Supplement to
Operationalizing Climate Targets under Learning: An Application of Cost-Risk Analysis

Delf Neubersch[*], Hermann Held[†], Alexander Otto[‡]

June 22, 2014

This online resource (electronic supplementary material) covers the following points:

Contents

[*]Research Unit Sustainability and Global Change, University of Hamburg and
International Max Planck Research School on Earth System Modelling
Grindelberg 5, Hamburg 20144, Germany
delf.neubersch@zmaw.de
[†]Research Unit Sustainability and Global Change, University of Hamburg – KlimaCampus
Grindelberg 5, Hamburg 20144, Germany
[‡]Environmental Change Institute, University of Oxford
South Parks Road, Oxford, OX1 3QY, United Kingdom

1 Threshold function and climate sacrifice

The probability of violating the target could be used as a risk metric. This is the intuitive way that a climate risk would be formulated if the concept of a guardrail is supported and also the main method suggested in Schmidt et al (2011):

$$R(T) = \Theta(T - T_{\mathrm{g}}) \tag{1}$$

Here Θ is the Heaviside function. This would imply that above T_{g} the risk is unity and by aggregating over states of the world a probability of violating the guard rail emerges in that time instant. The overall probability can then be calculated by two methods further discussed in Section 8.2.

However, for our set-up, this option leads to a so far unreported behavior in the case of learning, more specifically, perfect learning. After a perfect learning event the decision maker will choose the optimal path for each of the states of the world separately. The trade-off is assessed individually and so simplifies to the trade-off between a constant penalty of crossing the guard rail and the cost of mitigation. If the cost of mitigation is too high, in the case of a high climate sensitivity, it is cost effective to accept the constant penalty and abandon mitigation. This would sacrifice the climate if it is too expensive to return below the guard rail. Figure 1 shows example temperature paths for the model MIND-L with and without perfect learning.

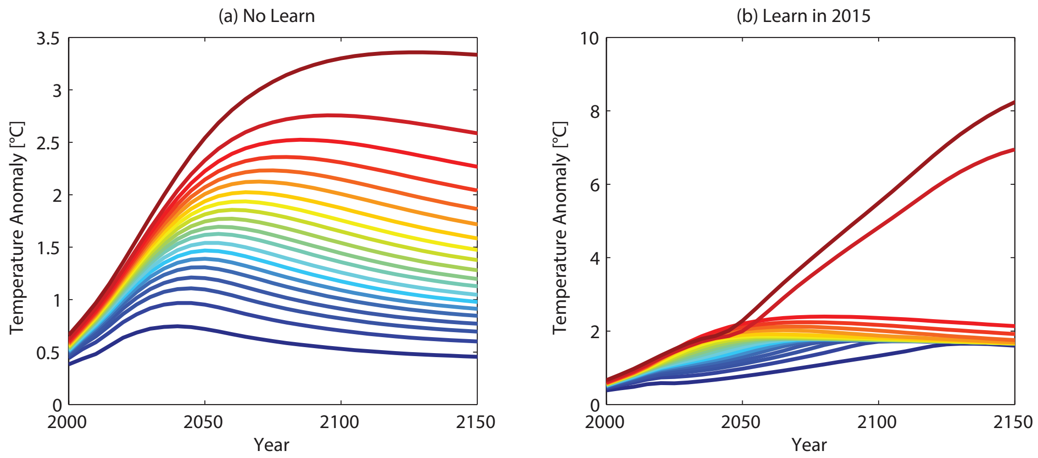

Figure 1: Temperature paths for a scenario with (a) no learning and (b) perfect learning in 2015. The color scale indicates the climate sensitivity (red: high, blue: low). The climate sacrifice when learning that the climate sensitivity is high is clearly shown in panel (b).

2 Conceptual Difference between CBA and CRA

Investigation of the value of information in an integrated assessment model in the past was done by using traditional CBA with damage functions (Nordhaus and Popp, 1997; Webster et al, 2008). Our results can be loosely compared to these studies, however, CRA is a conceptually different approach to tackling the climate problem under uncertainty and learning. A difference originates from the different methods of closing the impact loop. CRA does not influence the economy directly and instead closes the loop on the utility level.

On the conceptual level, CRA encapsulates an implicit risk perception for transgressing the target, including a plethora of potential, ill-described damages. The vagueness of this knowledge serves as an argument to trade that risk perception against mitigation costs in a most simplistic manner, by subtracting its expected value from utility.

To clarify the conceptual difference we consider the case without learning: in a CBA with the same set-up in terms of uncertainty like ours, the damages will indirectly become uncertain because they are based on uncertain temperature. This will lead to uncertain consumption paths. In a CRA however, the consumption path is deterministic because the uncertainty comes into the equation after utility from consumption has been calculated, i.e. in the utility equation.

