Optimally Climate Sensitive Policy under Uncertainty & Learning

Svenn Jensen* Christian Traeger†

November 2013

Abstract

The equilibrium response of the global temperature to greenhouse gas emissions is highly uncertain. We derive the optimal climate policy under uncertainty, acknowledging Baysian uncertainty, passive and active learning, and temperature stochasticity. Our analysis employs a stochastic dynamic programming implementation of the integrated assessment model DICE (Nordhaus, 2008). We find that the stochasticity of temperatures induces precautionary savings, while Bayesian uncertainty over the climate's sensitivity to greenhouse gas emissions increases the optimal present day carbon tax by approximately 25%. Currently, the scientific community does not agree on the correct Bayesian prior or even its expected value. We therefore re-evaluate optimal policy using a model of smooth ambiguity aversion, acknowledging low confidence into the Bayesian prior. We find that neither ambiguity, nor the anticipation of learning change the optimal policy.

JEL Codes: Q54, Q00, D90, C63
Keywords: climate change, uncertainty, ambiguity aversion, smooth ambiguity model, Bayesian learning, recursive utility, dynamic programming, integrated assessment, DICE

*Department of Agricultural & Resource Economics, UC Berkeley, 207 Giannini Hall #3310, Berkeley, CA 94720-3310, USA; svenn@berkeley.edu.

†Department of Agricultural & Resource Economics, UC Berkeley, 207 Giannini Hall #3310, Berkeley, CA 94720-3310, USA; traeger@berkeley.edu

1 Introduction

Anthropogenic greenhouse gas emissions are presently changing the energy balance of our planet. Various climatic feedbacks make the resulting warming over the next decades and centuries highly uncertain. The economic losses from climate change are convex in the temperature and, thus, the optimal mitigation policy is very sensitive to the parameterization of this climatic response to anthropogenic emissions. How should we tax (or cap) greenhouse gases today acknowledging the climatic uncertainty?

Climate sensitivity captures the equilibrium warming from doubling the CO_2 concentration with respect to preindustrial levels.[1] The scientists on the Intergovernmental Panel on Climate Change (IPCC) find that "equilibrium climate sensitivity is likely in the range 1.5°C to 4.5°C", adding that "no best estimate for equilibrium climate sensitivity can now be given because of a lack of agreement on values across assessed lines of evidence and studies" (IPCC, 2013). The economic losses from climate change differ largely across this likely interval: the loss of global production with a 4.5°C warming is 7 times higher than in the case of a 1.5°C warming, using the wide-spread damage function of the DICE model (Nordhaus, 2008). Moreover, climate sensitivity values ranging 4.5°C-10°C still carry considerable probability mass (Meinshausen et al., 2009). We introduce uncertainty about climate sensitivity into a stochastic dynamic programming integrated assessment model of climate change. Our integrated assessment model closely resembles the DICE model by Nordhaus (2008), which integrates emissions and climate into a Ramsey growth model. Economic production results in greenhouse gas emissions, which accumulate in the atmosphere, cause lagged warming, and ultimately impact economic output. The decision maker controls the mitigation policy and investment into produced capital. Like most integrated assessment models, the original DICE model is deterministic and cannot determine the optimal policy unless the decision maker knows (or is assumed to know) the climate's response to emissions with certainty. Our stochastic dynamic programming implementation enables us to derive optimal policies in the face of

[1] We are currently 70% toward such doubling, evaluating the different anthropogenic greenhouse gases by their 100 year global warming potential CO_2 equivalents.

uncertainty.

Our model captures four different components of uncertainty that affect the optimal policy choice: climate sensitivity uncertainty, passive learning, active learning, and temperature stochasticity. First, incorporating uncertainty about climate sensitivity allows us to pin down a value of the optimal carbon tax, even if we are unable to pin down a value for climate sensitivity. The interaction between the economy and the climate system is non-linear and complex. A priori, it is unclear whether uncertainty implies a tax above or below the optimal policy that would prevail in a deterministic setting assuming a climate sensitivity of 3°C (in middle of the IPCC's likely parameter range). Second, we capture passive learning: the future decision maker has better information about the climate's response to emissions. This aspect of learning is frequently invoked in the policy debate to argue in favor of a wait and see policy. Third, we capture active learning: by increasing emissions, the policy maker can speed up the learning about climate sensitivity. Active learning always increases the emission level, but the magnitude of the policy effect is unclear. Fourth, temperature is stochastic over time. Temperature observations help our Bayesian decision maker to learn about climate sensitivity. The stochasticity of temperature determines the speed of learning. In addition, this temperature volatility affects economic damages and the optimal policies directly.

We solve the model numerically using projection methods to determine the value function and control rules on a six dimensional state space. We find that short-term temperature stochasticity has no impact on the optimal abatement, but triggers precautionary savings. In contrast, Bayesian uncertainty increases the optimal carbon tax by 25% and reduces emissions by 15%, both compared to the policy that assumes expected values for the climate sensitivity parameter. Second, we show that both passive and active learning play a very minor role in determining optimal present day policy: varying the speed of learning has a negligible impact on optimal present day emissions. The speed of learning is a major focus of earlier studies on climate sensitivity uncertainty.

