) members.

The combined sample provides the basis for the numerical evaluation of objective function and constraint in the probabilistic optimisation problem (5)-(6). A numerical implementation requires to replace (i) the integral over time with a sum over discrete time steps (in our case 62 time steps of $\Delta t = 5$ years from 1995 to 2300) and (ii) the expectation operation on the uncertain parameter space with an estimator based on the sum of utilities over a parameter sample. The probability of observing the temperature guardrail will also have to be estimated from the sample of model realizations, using an indicator function for the compliance of the temperature trajectory in each model realization (0: violation of guardrail, 1: compliance). Noting that consumption depends only on the choice of fossil resource base and learning rate of renewable energy technology (i.e., the sum over utilities is indexed only by k,l, not by the climate parameter index m; no climate damages are taken into account in our constrained welfare maximization), and temperature only on (CS,α) and fossil resource base (which affects emissions for a given investment stream in the extraction sector; i.e. the indicator function is indexed only by k,m), we find the following expression as discretised versions of Eqs. (5) and (6).

$$\max_{I(\cdot)} \quad \sum_{k,l} \Delta t \sum_{t=1995}^{2300} L_t \ \ln\left(\frac{C_{klt}(I(\cdot))}{L_t}\right) \ \mathrm{e}^{-\rho(t-1995)} \tag{7}$$

$$\text{subject to} \quad \sum_{k,m} w_m \, B(T_{km}) \geq \left(K \sum_m w_m\right) \ P^* \tag{8}$$

$$\text{with} \quad B(T_{km}) = \begin{cases} 1 & \forall_{t\in[1995,2400]} \ T_{kmt}(I(\cdot)) \leq 2°\mathrm{C} \\ 0 & \exists_{t\in[1995,2400]} \ T_{kmt}(I(\cdot)) > 2°\mathrm{C} \end{cases}$$

the indicator function whether or not a particular temperature trajectory T_{km} has violated the temperature guardrail at some time. If the probability measure on our four parameters underlying the sampling of the parameter space is updated with a likelihood function derived from additional observational data (e.g. paleo data), this does not necessarily require to resample the parameter space. The likelihood information can be included by adding likelihood weights w_m to the sum of indicator functions over our (fixed) sample. In the remainder of the article, w_m will be utilised to re-weight the measure on CS when incorporating new information from analyses of the last glacial maximum (LGM; see Fig. 2, solid vs. dashed lines, Appendix D and Schneider von Deimling et al. (2006)).

Due to the requirement that the temperature trajectory needs to observe the guardrail at all times, the indicator function B depends only on the maximum temperature obtained along a trajectory. Since the use of the maximum operator in optimization problems should be avoided, we introduce an auxiliary variable T^*_{km} (see Eq. 9) that will converge to the maximum temperature of each trajectory in the optimal solution. Finally, to remain in the class of nonlinear continuous programming problems, we smooth the binary indicator function by approximating it with an error function that rapidly changes sign at the point of non-compliance (alternatively an arcustangens function may be used).

$$\forall_{t\in[1995,2400]} \ T^*_{km} \quad \geq \quad T_{kmt} \tag{9}$$

$$B(T_{km}) \quad = \quad 0.5 \ \left(\mathrm{Erf}\left(\frac{2°\mathrm{C} - T^*_{km}}{\epsilon}\right) + 1\right) \tag{10}$$

$$\text{with} \quad \mathrm{Erf}(x) \quad := \quad \int_{-\infty}^{x} \mathrm{d}x' \ \frac{1}{\sqrt{2\pi}} \ \mathrm{e}^{-\frac{1}{2}x'^2} \ , \tag{11}$$

The parameter ϵ controls the non-linearity of the approximation, and has to be chosen carefully to balance accuracy of the indicator function against stability of the optimisation. In this study, we have used a value of $\epsilon := 0.01$.

We have implemented the probabilistic formulation of the model MIND with objective function and chance constraint as expressed in Eq. (7) in the GAMS utilising the nonlinear solver CONOPT3 (Drud, 1992). Of course, the approximation of the continuous probabilistic optimisation problem (5) by sample estimators for expected utility and probability of observing the temperature guardrail (compare Eq. (7)) will give rise to sampling error that may affect the optimal policy. We will have to investigate the influence of sampling error in greater detail in future work. However, previous sensitivity analyses (Bauer

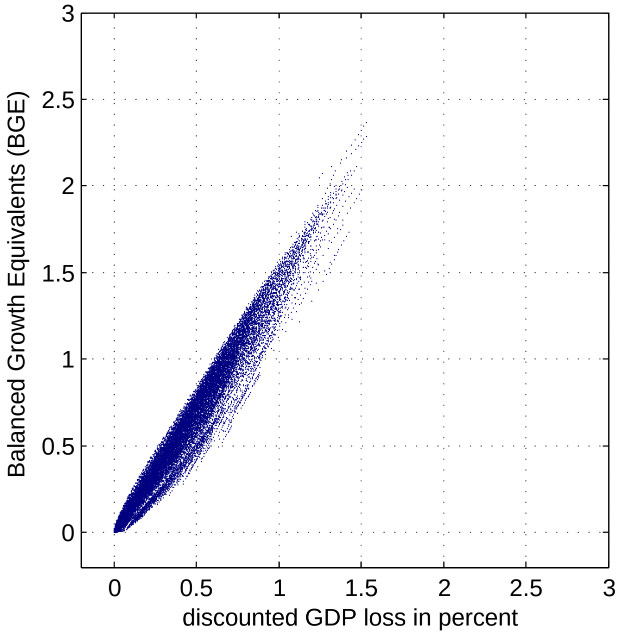

Figure 3: Two measures of economic loss: net present value GWP loss (discounted at 5% per year) and balanced growth equivalents (BGE) for a sensitivity analysis of efficient climate policies over the 30100 member ensemble of climate and economic parameter values.