On a formal level, CRA resembles CBA in the sense that both employ an expected utility framework with a function that expresses the costs due to temperature rise. However, the methods diverge in their reasoning: CBA tries to account for observable damages, whereas CRA calibrates the trade-off to fit normative expectations.

3 Perfect Learning

Perfect learning is modeled by assuming that the decision maker receives one of S messages at the learning point. If message m_s is received, the state of the world s attains a probability of 100%. The probability of receiving the message m_s is equal to the *a priori* probability of that state of the world s , i.e. p_s. This is referred to as consistent learning. In essence, the model interpretation changes from one uncertain world with S possible outcomes to S certain worlds each with a probability of occurring. From the next time step onwards, after the learning event, the decision maker can optimally adjust the investments in each certain world.

Learning about climate sensitivity from temperature observations is a gradual process and Kelly and Kolstad (1999) have estimated it to take at least 90 years. This could be shortened by observations of other variables of interest and usage of better models. Webster et al (2008) find that uncertainty in climate sensitivity can be reduces by 20-40% in the next two to five decades. The necessary investment to benefit from this information 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, where the late stage represents the learning scenario where information only originates from new observations. To test for the sensitivity of the late learning assumption we do the calculations with two reference years for late learning: 2050 and 2075.

It appears plausible for us that a business-as-usual, i.e. un-accelerated scientific progress will reduce climate response uncertainty in a couple of decades. Hence, assuming a much smaller uncertainty in 2050 or 2075 compared to today's appears plausible. 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 a good approximation as an intermediate value. While climate sensitivity values that are at the low end will be learned in a couple of decades, those at the high-end require at least centennial scale (as low climate sensitivity values come with short intrinsic time scales – see e.g. Lorenz et al (2012) – these require less time for being learned). The former would allow for higher emissions, activated before 2050, however, not immediately, but maybe in 2030. As we show, a time-shift of the later learning by one or two decades does not change the result substantially, here the more so as this effect applies only to a small proportion of the states of the world would contribute to the expected value. On the other hand, for large values of climate sensitivity, effective learning would appear rather late and hence would compensate the effects of early learning for small values. In total, perfect learning in 2050 or 2075 seems a good reference point.

We recognize that it can not always be clearly stated in which direction the incorporation of learning affects the optimal policy, although evidence is accumulating that it does not support postponing mitigation (Ha-Duong et al, 1997; Ha-Duong, 1998; Lange and Treich, 2008). We do not elaborate on this discussion, but rather focus on the novel method of analysis.

4 Reasons for Calibration with NOLEARN

The reader may wonder why we prefer the "no-learning" case to the "anticipated learning" case for the calibration. While in principle the calibration could have been performed for the latter case, to our impression the following two arguments would individually lead to a preference for calibrating to the "no-learning" case. (i) While both the CBA- as well as the CEA-based literature discusses anticipated learning (see e.g. Kelly and Kolstad (1999) and Webster et al (2008) respectively) this academic discussion had little influence on the scientific basis of the "Copenhagen Diagnosis" (Allison et al, 2009) that provided a condensed update after IPCC-AR4 for the decisive COPs 2009 and 2010. In the diagnosis there is no mention of learning or a reduction of uncertainty in the future. Moreover, if anticipated future learning had been a major issue during those negotiations, the aforementioned limitations of probabilistic targets would have become apparent, and some tool such as CRA would have been requested. However this was not the case. (ii) If one decided to calibrate a CRA-model to 66% compliance under anticipated learning, one would face a severe further conceptual problem: after learning, for most states of the world temperature paths tend to cluster around the guardrail. This implies that the diagnosed compliance probability becomes a very strong function of the numerical guardrail value. This is truly out-of-touch with the intuition of the concept of a compliance probability.

Figure 4 illustrates point (ii) why a calibration with the NOLEARN case is more desirable that with the LEARN case although the result is not very different. The calculation of the safety is difficult for LEARN cases.

5 Certainty and Balanced Growth Equivalents

The concept of Certainty and Balanced Growth Equivalents (CBGE) has been discussed at length in Anthoff and Tol (2009) and was introduced by Mirrlees and Stern (1972). Comparing a welfare against a reference welfare is done by calculating the level of initial consumption levels that are necessary to reach the levels of welfare under the assumption of constant growth. The relative change of initial consumption is calculated and used to compare the two scenarios. Suppose the constant growth rate is represented by a, the balanced growth equivalent initial consumption can be found by replacing the consumption in the CRRA utility equation by $e^{at}C_0$ representing a constant growth from an initial level C_0. All derivations here are done for $\eta \neq 1$, for the special case see Anthoff and Tol (2009).

$$W = \sum_{t=0}^{t_{\text{end}}} \sum_{s=1}^{S} p_s \cdot P_t \frac{1}{1-\eta} \cdot \left(\frac{e^{at}C_0}{P_t}\right)^{1-\eta} e^{-\Delta t \cdot \delta \cdot t} \tag{2}$$