Our study employs a standard Bayesian learning model: the decision maker has a subjective prior over climate sensitivity and the likelihood function governs stochastic temperatures. However, the scientific community does not even agree on a current subjective prior over climate sensitivity. The scientific working group of the IPCC does not publish an expected value of climate sensitivity (IPCC, 2013), and Meinshausen et al. (2009) reviews different probability distributions for climate sensitivity in the scientific literature, all of which differ in the expected values of the current prior. The decision-theoretic literature distinguishes between known risks or unique subjective beliefs, and situations of deep or hard uncertainty, usually referred to as situations of ambiguity. The behavioral literature finds that decision maker's behave differently in situations where they know the risk as opposed to situations whether they cannot judge the underlying probability distribution. A prominent model of ambiguity aversion that relates closely to the standard Bayesian model is the smooth ambiguity model by Klibanoff et al. (2009). Here, the decision maker evaluates the subjective prior with a higher uncertainty aversion than the risk aversion that applies to the evaluation of the objective risk of the likelihood function.

Acknowledging the low confidence of and disagreement on the climate sensitivity prior, we employ the smooth ambiguity model to analyze ambiguity and ambiguity aversion with respect to the ignorance over the climate sensitivity prior. Traeger (2011) provides a normative foundation of our approach. He shows that even in the von Neumann-Morgenstern framework, a fully rational decision-makers can have different degrees of aversion with respect to known probabilities and with respect to low confidence priors. We find that ambiguity aversion has a negligible effect on optimal policies. Thus, neither anticipated learning, nor ambiguity with respect to the true climate sensitivity should hold a policy maker from setting the optimal policy to the level optimal under our current subjective prior over climate sensitivity. This policy implies a significantly higher carbon tax than a deterministic climate change assessment relying on an expected value of climate sensitivity.

Relation to the Literature

Closest to our analysis is the seminal work by Kelly and Kolstad (1999), who investigate Bayesian learning about climate sensitivity in a similar model. While they analyze learning time in detail, they do not consider the separate contributions of uncertainty, learning and stochasticity on near term optimal policies. Analyzing initial beliefs, they find that a decision maker does not hedge against bad outcomes: an uncertain decision maker's choice of abatement is closer to the optimal policy under certain low climate sensitivity than certain high climate sensitivity. We find the opposite, possibly from our more symmetric comparison as well as our higher numerical precision.

Leach (2007) expands the work by Kelly and Kolstad (1999) by modeling a second climate parameter, the warming delay, as uncertain. He finds that modeling more than a single parameter as uncertain may practically prohibit learning. His findings hence complement our result that the speed of learning has very little influence on the currently optimal policies, much less than the level of stochasticity and the prior uncertainty. Leach (2007) also briefly considers the effect on optimal abatement in his setting, suggesting that a decision maker may lower abatement rates in order to speed up learning, which contrasts with our findings.

More recently, Kelly and Tan (2013) investigate uncertainty about climate feedbacks in a numeric integrated assessment framework. They focus on the impact of catastrophic damages resulting from a fat tailed probability distribution and highly convex damages. Their results suggest that the uncertainty effect is considerable in the first decade but wears off quickly as the probability mass in the tail shrinks (and the tail thin) and a catastrophic outcome becomes highly improbable.

Lemoine and Traeger (2013) model abrupt and irreversible changes in climate sensitivity resulting from the climate system crossing an unknown temperature threshold. The learning in their model consists of realizing that any temperature level reached without crossing the threshold is safe. Before and after crossing the threshold the climate sensitivity is known deterministically. In contrast, our decision maker learns the climate sensitivity smoothly over the course of decades and centuries. Lemoine and Traeger (2013) capture an extreme of sudden irreversible changes due to highly non-convex feedback processes. There, learning ahead of time is impossible. In contrast, we capture a world with smooth feedbacks and continuous learning.

Millner et al. (2013) relate to our analysis in that they model ambiguity aversion in the context of climate sensitivity. They assume that the decision maker has a prior over which climate model governs the warming of the world. The different possible models differ by climate sensitivity distributions. The authors find that ambiguity aversion has small welfare effects given the standard DICE damage function and large welfare effects when employing a more convex damage function. In contrast to Millner et al. (2013) we do not analyze the welfare effect of a given policy, but derive the optimal policy under uncertainty. Moreover, our decision maker behaves as a fully consistent Bayesian learner.

Another related strand of literature compares the performance of different climate policy instruments under various uncertainties (Hoel and Karp, 2001; Kelly, 2005; Karp and Zhang, 2006; Fischer and Springborn, 2011). Closest to our work are Karp and Zhang (2006) who compare taxes and quotas under learning about the relation between greenhouse gas concentrations and economic damages in a analytic integrated assessment model. They supplement their analytic results with a numerical application which suggests that anticipated learning reduces initial abatement levels considerably. In contrast, we find that learning does not affect near term abatement rates. The main differences compared to our model are their choice of linear-quadratic functional forms and the level of detail with which the climate and the economy are modeled (their model features three, ours six state variables).