et al., 2005; Edenhofer et al., 2006) of the model response surface as well as preliminary analyses using a hierarchy of samples indicate that sampling error has only a second order effect on the optimal solution – see also the last paragraph before our conclusions section.

4 Economic impact of climate parameters

Before analysing the results from the expected utility maximisation, we investigate the sensitivity of efficient climate policies and their associated mitigation costs in the deterministic version of MIND1.1 to the parameter uncertainties discussed above.

Sensitivity studies of mitigation costs with respect to climate parameters have been published by Gerlagh and van der Zwaan (2004) who varied the temperature guardrail for fixed climate sensitivity. In their approach, this is approximately equivalent to a varying CS for fixed temperature guardrail, and their range of variation covered an equivalent of CS=2...4°C. They find an increase of the costs of mitigation from 0.06% to 0.29% when moving to stricter temperature guardrails (equivalent to moving to higher climate sensitivities)[5].

[5]Their rather low mitigation costs are due to the fact that their baseline scenario, i.e. BAU, already contains a significant share of renewable source such that imposing a guardrail does

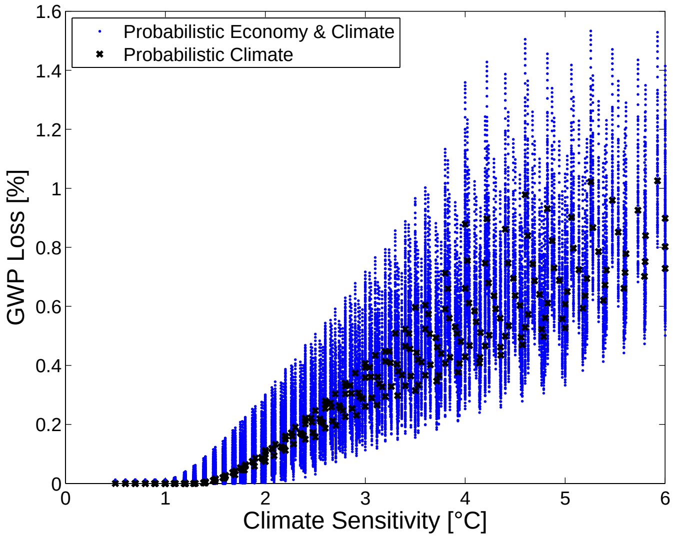

Figure 4: Sensitivity analysis of net present value (NPV) GWP loss of observing the 2DC target (discounted at 5% per year, in percentage of NPV GWP in BAU case) with respect to climate sensitivity for (i) the ensemble of climate response parameter values CS,α with economic parameters fixed at MIND version 1.1 default values ($\chi_3 = 3500$ GtC, $\tilde{lr} = 0.15$, $c_{\text{floor}} = 500\$/\text{kW}$, $n = 301$ – indicated by crosses), (ii) the full ensemble across all climate and economic parameter values in the sample ($n = 30100$ – indicated by grey dots). Note that for a large fraction of parameter settings, solutions are infeasible under the constraint and no costs estimates can be given. This further depicts the conceptual limitations of such a simple sensitivity study under a deterministic constraint, that can only fully be resolved by chance constrained programming.

As outlined in the introduction, we define mitigation costs as net present value (NPV) GWP loss in the climate policy case relative to the business as usual (BAU) case without temperature constraint, and express these costs in percentage of NPV GWP in the BAU case. We assume a constant interest rate of 5 percent to discount GWP losses, because there is no explicit interest rate in MIND. Instead two proxies were checked for consistency with this assumption. The MIND model exhibits an average marginal productivity of capital at 4.6 percent and a shadow price of capital accumulation at 6.5 percent (both for the period 2015 to 2115 where they are approximately constant).

In addition to NPV GWP losses, we have considered a welfare loss criterion based on balanced growth equivalents (BGEs Mirrlees and Stern, 1972). The BGE of a welfare maximising consumption stream is described by the initial value of consumption that when extended into the future with a prescribed constant growth rate leads to the same intertemporally aggregated social welfare. Based on the optimal consumption streams calculated in MIND1.1 we chose a constant growth rate of 2.1% per year to derive the BGEs for the welfare maximising climate policy and BAU solutions. The welfare losses were then calculated – in complete analogy to NPV GWP losses – by taking the difference in BGEs between BAU and policy case, and expressing it as percentage of the BGE in the BAU case. We ask how the two loss measures may differ across the ensemble of welfare maximising GWP and consumption streams under a 2DC temperature guardrail that can be obtained when varying the climate and economic parameter values in MIND1.1 over the 30100 members of our sampled parameter uncertainty. The results are displayed as a scatter plot in Fig. 3. Obviously, the two are strongly correlated, and BGE losses appear approximately 1.5-times larger than NPV losses. Due to the close correlation between the two we will use only one cost metric, i.e. NPV GWP loss, in the remainder of the paper.