Rearranging for C_0 and applying $\sum_{s=1}^{S} p_s = 1$:

$$C_0 = \left[(1-\eta)\, W\right]^{\frac{1}{1-\eta}} \left[\sum_{t=0}^{t_{\text{end}}} P_t \left(\frac{e^{at}}{P_t}\right)^{1-\eta} e^{-\Delta t \cdot \delta \cdot t}\right]^{-\frac{1}{1-\eta}} \tag{3}$$

Calculating the initial consumption for two scenarios in this way the relative change can be calculated.

$$\Delta\text{CBGE} = \frac{C_0 - C_{0,\text{ref}}}{C_{0,\text{ref}}} \tag{4}$$

Two scenarios that have the same labor development and utility functions can be compared directly without calculating the initial consumption explicitly. The growth rate a cancels out leaving:

$$\Delta\text{CBGE}(W_1, W_{\text{ref}}) = \left(\frac{W_1}{W_{\text{ref}}}\right)^{\frac{1}{1-\eta}} - 1 \qquad \text{for}\, \eta \neq 1 \tag{5}$$

This is calculated in the stochastic regime, i.e. W is seen as the expected welfare. The CBGE can be interpreted as a permanent relative loss to consumption every year.

6 Comparison of Scenarios

We define two symbols L_1 and L_2 representing two learning points where L_1 is earlier than L_2, i.e. $L_1 < L_2$. We introduce two ΔCBGE quantities:

- **Expected Welfare Loss due to Climate Change**: the cost created by including the calibrated risk metric. This cost includes the cost of mitigation as well as the loss from the risk metric. Calculated by the percentage change of CBGE between a LEARN(L) or NOLEARN scenario and the BAU scenario (using eq. 5 above). Note that if the reference is the BAU case, the resulting value will be negative, therefore, the welfare loss is defined as the negative change in CBGE.
- **Expected Value of Perfect Information** EVPI(L_1, L_2): calculated by the percentage change of CBGE between LEARN(L_1) and LEARN(L_2) scenarios. Gives the value created by receiving the information in year L_1 instead of year L_2. The learn scenario with the latest learning point is chosen as the reference.

7 The Model 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 (2012). 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 integrated assessment model 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), MESSAGE (Messner and Strubegger, 1995), REMIND (Luderer et al, 2013). MIND-L is a forward-looking Ramsey-type macro-economic growth model that comprises induced technological change in the energy sector, the latter consisting of a renewable and a fossil sector. An energy balance model represents the climate module within MIND-L that links emissions to global mean temperature change. MIND-L is implemented in the optimization language GAMS.

All the simulations run until 2200. For all non-CO_2 GHGs a stringent mitigation trajectory is assumed (RCP 3-PD scenario (van Vuuren et al, 2007)). Finally, the model has been constrained to CO_2 emission data up to year 2010. To cope with the finite horizon effect the last element of the discount vector was adjusted. See Section 8 for more detail on numerical issues.

7.1 Uncertainty

The uncertainty that is considered in the simulations stems from uncertainty about the climate sensitivity correlated to the ocean heat uptake. The log-normal form: pdf(CS) = $\mathcal{LN}(0.973, 0.4748)$ (Wigley and Raper, 2001) with indications of the samples that where take is shown in Figure 2. The samples where taken in 5% quantiles. The ocean heat uptake is calculated from the climate sensitivity by assuming a correlation presented in the Appendix 1 of Lorenz et al (2012).

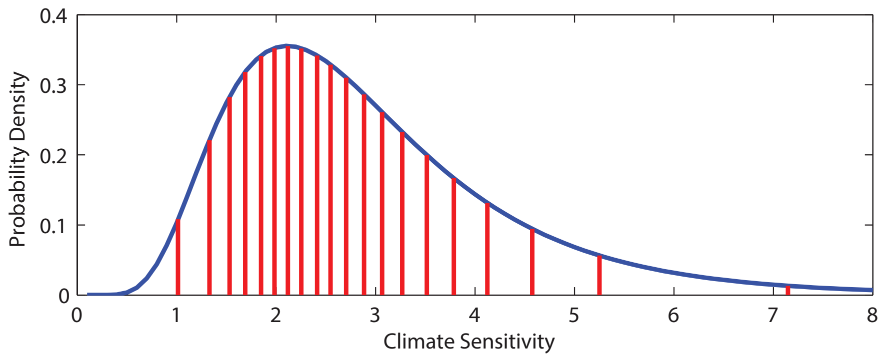

Figure 2: The probability distribution log-normal form: pdf(CS) = $\mathcal{LN}(0.973, 0.4748)$ (Wigley and Raper, 2001) with red lines indicating the sampled values