2 Model

Our model is a recursive implementation of the integrated assessment model DICE (Nordhaus, 2008). It interacts a Ramsey growth economy with a simple climate model. We produce a single good by combining labor, capital and productivity in a Cobb-Douglas function. Production causes emissions which, if unabated, increase the stock of greenhouse gases in the atmosphere and, with a delay, the global surface temperature. The precise relation between greenhouse gas concentrations and temperatures depends on the climate sensitivity, which the social planner does not know. Temperatures above pre-industrial level (year 1900) hurt production. The social planner can reduce future damages by purchasing emission abatement. Alternatively, she can invest the produced good in

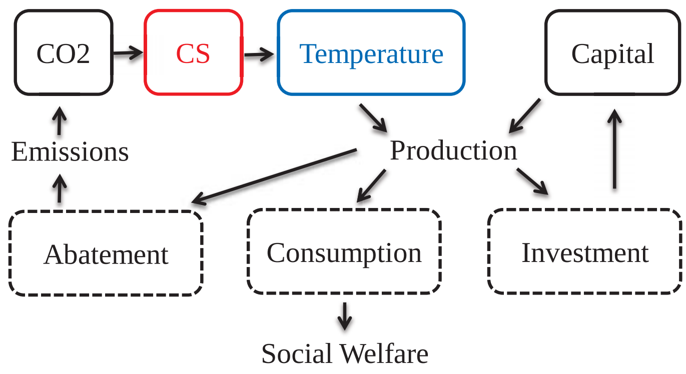

Figure 1: The main relations in the climate-enriched economy model. Control variables are represented by dashed rectangles. Main state variables are depicted by solid rectangles. Climate sensitivity ('CS') is uncertain. The decision maker has a prior over its value (2 state variables). Temperature is stochastic.

physical capital or allocate it to consumption. Appendix A contains a complete, formal description of our model, and Figure 1 depicts a stylized representation.

We formulate our model as a discrete time, infinite horizon dynamic programming problem. With this structure we can introduce stochastic temperature shocks in each period, model the social planner's prior over climate sensitivity by state variables, and introduce recursively defined ambiguity aversion preferences. We simplify the ocean feedbacks in DICE to reduce the computational costs of the recursive approach. Instead of the ocean carbon sink and the ocean temperature, we calibrate a decay rate for atmospheric carbon and a heat transfer between the atmosphere the ocean respectively. This adjustment saves us two computationally costly state variables. Our calibration matches the baseline DICE scenario under certainty closely. Please refer to Traeger (2012a), who derives the recursive implementation in detail.

2.1 Temperature stochasticity

For given climate sensitivity s, temperature T_t evolves according to

$$\tilde{T}_{t+1} = (1-\sigma)T_t + \sigma\,\mathrm{s}\left[\frac{\ln \frac{M_t}{M_{pre}}}{\ln 2} + \frac{EF_{t+1}}{\lambda}\right] - \sigma_{ocean}\Delta T_t + \tilde{\epsilon}_t\ . \tag{1}$$

Temperature in the next period T_{t+1} depends on the current temperature T_t, radiative forcing from atmospheric carbon stocks M_t (above pre-industrial level M_{pre}) as well as other, exogenous forcing EF_t, and heat exchange with the ocean ΔT_t. The parameters σ and σ_{ocean} capture delays, while s is the climate sensitivity[2]. Each year random events $\tilde{\epsilon}$ shock temperature. These "weather fluctuations" are normally distributed with mean zero. For a given value of climate sensitivity, the next period's temperature is then normally distributed

$$\tilde{T}_{t+1} \sim \mathcal{N}(\mu_{T,t+1}(s), \sigma_T^2) \qquad \text{with} \qquad \sigma_T^2 = 0.042,\ 0.2,\ 0.7\ .$$

The temperature mean from taking expectations in equation (1) is

$$\mu_{T,t+1} = s\ \chi_t(M_t, t) + \xi(T_t, t)$$

where

$$\begin{aligned} \chi_t(M_t, t) &= \sigma_{forc} \frac{\log \frac{M_{t+1}}{M_{pre}}}{\log 2} + \frac{EF_t + 1}{\eta_{forc}}\ , \quad \text{and} \\ \xi(T_t, t) &= (1 - \sigma_{forc}) T_t - \sigma_{ocean} \Delta T_t\ . \end{aligned}$$

The variance σ_T^2 is exogenous. Empirical estimates suggest annual volatility in global mean temperature in $\sigma_T^2 = 0.042$.[3] For our analysis we will also use considerably larger values. First, this estimate measures only global averages, whereas the within-country fluctuations are significantly larger, closer to our next higher value of $\sigma_T^2 = 0.2$. The effective damage increase of stochastic temperatures is captured better by a country's temperature volatility. Third, and that motivates our value of $\sigma_T^2 = 0.7$, stochasticity this high inhibits learning almost completely, which allows us to isolate the effect of uncertainty about climate sensitivity from the learning effect (see next section).

[2]Table B at the end of the paper contains all parameter definitions and values.

[3]Kelly and Kolstad (1999) and Leach (2007) both use $\sigma_T^2 = 0.1$. Averaging temperatures over 174 countries and estimating yearly fluctuations with respect to a common trend over 109 years results instead in the lower $\sigma_T^2 = 0.042$. We thank Christian Almer from the University of Bern for the estimates.

2.2 Uncertainty and Bayesian learning about climate sensitivity

The social planner is uncertain about the value of climate sensitivity and holds the following initial prior $\Pi(s)$

$$\tilde{s}_0 \sim \Pi(s) = \mathcal{N}(\mu_{s,0}, \sigma^2_{s,0}) \qquad \text{with} \qquad \mu_{s,0} = 3 \qquad \sigma^2_{s,0} = 1,\ 2,\ 3\ .$$

Most commonly, estimates of climate sensitivity take fat-tailed distributional forms such as the log-normal. To simplify the characterization of learning, we assume a normal distribution. Given this limitation, $\sigma^2_{s,0} = 3$ is a rounded-up empirical approximation to the set of distributions found in IPCC (2013)[4] To analyze the dynamics of learning, we vary the prior variance $\sigma^2_{s,0} = 1,\ 2,\ 3$.