Fig. 4 shows the mitigation costs as a function of climate sensitivity across the ensemble of sampled climate and economic parameter values. Our study covers a much larger range of climate sensitivity than considered in Gerlagh and van der Zwaan (2004). The NPV GWP losses for implementing climate policies observing the 2DC guardrail increase significantly with climate sensitivity, reflecting the fact that the more sensitive the climate responds to carbon dioxide forcing, the more binding a 2°C guardrail acts on the energy sector. Concentrating first on the sensitivity with respect to variations in the climate response parameters CS,α for fixed economic parameters (crosses), we find an increase in loss of roughly 1/4 %/°C in CS.[6]

Note that the larger CS, the larger the fraction of parameter combinations for which feasible GAMS-solutions cease to exist, i.e. no matter how fast carbon dioxide emissions were reduced, the 2DC target would still be violated. This is not a numerical effect, but is intuitively to be expected due to the combination of two factors, the (i) additional fixed forcing from other greenhouse gases (OGHGs) and (ii) the warming commitment from the carbon stock already in the atmosphere.

When considering the uncertainty in GWP losses across the entire ensemble of climate and economic parameters (grey dots), it can be seen that climate

not demand as much emission reductions as in other models.

[6]If instead of the aggressive mitigation scenario for OGHGs we use a standard OGHGs forcing following the SRES B2 scenario, then the loss-to-CS-ratio increases to 1/2 %/°C.

sensitivity is a crucial factor in determining the mitigation costs, but that the economic parameters also account for a major fraction of overall uncertainty (populated vertical lines for fixed CS), and that variations in response rate α have a smaller influence (vertical spread in crosses). In contrast to the feasibility limit imposed by the climate component, the economic module responds rather smoothly to CS-tuning due to the flexibility in the energy sector to mitigate carbon dioxide. In this context, it is interesting to note that the model calculates much higher costs (up to 6.5% NPV GWP loss) of climate protection if both CCS and an increase of capacity in the renewable energy sector are switched off (Edenhofer et al., 2005).

Our sensitivity analyses revealed that within the chosen parameter ranges optimal emission paths vary from BAU emissions to almost immediate cessation of emissions. It seems therefore infeasible to identify robust features of optimal controls, even on a conceptual level, from such a sensitivity study. Hence we strongly argue in favour of a full-fledged *optimisation under uncertainty* in which the normative issue of how to deal with such uncertainty is explicitly conceptualised in the formulation of the optimisation problem. In this study, Eqs. (5)-(6) represent our choice for optimisation under uncertainty.

5 Results from probabilistic optimisation

In the following we investigate the climate policy under uncertainty, and its associated set of model futures, that emerges as solution of the numerical approximation (7)-(8) of the optimisation problem (5)-(6) for the sample described in Section 3. To the best of our knowledge it is the first time that a probabilistic optimisation of a coupled climate-macroeconomic growth model including a chance constraint has been performed within a GAMS environment, suitable for models that cannot be solved in analytic form. We note that the optimisation process required more than a week of CPU time (IBM 1.1 GHz Power 4). The resulting curves are less smooth than one would expect (optimality usually implies smoothness). We know from selected experiments that tightening the optimality tolerance improves the smoothness but without qualitative impact on the results. We have thus refrained from the additional computational burden. The results of a model run consists of a single set of investment streams. Due to the modelled economic uncertainty in resource extraction and learning of the renewables the optimal investments give rise to 10×10 economic futures that differ e.g. in resource use, energy mix, and emissions. And taking climate uncertainty into account, each of the emissions futures (10 per optimal investment policy) gives rise to 301 temperature reactions. Out of these 3010 temperature paths, a fraction P^* observes the temperature guardrail. It should also be recalled that the solutions to stochastic optimisation problem (7)-(8) are based on sample estimates of expected utility and the probability to observe the temperature guardrail, and therefore are somewhat contingent on the choice of the parameter sample (see Appendices D&E). We will discuss the possible influence of sampling error at the end of the next Subsection.