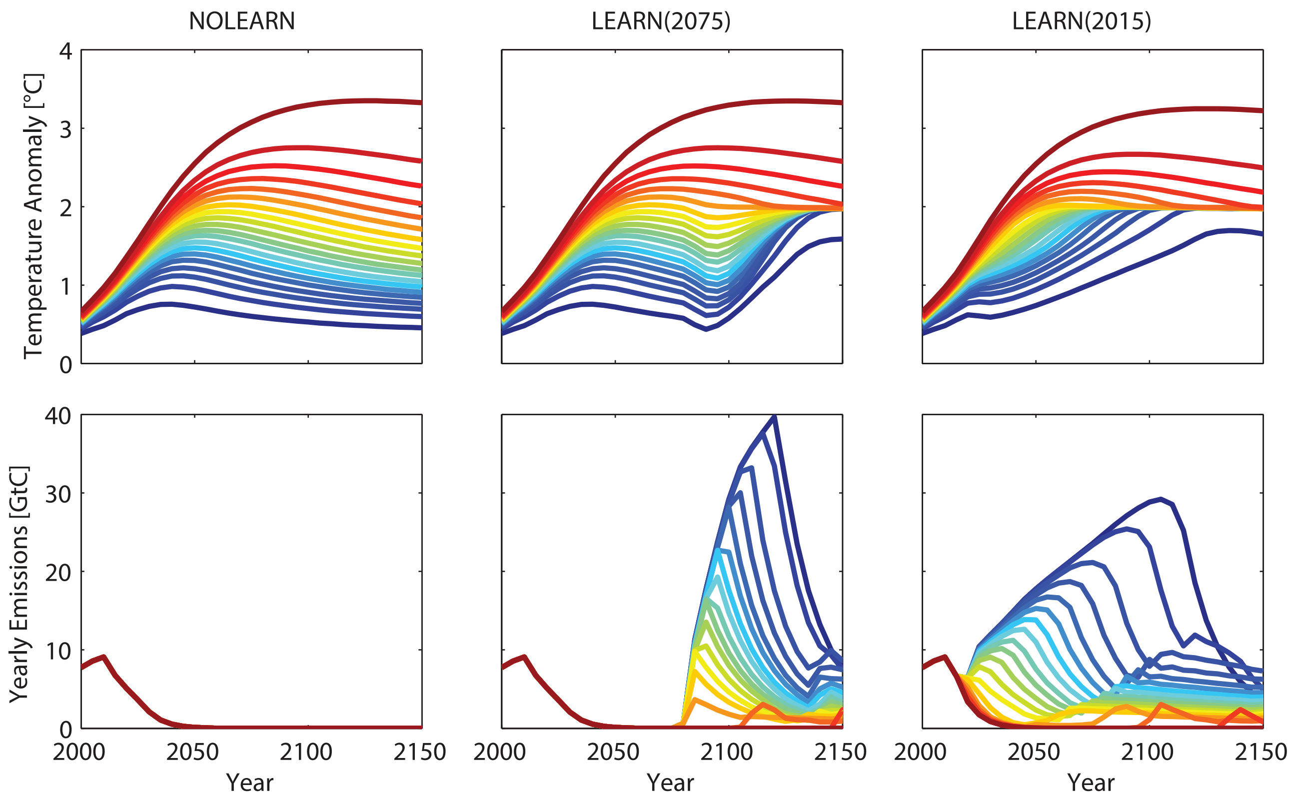

Figure 3: Temperature (top) and emission (bottom) paths for NOLEARN, LEARN(2075) and LEARN(2015) (from left to right respectively). The color indicates the value of CS (red: high, blue: low)

7.2 The Base Case

Figure 3 shows temperature and emission paths of the base model setup for NOLEARN, LEARN(2075) and LEARN(2015). The base setup consists of a constant relative risk aversion of 2, a pure rate of time preference of 1% and a calibration to a safety of 66% for a target of 2°C in the NOLEARN scenario.

In a NOLEARN scenario the control (emissions) is equal for all state of the worlds (see bottom left panel in Figure 3). Emissions drop to virtually zero around 2050 to be able to comply with the 2°C target. After the learning point, each state of the world can be controlled separately and so all worlds will strive towards the guard rail which acts as an attractor if the risk metric is not too concave. For high climate sensitivity (red lines) learning provokes slightly stronger mitigation which is more expensive. Learning that the climate sensitivity is low brings benefit compared to a case without learning by allowing for more emissions (blue lines).

8 Numerical Considerations

To reduce the risk that the solution is a local optimum the model is always run multiple times, with different initial guesses. The idea is to initialize learning cases with a solution from BAU and a solution from NOLEARN. This encourages the solver to search for a solution starting at high temperatures and starting at low temperatures respectively. If the solution is equal, one can be fairly certain that it is globally optimal. If it is different, the better solution is chosen.