We can learn the value of climate sensitivity from observing the CO_2 stock and temperatures over time. All the feedbacks that are not part of the climate sensitivity are known. Every period the decision maker foresees what a future realization of the temperature teaches her about climate sensitivity distribution and updates her prior accordingly.

Her posterior in period t is the prior conditional on historic temperature realizations $\Pi(s|\hat{T}_1, ..., \hat{T}_t)$. This posterior also depends on the historic CO_2 stock information which we suppress for notational convenience. Given the current stock M_t, a realization of temperature $\hat{T}_{t+1}$ in the subsequent period results in the updated posterior $\Pi(s|\hat{T}_1, ..., \hat{T}_{t+1})$. In Appendix B we show that the updated posteriors are again normally distributed so that at all times $\Pi(s|\hat{T}_1, ..., \hat{T}_t) = \mathcal{N}(\mu_{s,t}, \sigma^2_{s,t})$ for some $\mu_{s,t}$ and $\sigma^2_{s,t}$. Moreover, we prove the following updating rules for the expected value

$$\mu_{s,t+1} = \frac{\chi_t^2 \sigma^2_{s,t} \frac{\hat{T}_{t+1} - \xi_t}{\chi_t} + \sigma^2_T \mu_{s,t}}{\chi_t^2 \sigma^2_{s,t} + \sigma^2_T}$$

[4]IPCC (2013) considers climate sensitivity values between 1.5 to 4.5 degrees Celsius as likely and values below 1 degree and above 6 degrees Celsius as extremely and very unlikely respectively. The AR5 does not provide a best estimate due to lack of scientific consensus. A normal distribution with mean 3 and variance 3 has the one standard deviation bands [1.27,4.73].

and the variance

$$\sigma_{s,t+1}^2 = \frac{\sigma_T^2 \sigma_{s,t}^2}{\chi_t^2 \sigma_{s,t}^2 + \sigma_T^2} \ . \tag{2}$$

The new expected value of the parameter s is a weighted mean of the previous expected value and the inferred "climate sensitivity observation", $\frac{\hat{T}_{t+1}-\xi_t}{\chi_t}$. The weight on the new observation is proportional to the precision (the inverse of the variance) of the temperature and the magnitude of the multiplicative factor χ_t, which increases in the carbon stock. The decision maker learns faster the lower the temperature stochasticity and the larger the carbon stock. This insight follows from observing that the first summand in the bracket in equation (2) grows in $1/\sigma_T^2$ and in χ_t. With uncertainty about climate sensitivity, temperature realizations are themselves uncertain (not only stochastic) and governed by the predictive distribution $\tilde{T}_{t+1} \sim \mathcal{N}(\xi_t + \chi_t \mu_{s,t}, \chi_t^2 \sigma_{s,t}^2 + \sigma_T^2)$. We can conveniently use this distribution (derived in Appendix B) to evaluate the uncertainty in the optimization.

2.3 Welfare specification and Bellman equation

The social planner maximizes her value function subject to the exogenous and endogenous equations of motion for the economy and the climate. The physical state variables describing the system are capital k, CO_2 stock M, and temperature T, and the informational state variables for the climate sensitivity prior $\Pi(s)$ are $\mu_{s,t}$ and $\sigma_{s,t}^2$. Time t captures all exogenously evolving processes, such as population and technology growth and temperature feedbacks. For numerical reasons, we express capital (and production and consumption) in effective labor units, i.e. $k_t = K_t/A_t L_t$ where A_t is technology level and L_t population at time t.[5] The Bellman equation reads

$$V(k_t, M_t, t, T_t, \mu_{s,t}, \sigma_{s,t}^2) = \max_{c_t, \mu_t} \frac{(c_t)^{1-\hat{\eta}}}{1-\hat{\eta}} \tag{3}$$
$$+\beta_t \, \mathrm{E}\left[V(k_{t+1}, M_{t+1}, t+1, \tilde{T}_{t+1}, \tilde{\mu}_{s,t+1}, \tilde{\sigma}_{s,t+1}^2)\right]$$

[5] This variable change avoids unbounded growth of the state space in the capital dimension, see Traeger (2012a) for the transformation of the Bellman equation.

For any given state of the system, maximum obtainable welfare today V_t is the sum of instantaneous welfare from current consumption c_t and expected discounted maximized future welfare V_{t+1}. The social planner has constant relative risk aversion (CRRA) preferences with $\eta = 2$ (Nordhaus, 2008). This instantaneous utility function describes her risk aversion as well as her desire to smooth consumption over time.[6] The discount factor β_t contains a pure rate of time preference of 1.5 percent.[7], and the decision maker takes expectations over the uncertain climate sensitivity prior and temperatures using the predictive distribution. By choosing the consumption level c_t, she balances immediate consumption against future physical capital stocks, while her abatement decision μ_t trades off immediate consumption and lower future carbon concentration.