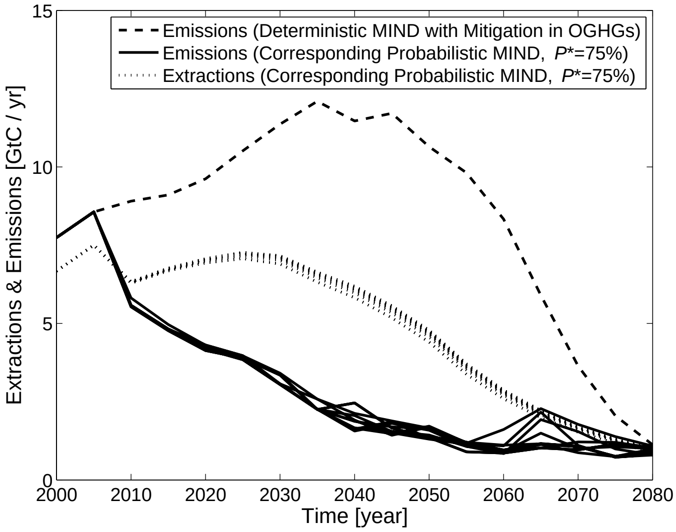

Figure 5: Optimal carbon dioxide emission paths that observe the 2°C guardrail. Solid lines: 10 emissions paths (for the 10 different values of fossil resource base) from the probabilistic optimisation with $P^* = 75\%$. The trajectories include an exogenous scenario for CO_2 land use change emissions (A1T). Dotted lines: Corresponding resource extraction. Dashed line: single emission path from the deterministic application of MIND version 1.1 with default values for economic and climate parameters ($\tilde{lr}$=0.15, χ_3=3500 GtC, CS = 2.8°C, α=0.016855/yr). Emissions need to be reduced to 2 GtC/yr in 2040 while extraction peaks in 2030. This difference is mainly due to carbon capture and sequestration. The mitigation requirements are much more stringent under the chance constraint as compared to the standard deterministic case.

5.1 Optimal emission paths

We cannot expect to find a feasible solution of optimisation problem (7) if the chance constraint of observing the guardrail is set too high. In our model setting, we find a value of $P^* = 80\%$ to be the most stringent probability constraint that still allows to find a feasible solution, i.e., a set of investment streams in the various economic sectors that complies with the probabilistic temperature guardrail. For higher values of P^*, such a solution ceases to exist[7]. For our analysis, we decided to use a compliance probability of $P^* = 75\%$ for the following two reasons: (i) it is intuitively accessible (the odds of compliance are 3 out of 4), and (ii) it is located close to, but not directly at the boundary of the feasibility region where extreme climate policy scenarios will be realized.

Fig. 5 displays the optimal carbon dioxide emissions that observes the 2° guardrail with $P^* = 75\%$, along with the underlying resource extraction. Carbon capture and sequestration allows to increase extraction of fossil resources while at the same time emissions decline. Note that a unique optimal solution in terms of investment streams defines a set of 10 emissions and extraction

[7]This means that, within GAMS, even if the energy sector reduced emissions at the fastest rate possible, the probability of $T > 2°C$ would still transgress P^*.

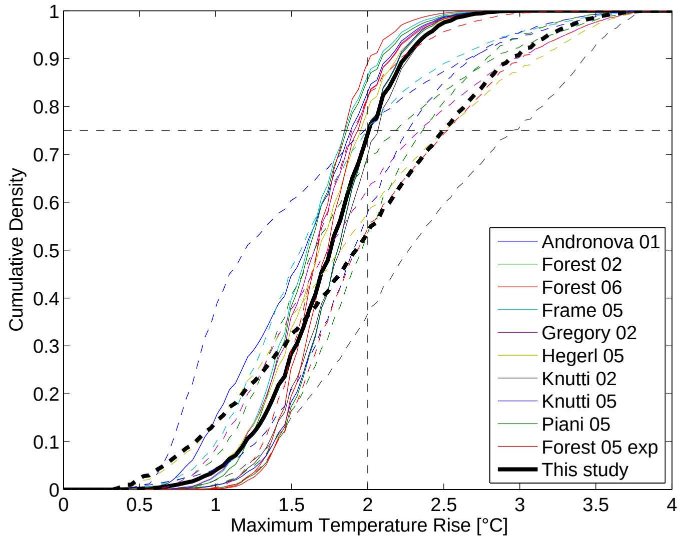

Figure 6: Temperature response driven by the emission paths depicted in Fig. 5. We display the probability that a maximum temperature is not exceeded at any time. By construction our solution does so with $P^* = 75\%$ for a maximum temperature rise of 2°C according to the probability measure that we used in the probabilistic optimisation (solid line, including paleo data (from LGM)). Most probability estimates from the IPCC lie above when updated with the same paleo data (thin solid lines).

paths due to the uncertainty about the efficiency of the investments in the fossil extraction sector which is influenced by the fossil resource base parameter χ_3 (see Appendix D). For comparison, Fig. 5 also shows the CO_2 emissions path derived from the deterministic application of MIND1.1 with default parameter values, including MIND's standard climate sensitivity of 2.8°C. It is apparent that much earlier and drastic emission cuts are necessary in the probabilistic case, reflecting the fact that the specification of the chance constraint is pushing the model to take into account the upper tails of the probability distributions for the uncertain parameters (e.g., the 75% quantile of the marginal pdf for CS is 3.6°C).

In Fig. 6, bold solid line, we display the probability that given those emission paths a prescribed temperature is not exceeded at any time. By construction, according to our joint probability measure for CS and the time scale of the temperature response, the 2°C guardrail is observed with probability $P^* = 75\%$. We ask further which targets would be observed according to the probability distributions reported by IPCC 2007b. As IPCC 2007b does specify only marginal distributions in CS, we pragmatically keep the correlation structure of our joint probability density function (pdf), while adjusting the marginals along CS (for details see Appendix E). With and without paleo data, our probability distributions for temperature is more on the conservative (i.e. lower) side of distributions that can be derived from the IPCC estimates. Without consideration of paleo information (dashed solid lines), the compliance probability can be as low as