Restarting a model from a previous solution that was flagged by GAMS/CONOPT as non-optimal will lead to a refinement of the solution. In the analysis presented here about 4 to 6 iterations, restarting GAMS/CONOPT with the last solution, are necessary to converge on a solution.

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 which can be justified by pointing out that measurement errors are on the order of 0.1°C which is of the same order of magnitude as the smoothing by the error functions.

8.1 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.

8.2 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 curves cannot cross, 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 the looks for the maximum:

$$P_{\text{soft}} = \max_t \left\{ \sum_{s=0}^{S} p_s \Theta \left(T(t,s) - T_{\text{g}} \right) \right\}. \tag{6}$$

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

$$P_{\text{hard}} = \sum_{s=0}^{S} p_s \Theta \left(\max_t \left\{ T(t,s) \right\} - T_{\text{g}} \right) \tag{7}$$

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

For 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 4).

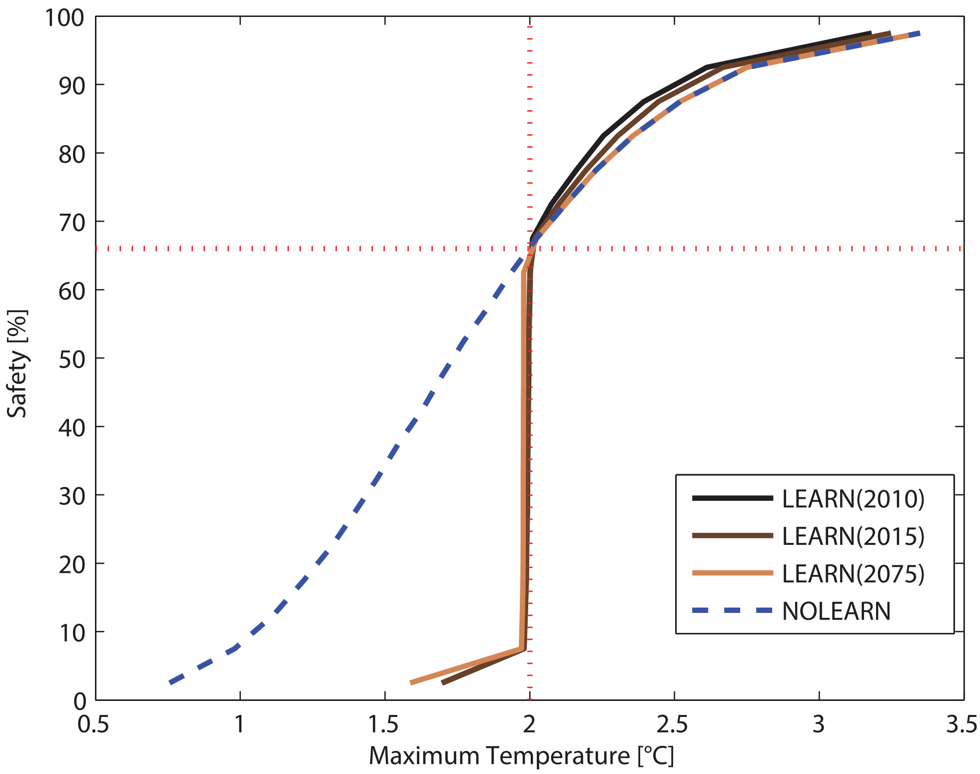

Figure 4: Cumulative distribution function for the maximum temperature for NOLEARN and LEARN case and other parameters kept at base values. It is evident that for NOLEARN a clear safety for 2°C can be determined whereas for LEARN cases it is a very strong function of temperature