The social cost of carbon is the discounted sum of future welfare costs caused by the marginal emission unit. We can recover the optimal social cost of carbon from the value function as the ratio of the marginal value of a ton of carbon and the marginal value of a unit of the consumption good

$$SCC_t = \frac{\partial_{M_t} V(\cdot)}{\partial_{k_t} V(\cdot)} A_t L_t \ .$$

Another useful welfare measure is the so called balanced growth equivalent. It measures welfare effects of a set of optimal abatement and consumption policies (Mirrlees and Stern, 1972; Anthoff and Tol, 2009). It is the per capita consumption $\bar{c}$ that, growing at some fixed rate g, would yield the same welfare as the (optimal) policy $\mathcal{A}$

$$\bar{c}^{\mathcal{A}}(\cdot) = \left[\frac{(1-\eta) V^{\mathcal{A}}(\cdot)}{\frac{L_\infty}{1-\exp[(1-\eta)g-\delta_u]} - \frac{L_\infty - L_0}{1-\exp[(1-\eta)g-\delta_u-g_L^*]}} \right]^{\frac{1}{1-\eta}} \ .$$

With the balanced growth equivalent we can conveniently compare the policies under two alternative scenarios $\mathcal{A}$ and $\mathcal{B}$ by the percentage difference in their respective consumption paths

$$\Delta^{\mathcal{AB}} \bar{c}(\cdot) = \frac{\bar{c}^{\mathcal{A}} - \bar{c}^{\mathcal{B}}}{\bar{c}^{\mathcal{A}}} = 1 - \left[\frac{V^{\mathcal{B}}(\cdot)}{V^{\mathcal{A}}(\cdot)} \right]^{\frac{1}{1-\eta}} \ .$$

[6] Those two preferences are a priori unrelated and could be disentangled, see Traeger (2012b).

[7] Due to our transformation it in addition comprises labor and technological growth rates.

2.4 Numerical implementation

We solve the dynamic programming equation (3) by function iteration, using the collocation method to approximate the value function. As basis functions we choose Chebychev polynomials with $22,400$ Chebychev nodes and coefficients[8]. The normal distributions for temperature stochasticity and the climate sensitivity prior are approximated by Gauss-Legendre quadrature with 3 nodes[9]. Our convergence criterion is a change in the value function coefficients of less than 10^{-4}. The code is written in Matlab, we use the CompEcon toolbox by Miranda and Fackler (2002) to generate and evaluate the Chebychev polynomials, and let the solver KNITRO to carry out the optimization.

3 Results

In this section we present the results for three different scenarios that build upon each other: Pure temperature stochasticity, climate sensitivity uncertainty and learning. We introduce and discuss ambiguity aversion separately in Section 4.

3.1 Temperature stochasticity

Figure 2 presents our results for pure temperature stochasticity, and Table 1 contains and compares our main findings in this paper. In this scenario, the decision maker knows the climate sensitivity. The four panels show the abatement rate, the social cost of carbon, the investment rate and emissions over the current century. We distinguish four scenarios: deterministic temperature ('certainty') and three levels of stochasticity ($\sigma_T^2 = 0.042, 0.2, 0.7$). To generate the stochastic scenario paths, we draw the expected value of temperature in each period, such that each period the shock is zero. This 'expected path' procedure insures that for a given set of abatement and investment policies, temperatures coincide under certainty and stochasticity, and the difference we initially observe is a direct result of the optimal reaction to stochasticity.

[8]Along each dimension of the state space we use the following node numbers: $k = 7, M = 4, t = 8, T = 4, c = 5, s = 5$. The results are robust to increasing time nodes to 12 and