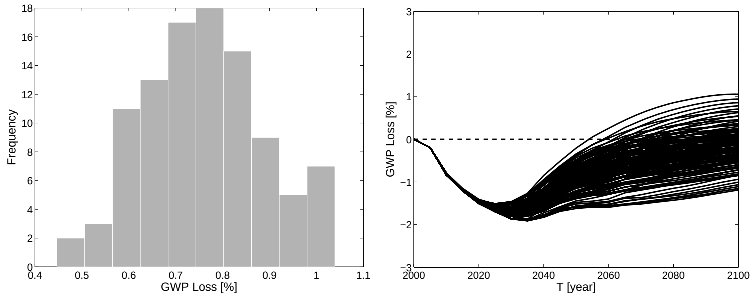

Figure 7: Left graph: Histogram of net present value of GWP losses (discounted by 5%; 2DC policy vs. BAU in units of BAU output) under a 75%-2°C guardrail. Right graph: GWP losses over time. The losses culminate in 2035. For some ensemble members net gains are seen after 2060 for parameter combinations that favour renewable energy technologies and assume scarce fossil resources.

50%, and even for the lowest pdf of CS reported in the IPCC barely reaches up to $P^* = 75\%$.

However, it should be noted that P^* increases steeply at CS=2°C. Hence, a rather small relaxation of the guardrail would result in a large gain in probability of observing the guardrail.

One may ask how robust our results are with respect to sampling error of our subjective probability distributions. When switching to an extremely reduced climate parameter ensemble (23 instead of 301) we find that carbon dioxide emissions paths calculated on the basis of the reduced ensemble oscillate around the bundle derived from the large ensemble with a spread of up to ±1GtC/yr. We suggest to use this number as a very conservative error bar.

5.2 Chance constraint-induced GWP losses

We now analyse the costs of the 2DC climate policy under uncertainty. Since output from the industrial sector directly depends on the fossil resource base and the learning potential of renewable energy technologies (sampled with KL = 100 different parameter constellations), we have to calculate GWP loss for each of those sample members. In contrast, output depends not directly on the temperature response, which affects it only indirectly via the influence of the probabilistic climate guardrail on the choice of optimal policy. The left graph in Fig. 7 displays the distribution of NPV GWP loss. It can be seen that GWP loss lies in the range $0.76\% \pm 0.13\%$ (one standard deviation across the ensemble).

The right graph in Fig. 7 disaggregates the GWP losses along the time axis for each of the 100 ensemble members. The losses peak around 2035 at 2% GWP loss, and then decline towards the end of the century. For a significant fraction of ensemble member, net gains in 2100 are realised.

In order to analyse the cause of the ensemble spread in GWP losses in the year 2100, we generated a reduced ensemble by averaging over learning rate and floor costs in the renewable energy sector. This is possible due to the factorial design we have chosen. We find that the averaging removes half of the ensemble spread, and with it the majority of “net gain paths” in 2100. Accordingly,

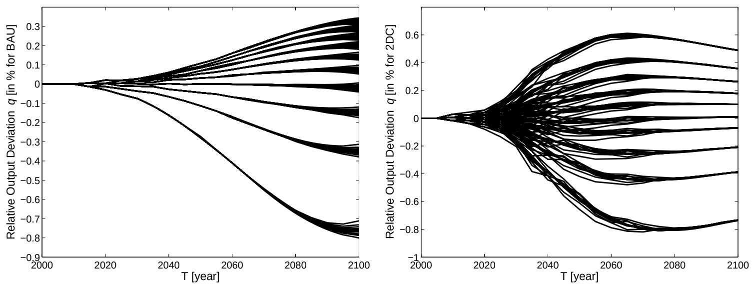

Figure 8: Intra-ensemble spread of output in terms of GWP, in units of ensemble mean output. Left graph: BAU case. Right graph: 2DC case. The spread expands when the transformation of the energy system occurs.

a high learning potential of renewable energy technologies in combination with scarce fossil resources are responsible for net output gains in the future if binding climate targets were applied.

To further disentangle the effect of the individual parameter uncertainties on the mitigation costs, we ask for the relative GWP gain across the ensemble of output streams with respect to the ensemble mean for both the BAU and 2DC cases. For this purpose we consider the ratio

$$q_{Y.n}(t) := \frac{Y_n(t) - < Y > (t)}{< Y > (t)} \tag{12}$$

whereby Y is output, n the ensemble member index and $< . >$ denotes the ensemble mean per time slice. We consider the BAU case first. The left graph in Fig. 8 reveals an increasing uncertainty in output during the century in terms of groups and sub-groups, the maximum showing $\approx 0.8\%$ loss with respect to the ensemble mean. This hierarchy of structures suggests that one of the two economic parameters (fossil resource base vs. the duplet of learning rate for reducible costs and floor costs in the renewable energy sector, henceforth summarised as "learning potential") may dominate output. Averaging over the learning potential eliminates the sub-structure so that the main variance in output must be due to uncertainty in the fossil resource base.

We follow an analogue procedure for the 2DC case and derive the right graph in Fig. 8. Variance expands around 2030 and shrinks after 2070. The maximum losses are 0.8% of GWP relative to the ensemble mean. As in the BAU case, we observe a hierarchy of structures. Averaging over the fossil resource base parameter eliminates the sub-structure. Hence for the 2DC case, the learning potential in the renewable energy sector is decisive while for the BAU case the resource base was. It should be noted, however, that the mitigation costs depicted in Fig. 7 are based on the difference between BAU and 2DC case, and therefore affected by both economic parameters.