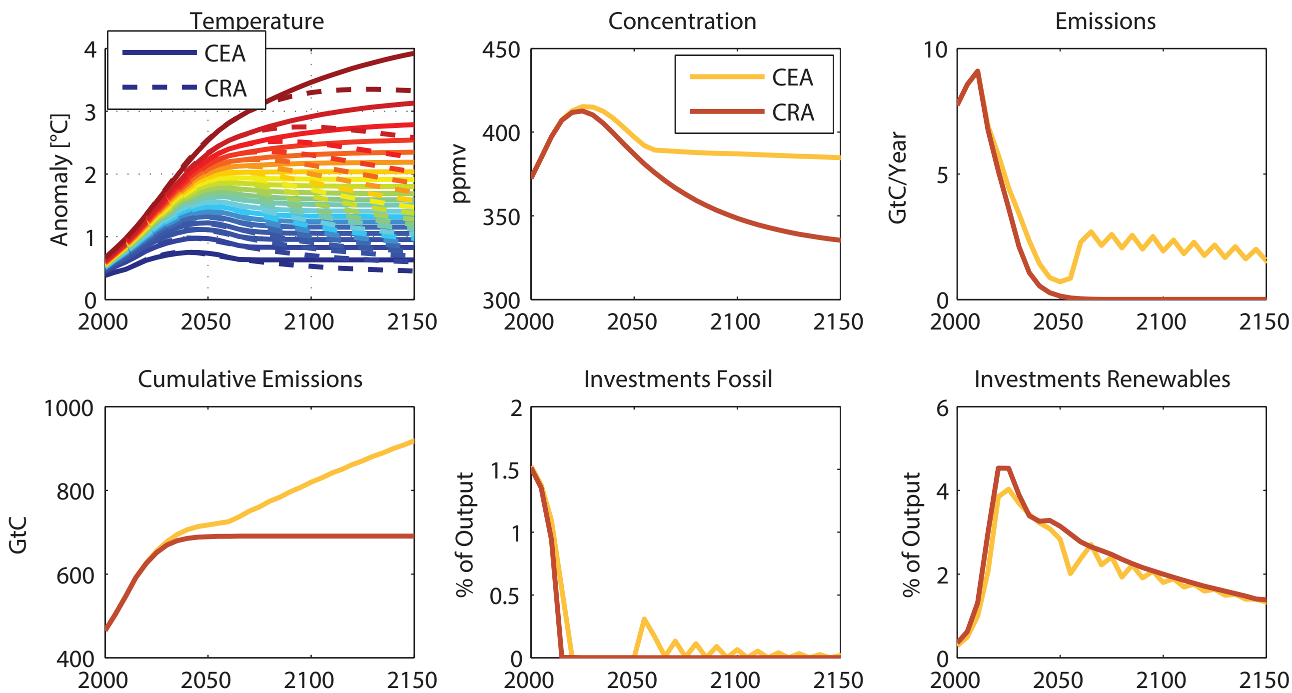

Figure 5: Time paths for different variables comparing CEA and CRA for a base case

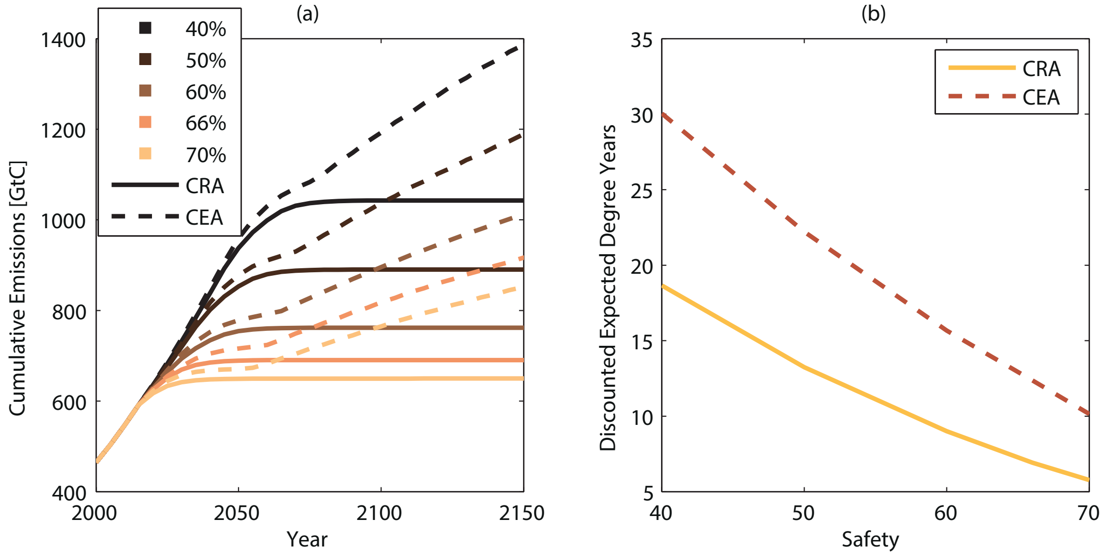

Figure 6: Comparing CEA and CRA by looking at (a) cumulative emissions until 2200 and (b) discounted expected degree years above the guard rail (2°C). The x-axis gives the safety for which the trade-off parameter was calibrated. It is clear that although CEA and CRA have the same safety, CRA mitigates more and has less risk than CEA.

9 Comparison to CEA

Figure 5 shows temperature, concentrations, emissions and investments over time for CRA and CEA simulations, both with a safety of 66% of staying below 2°C. Figure 6 shows cumulative emissions and discounted expected degree years for CRA and CEA runs with the same safety and guard rail (2°C). Panel (a) underlines the equality of the solutions up to 2050 and (b) shows that CEA always has a higher temperatures through less mitigation after 2050. This change can also be seen in Figure 5 where there is a clear cut regime change for CEA around 2060. This is because CEA will stick to 2°C 66% as close as possible until the end of the simulation and therefore is emitting to compensate for a natural reduction in greenhouse gases, for example, by ocean uptake.