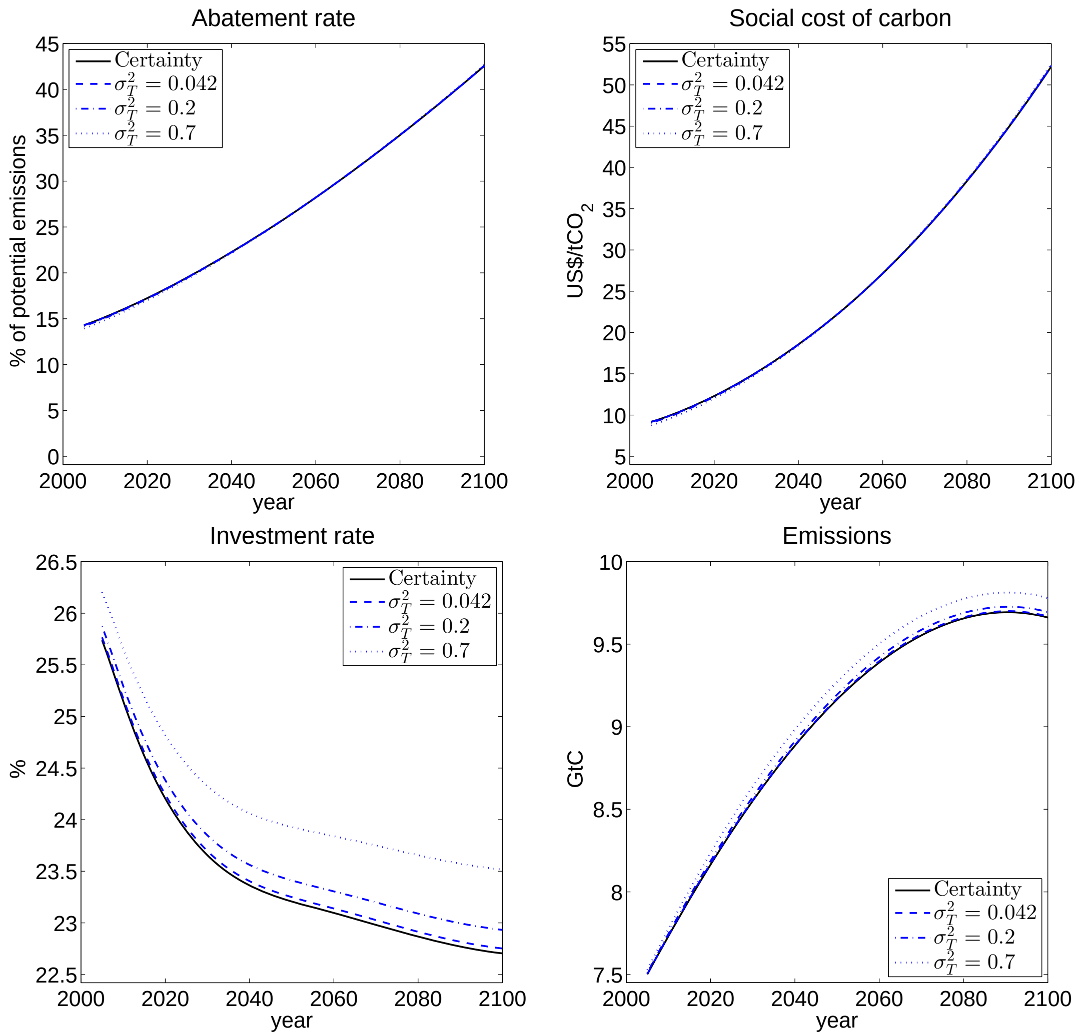

Figure 2: Optimal abatement rate, social cost of carbon, emissions and investment rate for the current century. Temperature is either deterministic or stochastic with one of the three variances $\sigma_T^2 = 0.042$, $\sigma_T^2 = 0.2$ and $\sigma_T^2 = 0.7$.

We find that temperature stochasticity has no economically significant effect on the optimal abatement policy and the associated social cost of carbon. Individual shocks have no direct lasting impact on the climate system, so the decision maker sees no need to accomodate them by adjusting abatement. Investment in manmade capital however slightly increases: we observe a modest precautionary

temperature nodes to 8.

[9]Results are unaffected by increasing the number of Gauss-Legendre nodes to 15.

	Abatement		Tax		Investment		Emissions	
	Rate	*Change*	*US$*	*Change*	*Rate*	*Change*	*GtC*	*Change*
Certainty	16		10.9		24.7		7.91	
Stochasticity ($\sigma_T^2 = .042$)	16	0 %	10.9	0 %	24.8	0 %	7.91	0 %
Stochasticity, medium ($\sigma_T^2 = .2$)	15.9	-0.6 %	10.8	-1 %	24.9	1 %	7.93	0 %
Stochasticity, high ($\sigma_T^2 = .7$)	15.7	-1.9 %	10.5	-3.7 %	25.3	2.4 %	7.97	0.8 %
Learning ($\sigma_{s,0}^2 = 3, \sigma_T^2 = .042$)	17.8	**11.3** %	13.2	**21.1** %	24.7	0 %	7.76	**-1.9** %
Slow Learning ($\sigma_{s,0}^2 = 3, \sigma_T^2 = .2$)	18	12.5 %	13.5	23.9 %	24.8	**0.4** %	7.75	-2.0 %
Ambiguity, RRA=10 ($\sigma_{s,0}^2 = 3, \sigma_T^2 = .042$)	17.8	11.3 %	13.3	22.0 %	24.5	-0.8 %	7.76	**-1.9** %
Ambiguity, RRA=50 ($\sigma_{s,0}^2 = 3, \sigma_T^2 = .042$)	18	12.5 %	13.5	23.9 %	24.5	-0.8 %	7.74	-2.1 %

Table 1: Optimal 2014 values of abatement (% of business as usual emissions), carbon tax (in US$ per ton of CO_2), investment rate (in % of production) and emissions (in GtC). The second value in each column shows the percentage change relative to the certainty scenario. Yellow highlighting indicates main effects.

savings effect. All else equal, a high temperature realization causes high damages for one period. Production falls, and investment (in absolute terms) is lower. The single shock is propagated via the capital stock and remains in the economy for multiple time periods. To insure against this expected welfare loss, the decision maker invests more in manmade capital at any given time. The higher level of investment leads to a higher capital stock which eventually increases total emissions. In the present setting, temperature stochasticity alone does not influence the optimal social cost of carbon. Of course, this result crucially depends on the absence of non-linear, self-enforcing feedbacks (the melting of the Antarctic ice-sheet, or methane release from thawing permafrost, for example).

3.2 Climate sensitivity uncertainty

Figure 3 shows the evolution of the climate sensitivity prior variance for the current century for two different initial priors and three different values of temperature stochasticity. Temperature is realized at its expected value, therefore the climate sensitivity prior mean remains unchanged at $\mu_{s,0} = \mu_{s,t} = 3$. The

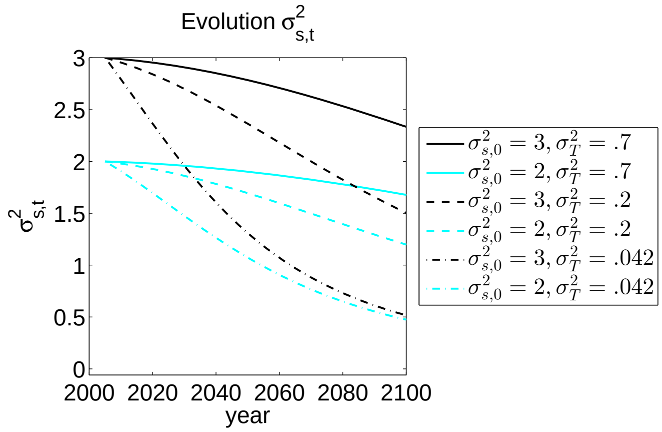

Figure 3: shows how the climate sensitivity prior variance $\sigma^2_{s,t}$ evolves over time for two initial values ($\sigma^2_{s,0} = 2$ and $\sigma^2_{s,0} = 3$) and three levels of temperature stochasticity ($\sigma^2_T = 0.042, 0.2, 0.7$). The mean $\mu_{s,t}$ remains constant at 3 as observations confirm this value in each period.