5.3 Optimal investment paths in renewable energy

Fig. 9 displays optimal investment paths into the renewable energy sector. The stream for chance constrained optimisation (thick solid line) peaks at 3% GWP in 2050. Accordingly it would be necessary to increase the share of investments

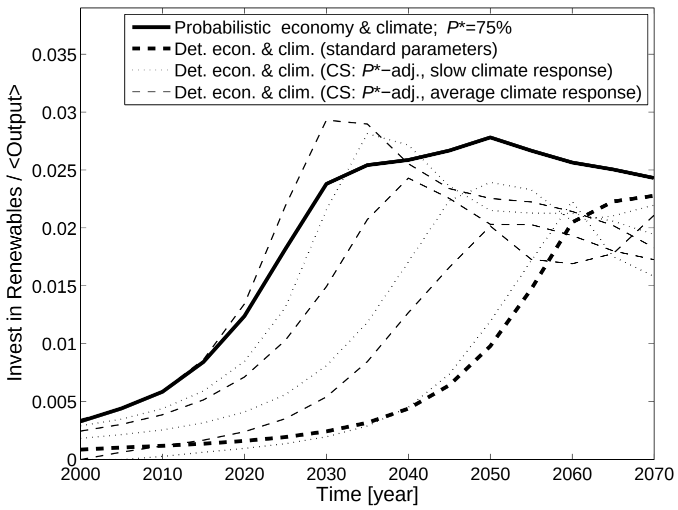

Figure 9: Optimal investment paths in the renewable energy sector, given our aggressive mitigation scenario in OGHGs; results for (a) chance-constrained programming (bold solid line), (b) standard deterministic MIND (bold dashed-line), (c) 3 curves for deterministic MIND, but with 75%-quantile-adjusted CS and probability-weighted average of climate response time scale (thin dashed lines), for 3 extreme technology parameter settings, from left to right: highest learning rate for renewables and scarcest fossil resources, centre: probability-weighted averages, right: lowest learning rate for renewables and most abundant fossil resources, (d) as (c), but for slowest climate response. Under the 75% chance constraint, investment into renewables must occur about 30 years earlier than in the standard case. Furthermore we compare the chance constrained case to possible deterministic substitutes (cases c, d), in particular to the average case (thin centre dashed line). The chance constrained case invests 10 years and cumulatively more. It resembles the "pro-renewable" case (left thin dashed curve) more than the "renewable-neutral" case. Note that the "fast climate" case leads to infeasibility and hence cannot be shown, further illustrating the limitations of an attempt to substitute for the full chance constrained solution.

into the renewable sector by more than an order of magnitude as against present-day values. Compared to standard MIND settings, investments into renewables are pulled forward by 30 years.

5.4 Prospects for bypassing chance constrained programming

Now we ask whether choosing a CS at the $P^* = 75\%$-quantile from the marginal distribution on CS would emulate our rather complicated probabilistic optimisation. Accordingly, CS would assume the value 3.6°C. In Fig. 9 we display investment streams for that elevated CS value, for the the following (3×3)-matrix of parameter settings (for the numerical parameter values, see Appendix F):

In the first dimension we let the learning rate vary from the smallest to the largest number represented in our ensemble and anti-correlate it with the extreme numbers of the fossil resource base. That way we cover the extreme cases of competition between renewable energy and fossil fuels. As centre values we choose the probability-weighted averages.

In the second dimension we vary the climate system response time scale, from the smallest to the largest value, given CS=3.6°C, supported by our joint distribution for CS and climate response time scale. We find that for the fastest climate response, there does not exist a feasible solution. This already shows that CS alone does not determine the feasibility boundary for the compliance probability P^* alone.

However, one may hope that the central deterministic combination (centre thin dashed line and centre of our (3×3)-matrix) results in a solution not too different from our chance-constrained one. First, we find that investments in renewables are underestimated and delayed by 10 years (see centre thin dashed curve). Correspondingly, emissions are higher than in the chance-constrained solution (not shown). For the deterministic solution, we find a probability of observing the 2°C target of 66% instead of 75%. Finally, the corresponding GWP loss is 0.49% instead of 0.76%. Hence the loss by the deterministic solution would have to be increased by 50% in order to emulate the loss by the chance constrained optimal solution.

In summary, the simple quantile-adjusted deterministic solution is not able to reproduce the chance constrained optimal solution to a satisfactory degree. Future work could show whether more sophisticated ways of parameter settings for deterministic solutions would obtain better matching results.

6 Conclusions

We have analysed the impact of uncertainty about (i) climate response to elevated greenhouse gas concentrations, and (ii) economic response to investments in the energy sectoron the optimal energy investment portfolio. For the analysis, we employed the state of the art model of induced technological change MIND, assuming uncertainty in fossil resource base, learning rate and floor costs of renewable energy technology, climate sensitivity and timescale of climate response. The analysis was implemented as an expected welfare maximisation with a probabilistic guardrail aiming to limit global mean temperature rise to

2°C with at least some minimum probability P^*. This constitutes a chance-constrained programming problem.