10 Sensitivity to Normative Parameters

Most plots in the article are plotted against the social discount rate d, here we explain why. We varied constant relative risk aversion η through [0.75, 1.5, 2, 3, 4, 5] and pure rate of time preference δ through [0.5%, 1%, 2%, 4%]. Risk aversion has two effects: discounting over time due to consumption growth as well as risk aversion in one time slice. Due to very similar consumption streams across state of the worlds, risk aversion has little effect. We therefore concentrate on the discounting aspect and recall the Ramsey formula for the social discount rate $d = \eta g + \delta$. This combines the discounting effect that originates from the growth of consumption g and the consumption smoothing tendency (inter-temporal elasticity of substitution) of the risk aversion parameter η with the pure time preference to an overall discount rate. The growth of consumption g is endogenous in the model and lies around 2.2% pa, again depending on discounting. This leads to social discount rates roughly in the range between 2% and 12%.

11 Origin of Value and Anticipation

Due to the additive nature of the risk metric in CRA, the question might arise which part of the welfare function, consumption utility change or risk change, is responsible for a particular value of information. For example, the EVPI for learning in 2015 instead of 2075 is 0.66% CBGE, but how much of this value is due to a reduction in temperature and how much is due to an increase in economic activity? The decision maker might even forfeit consumption to decrease temperature if information is received early enough or vice versa. Furthermore, to calculate the economic mitigation costs, a method is needed to separate the economic and the risk-related effect.

To uncover this information, we need to divide the value into its risk-related and economic part. However, due to a non-linear utility function, any partial utility change affects the CBGE change depending on the current level of utility. One way to meet this problem is to linearize the utility function at the points of interest. Lets assume that we have two scenarios with two welfare results W_1 and W_2. The CBGE change between them is denoted by Γ and calculated as shown in Section 5. We further know that $W_i = W_i^E + W_i^R$ where the parts stand

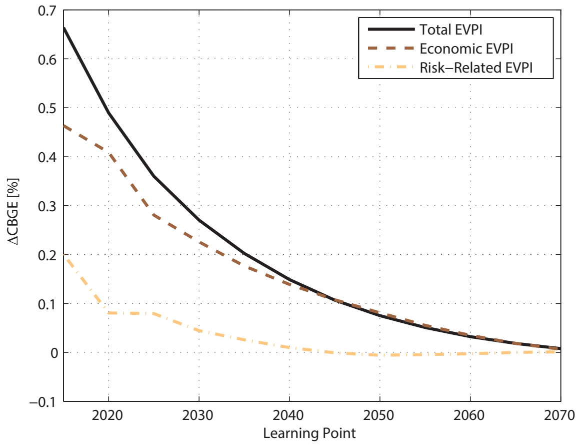

(a) EVPI

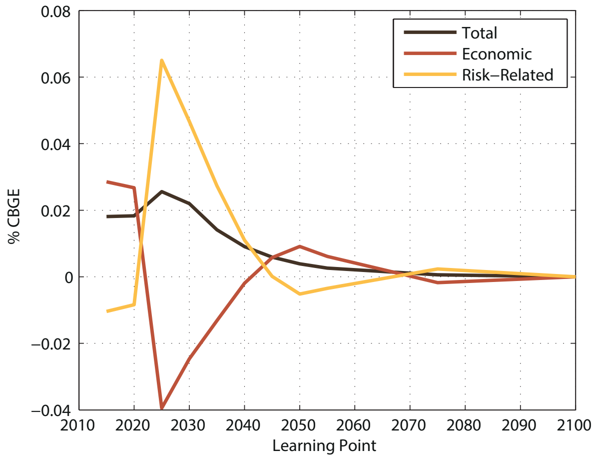

(b) Value of anticipation

Figure 7: Origin of Value and Anticipation for EVPI(2015,2075) in the base case and different learning time points decomposed into risk related and economic value

for the welfare from the economy and risk-related parts respectively. Disentangling the value into $\Gamma{=}\Gamma^E + \Gamma^R$ is done by dividing the total value into the same ratios as the differences in welfare:

$$\Gamma^E = \Gamma \cdot \frac{W_2^E - W_1^E}{W_2 - W_1} \quad \text{and} \quad \Gamma^R = \Gamma \cdot \frac{W_2^R - W_1^R}{W_2 - W_1} \tag{8}$$

Figure 7a shows the separate parts of EVPI(2015,2075) for different learning points. Figure 7b shows the value gained only by anticipating a learning event. It only makes up about 3% of the total EVPI. This is very small and discussing it further has no merit.

12 Other Studies

Knowledge of the monetary value of information is of great importance to decision makers when planning projects that have information as output, such as satellite missions, measuring networks or research programs.