expectations of the decision maker are confirmed with every single observation, yet for $\sigma^2_T = 0.7$ her confidence in her prior does increase only slightly over the first 100 years when temperature stochasticity is high. Therefore we interpret this case as representing uncertainty rather than learning. Only with lower levels of temperature volatility, in particular $\sigma^2_T = 0.042$, meaningful learning takes place.

Figure 4 shows the same set of graphs as Figure 2. Temperature is stochastic with $\sigma^2_T = 0.7$, whereas climate sensitivity is either known with certainty or subjectively uncertain with a prior mean of $\mu_{s,0} = 3$. We distinguish two possible prior variances, $\sigma^2_{s,0} = 2$ and $\sigma^2_{s,0} = 3$. Again we plot the paths along the expected values for temperature stochasticity. The decision maker's climate sensitivity prior is unbiased, so her expectation coincides with the true value.

Subjective uncertainty about the value of climate sensitivity modestly raises the abatement rate. and the social cost of carbon. When the climate sensitivity is known, we optimally abate 15.5% of emissions in 2014. With climate sensitivity uncertainty and an initial prior of $\mathcal{N}(3, 3)$, abatement in 2014 is 17.7%, or about 15% higher for the uncertainty case. The corresponding social cost of carbon is almost 30% higher.

The decision maker acts with precaution in the face of possibly very different realities: In comparison to stochasticity, subjective uncertainty means the "re-

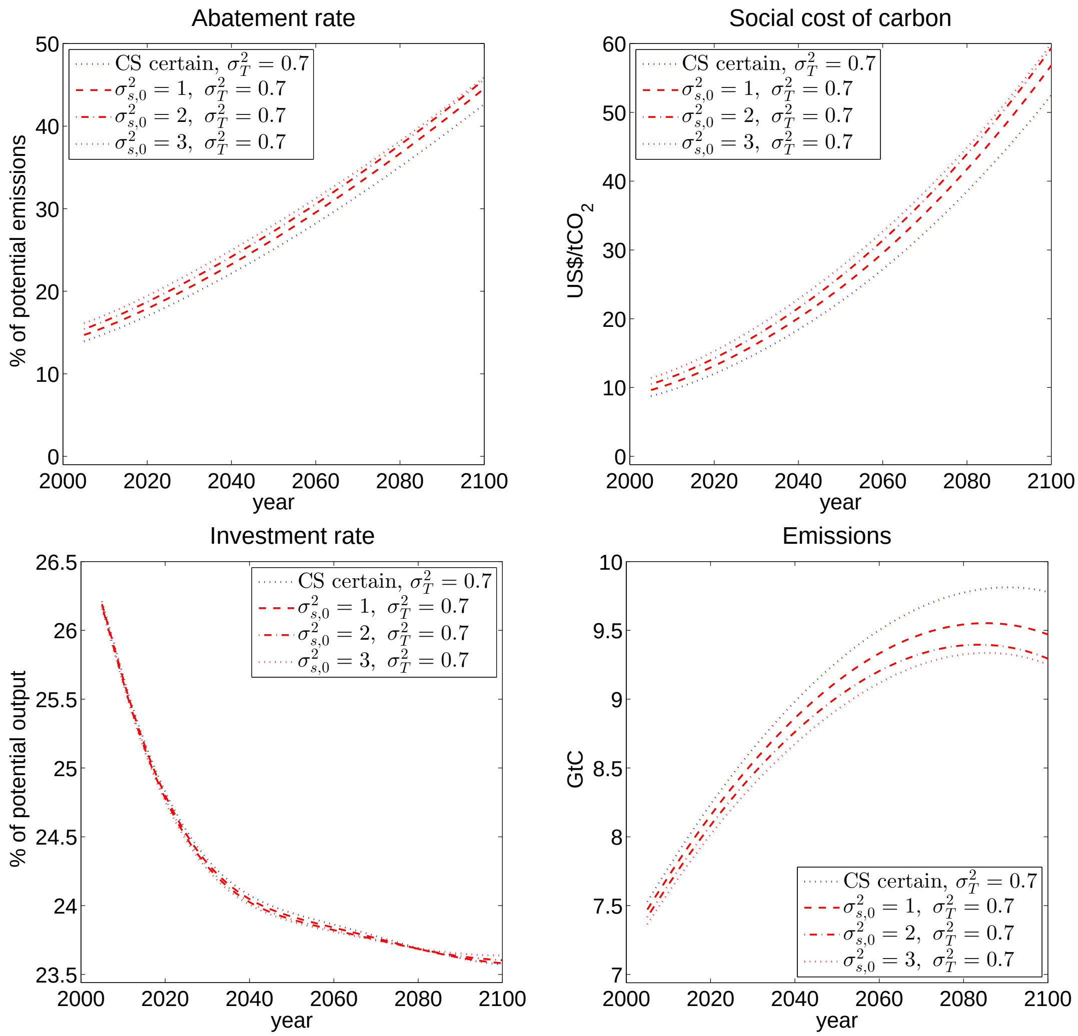

Figure 4: Optimal abatement rate, social cost of carbon, emissions and investment rate for the current century with stochastic temperature ($\sigma_T^2 = .7$), and known and uncertain climate sensitivity. The unbiased prior mean is $\mu_{s,0} = 3$. Initial prior variances are $\sigma_{s,0}^2 = 1,\ 2\ 3$.