Given the inertia within our economic module, we find that a probability of $P^* = 80\%$ to observe the guardrail can be achieved at maximum, even though we already assume a very aggressive mitigation scenario for other greenhouse gases (OGHGs). When optimising expected utility under the weaker constraint $P^* = 75\%$ we find much more stringent emission cuts than those obtained from the deterministic version of MIND with default parameter values. Investment paths into renewable energy sources peak at 3% GWP between 2030 and 2050, implying an increase in investment by more than an order of magnitude against present-day values.

Only the chance-constrained approach allows deriving the optimal policy for achieving a climate target under uncertainty, and the associated distribution of GWP loss, in a meaningful way. This is because for a significant fraction of parameter combinations in our sample, the 2° target cannot be achieved under any feasible emissions path. If one considers only the feasible deterministic solutions, one obtains a strong increase of mitigation costs with climate sensitivity, i.e. 1/4 % GWP loss per degree Celsius in climate sensitivity for the aggressive OGHGs scenario and 1/2 % GWP loss per degree Celsius in climate sensitivity for a OGHGs forcing following the SRES B2 scenario. For full-fledged chance-constrained solution, we obtain a GWP loss of 0.76% (with a standard deviation of 0.13%).

Our probabilistic approach can only crudely be mimicked by a computationally less demanding deterministic counterpart: When we replace standard climate sensitivity by its (marginal) 75% quantile, i.e. 2.8°C by 3.6°C, the probabilistic economy by an averaged economy, and probabilistic climate time scale by its CS-conditioned averaged value, we observe the following deficits: (i) GWP losses would have to be corrected for by 1/2, (ii) the chance of observing the guardrail is too low by 0.1, (iii) investment streams in renewable sources are delayed and cumulatively too low. An in-depth analysis of the "irreducibility" of our probabilistic optimisation must be left for future work.

The calculations were performed by using the optimisation software GAMS. Within the GAMS environment we invented an algorithm to implement chance-constrained welfare optimisation, suitable for state of the art models of induced technological change, based on a dozen control paths, dozens of prognostic variables and investment periods. Future work has to reveal the numerical advantages and disadvantages when compared to the method by Keller et al. (2004) and McInerney and Keller (2008) that is not based on GAMS.

The numbers presented should be interpreted in a semi-quantitative manner, as the model representation of the energy system, even though sufficient to resolve the major mitigation options, is highly stylised. Due to its generality our scheme could be applied to more sophisticated integrated assessment models with higher resolution of the energy sector and more complex climate modules. Ultimately, the scheme should be extended to include hedging while anticipating future learning.

7 Acknowledgement

We would like to thank G. Hegerl, coordinating lead author of IPCC (2007b), Ch. 9, and R. Knutti, lead author of IPCC (2007b), Ch. 10, for providing us with pdfs on climate sensitivity, furthermore M. Leimbach for helpful discussions and sophisticated hints regarding our implementation in GAMS. H.H., E.K. and K.L. acknowledge support by the Volkswagen Foundation under grant number II/78470. E.K. was supported by a Marie Curie Outgoing International Fellowship(#MOIF-CT- 2005-008758) within the 6th European Community Framework Programme. His research infrastructure was partly provided by the Climate Decision Making Center at Carnegie Mellon University (National Science Foundation SES-0345798).

8 Appendix

The Appendix is available upon request from the authors.

References

Ambrosi, P., J.-C. Hourcade, S. Hallegatte, F. Lecocq, P. Dumas, and M. H. Duong (2003). Optimal control models and elicitation of attitudes towards climate damages. *Environmental Modeling and Assessment 8*, 133–147.

Annan, J. D. and J. C. Hargreaves (2006). Using multiple observationally-based constraints to estimate climate sensitivity. *Geophys. Res. Lett. 33*, L06704.

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

Bauer, N., O. Edenhofer, H. Held, and E. Kriegler (2005). Uncertainty of the role of carbon capturing and sequestration within climate change mitigation strategies. In E. Rubin, D. W. Keith, and C. F. Gilboy (Eds.), *Proceedings of 7th International Conference on Greenhouse Gas Control Technologies (Vancouver, Sep 5-9, 2004)*, Volume 1: Peer-Reviewed Paper and Plenary Presentations, pp. 931–939. Cheltenham, UK.

Brooke, A., D. Kendrick, and A. Meeraus (1992). *GAMS: A User's Guide, Release 2.25*. San Francisco: Scientific Press.

Bürgenmeier, B., A. Baranzini, C. Ferrier, C. Germound-Duret, K. Ingold, S. Perret, P. Rafaj, S. Kypreos, and A. Wokaun (2006). Economics of climate policy and collective decision making. *Climatic Change 79*, 143–162.

Charnes, A. and W. W. Cooper (1959). Chance constrained programming. *Management Science 6*, 73–79.

Commission of European Communities (2007). *Limiting global climate change to 2 degrees Celsius: The way ahead for 2020 and beyond*. Commission of European Communities; Brussels.

Drud, A. S. (1992). CONOPT: A large-scale GRG code. *ORSA Journal on Computing 6*, 207–216.

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

Edenhofer, O., C. Carraro, J. Köhler, and M. Grubb (2006). *Energy Journal: Special Issue (27), Endogenous Technological Change and the Economics of Atmospheric Stabilization*. International Association for Energy Economics.