Converting the CBGE changes into monetary values is a risky business but could be done by multiplying with current and projected consumption worldwide every year. Assuming 71.83 trillion USD for 2012 and a consumption share of 75% would yield 355 billion USD(2012) for 0.66 CBGE% in 2012. This supports the notion that if we take a guard rail of 2°C seriously, the information about the climate becomes more valuable than currently acknowledged. This value includes the fictive value of reducing the risk from high temperatures and the economic value of reducing the cost of mitigation. The purely economic value will depend on the scenario and normative parameters. This will be discussed in the follow up article. For a 3°C guard rail we would get values of below 53 billion USD(2012) in 2012. All of these values will increase over time as consumption increases because they are given as percentages of consumption in each year.

There are only a few studies that try to quantify the value of perfect knowledge about the climate response. Nordhaus and Popp (1997) estimate the expected value of perfect information about climate feedback if the information is received in 2015 instead of 2045 as 7.8 billion USD(1990) net present value. However, they have a high probability of violating 2°C (around 80%) so the study could only be compared to a CRA with a higher guard rail of 3°C or 4°C. The purely economic value calculated by such a CRA would lie well below 53 billion USD(2012). Considering the increase in GWP from 1990 to 2012 we can argue that our approach lies in the same order of magnitude as Nordhaus and Popp (1997). In Held et al (2009) CBGE changes were shown to be linearly correlated with net present value with a factor of around 1.5 for similar discount rates used here. This reduces the gap even further.

Webster et al (2008) calculate values of perfect information in 2020 compared to no learning and find 24 billion USD(2005). They argue that this is very low due to the low damages in the model, equating it with a CEA with a target above 4°C. They then calculate optimal carbon tax for a 2°C target with a cost effectiveness analysis and a 100% probability of reaching the target. This is only possible because they do not include infeasible climate sensitivities that would lead to a inevitable violation of the target. They do not give the value of such information. This is exactly the case where the value can be negative and probably is.

References

Allison I, Bindoff NL, Bindschadler R, Cox P, de Noblet N, England M, Francis J, Gruber N, Haywood A, Karoly D, Kaser G, Quéré CL, Lenton T, Mann M, McNeil B, Pitman A, Rahmstorf S, Rignot E, Schellnhuber H, Schneider S, Sherwood S, Somerville R, Steffen K, Steig E, Visbeck M, Weaver A (2009) The Copenhagen Diagnosis, 2009: Updating the world on the latest climate science. Tech. rep., The University of New South Wales Climate Change Research Centre (CCRC), Sydney, Australia

Anthoff D, Tol RSJ (2009) The Impact of Climate Change on the Balanced Growth Equivalent: An Application of FUND. Environ Resource Econ 43(3):351–367, DOI 10.1007/s10640-009-9269-5

Bosetti V, Carraro C, Galeotti M, Massetti E, Tavoni M (2006) WITCH - A World Induced Technical Change Hybrid Model. University Ca' Foscari of Venice Economic Research Paper No 46/06 DOI 10.2139/ssrn.948382

Edenhofer O, Bauer N, Kriegler E (2005) The impact of technological change on climate protection and welfare: Insights from the model mind. Ecolog Econ 54(2-3):277–292, DOI 10.1016/j.ecolecon.2004.12.030

Ha-Duong M (1998) Quasi-option value and climate policy choices. Energy Econ 20(5-6):599–620, DOI 10.1016/S0140-9883(98)00011-5

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

Held H, Kriegler E, Lessmann K, Edenhofer O (2009) Efficient climate policies under technology and climate uncertainty. Energy Econ 31:S50–S61, DOI 10.1016/j.eneco.2008.12.012

Hope C (2011) The Social Cost of CO_2 from the Page09 Model. Economics Discussion Paper No 2011-39 DOI 10.2139/ssrn.1973863

Kelly DL, Kolstad CD (1999) Bayesian learning, growth, and pollution. Journal of Economic Dynamics and Control 23(4):491–518, DOI 10.1016/S0165-1889(98)00034-7

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

Lorenz A, Schmidt MGW, Kriegler E, Held H (2012) Anticipating climate threshold damages. Environ Model Assess 17(1-2):163–175, DOI 10.1007/s10666-011-9282-2

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

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. J Econ Theory 4(2):268–288, DOI 10.1016/0022-0531(72)90153-6

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

Nordhaus WD, Popp D (1997) What is the value of scientific knowledge? An application to global warming using the PRICE model. Energy J 18(1):1–45

Schmidt MGW, Lorenz A, Held H, Kriegler E (2011) Climate targets under uncertainty: challenges and remedies. Clim Change 104(3-4):783–791, DOI 10.1007/s10584-010-9985-4

van Vuuren DP, Elzen MGJ, Lucas PL, Eickhout B, Strengers BJ, Ruijven B, Wonink S, Houdt R (2007) Stabilizing greenhouse gas concentrations at low levels: an assessment of reduction strategies and costs. Clim Change 81(2):119–159, DOI 10.1007/s10584-006-9172-9

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

Wigley TM, Raper SC (2001) Interpretation of high projections for global-mean warming. Science 293(5529):451–4, DOI 10.1126/science.1061604