Edenhofer, O., K. Lessmann, and N. Bauer (2006). Mitigation strategies and costs of climate protection: The effects of ETC in the hybrid model MIND. *Energy Journal Special Issue: Endogenous Technological Change and the Economics of Atmospheric Stabilization*, 89–104.

Gerlagh, R. and B. van der Zwaan (2004). A sensitivity analysis of timing and costs of greenhouse gas emission reductions. *Climatic Change 65*, 39–71.

Hansen, J. E. and M. Sato (2001). Trends of measured climate forcing agents. *Proceedings of the National Academy of Science 98*, 14778–14783.

Hare, B. and M. Meinshausen (2006). How much warming are we committed to and how much can be avoided? *Climatic Change 75*, 111–149.

IPCC (2001). *Climate Change 2001: The Scientific Basis*. Contribution of Working Group I to the Third Assessment Report of the IPCC, Appendix II. Cambridge University Press, Cambridge.

IPCC (2007a). *Climate Change 2007: Mitigation*. Contribution of Working Group III to the Fourth Assessment Report of the IPCC. Cambridge University Press, Cambridge.

IPCC (2007b). *Climate Change 2007: The Physical Science Basis*. Contribution of Working Group I to the Fourth Assessment Report of the IPCC. Cambridge University Press, Cambridge.

Kann, A. and J. P. Weyant (2000). Approaches for performing uncertainty analysis in large-scale energy/economic policy models. *Environmental Modeling and Assessment 5*, 29–46.

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

Kleinen, T. (2005). *Stochastic Information in the Assessment of Climate Change*. PhD-Thesis, Potsdam University, Germany.

Knight, F. H. (1921). *Uncertainty, and Profit.* Boston, MA: Hart, Schaffner & Marx; Houghton Mifflin Company.

Kriegler, E. and T. Bruckner (2004). Sensitivity analysis of emission corridors for the 21^{st} century. *Climatic Change 66*, 345–387.

Kriegler, E., H. Held, and T. Bruckner (2007). Climate protection strategies under ambiguity about catastrophic consequences. In J. Kropp and J. Scheffran (Eds.), *Advanced Methods for Decision Making and Risk Management in Sustainability Science*, pp. 3–42. Nova Science, New York.

Lempert, R. J. (2002). A new decision sciences for complex systems. *PNAS 99*, 7309–7313.

Manne, A. and R. Richels (1997). On stabilizing CO_2 concentrations – cost-effective emission reduction strategies. *Environmental Modelling and Assessment 2*, 251–265.

Mastrandrea, M. D. and S. H. Schneider (2004). Probabilistic integrated assessment of 'dangerous climate change'. *Science 304*, 571–575.

McInerney, D. and K. Keller (2008). Economically optimal risk reduction strategies in the face of uncertain climate thresholds. *Climatic Change 91*, 571–575.

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

Nakićenović, N. and R. Swart (2000). *Emissions Scenarios. Special Report of the IPCC.* Cambridge University Press, Cambridge.

Nordhaus, W. D. (1994). *Managing the global commons; The Economics of Climate Change.* The MIT Press, Cambridge, Massachusetts.

Nordhaus, W. D. and J. Boyer (2000). *Warming the world.* The MIT Press.

Peterson, S. (2006). Uncertainty and economic analysis of climate change: A survey of approaches and findings. *Environmental Modeling and Assessment 11*, 1–17.

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

Richels, R. G., A. S. Manne, and T. M. L. Wigley (2004). Moving beyond concentrations: The challenge of limiting temperature change. *AEI-Brookings joint center for regulatory studies 04-11.*

Schneider von Deimling, T., H. Held, A. Ganopolski, and S. Rahmstorf (2006). Climate sensitivity estimated from ensemble simulations of glacial climates. *Climate Dynamics 27*, 463–483; DOI 10.1007/s00382–006–0126–8.

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

Syri, S., A. Lehtilä, I. Savolainen, and T. Ekholm (2007). Global energy and emission scenarios for effective climate change mitigation. *VTT Research Report VTT-R-03592-07.*

Valverde, L. J., H. D. Jacoby, and G. M. Kaufmann (1999). Sequential climate decisions under uncertainty: An integrated framework. *Environmental Modeling and Assessment 4*, 87–101.

van der Zwaan, B. and R. Gerlagh (2006). Climate sensitivity uncertainty and the necessity to transform global energy supply. *Energy 31*, 2571–2587.

van Vuuren, D., M. G. J. den Elzen, M. M. Berk, P. Lucas, B. Eickhout, H. Eerens, and R. Oostenrijk (2003). Regional costs and benefits of alternative post-kyoto climate regimes – comparison of variants of the multi-stage and per capita convergence regimes. RIVM report 728001025/2003, National Institute of Public Health and the Environment, Bilthoven, The Netherlands.

Yohe, G., N. Andronova, and M. Schlesinger (2004). To hedge or not against an uncertain climate future? *Science 306*, 416–417.

Yohe, G., M. Schlesinger, and N. Andronova (2006). Reducing the sirk of a collapse of the atlantic thermohaline circulation. *The Integrated Assessment Journal 6*, 57–73.