alized shock" lasts as long as the carbon stock, i.e. for several centuries. The temperature stochasticity $\sigma_T^2 = 0.7$ is so high that a single temperature observation recieves very little weight when the decision maker updates her prior, and the decision maker anticipates that she will learn very little. Unlike stochastic temperature, subjective uncertainty does not affect the investment rate, such that higher abatement rates translate without moderation into lower emissions. Short term fluctuations in temperature ("weather") and long term uncertainty about the severity of climate change hence require different optimal policy responses: Temperature stochasticity leads to precautionary savings in capital which over time increase emissions, whereas we meet climate sensitvity uncertainty optimally by increasing abatement, directly reducing emissions.

3.3 Learning about climate sensitivity

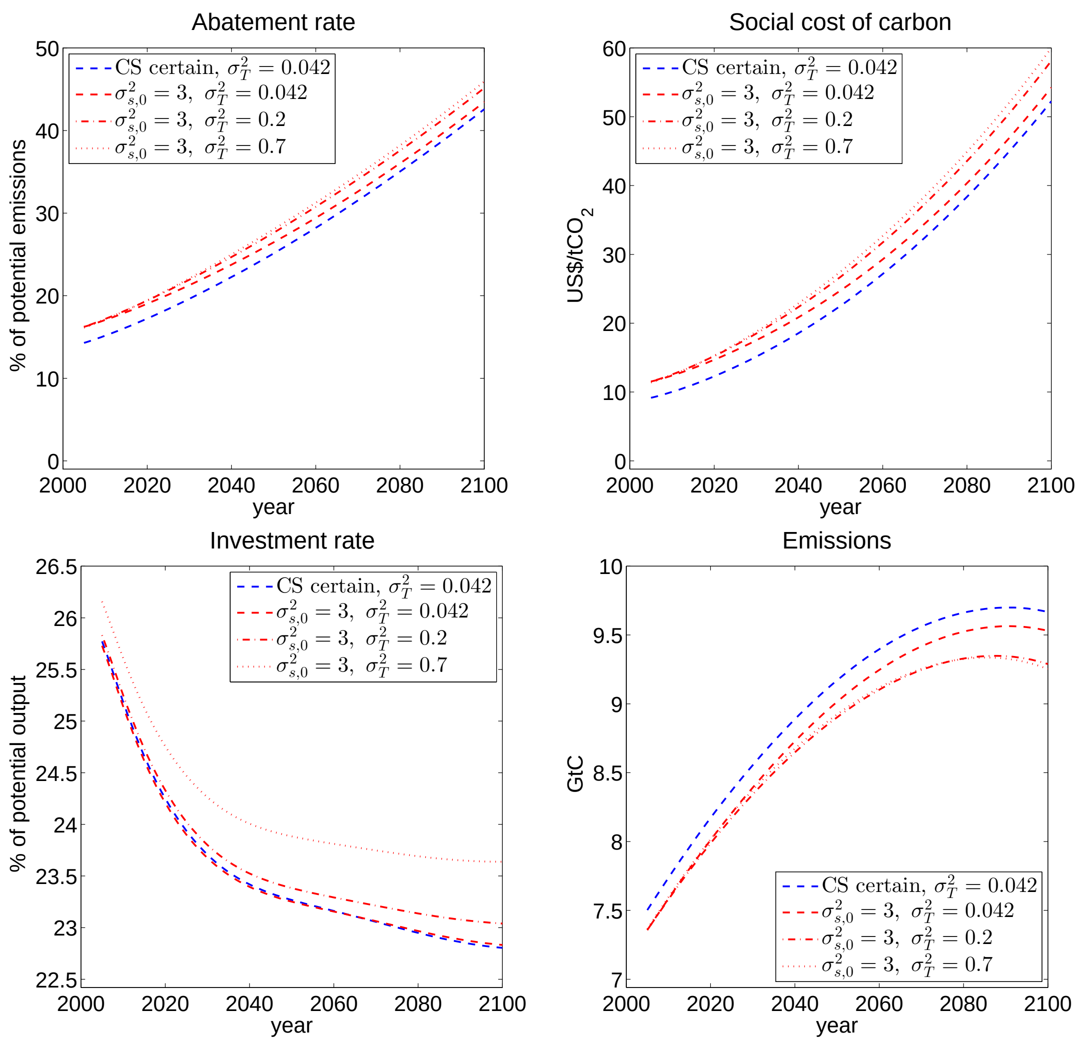

Figure 5: Optimal abatement rate, social cost of carbon, investment rate and emissions for the current century. Climate sensitivity is either certain or uncertain. In the latter case the decision maker holds an unbiased prior with a initial variance of $\sigma_{s,0}^2 = 3$. We show three different temperature stochasticities: $\sigma_T^2 = 0.042,\ 0.2,\ 0.7$.

How does learning at different speeds translate optimal abatement policies? We compare learning for the three different values of temperature stochasticity in Figure 5. Again we display the abatement rate, the social cost of carbon, investment and emissions. The paths for high temperature stochasticity ($\sigma_T^2 = .7$) and known climate sensitivity ('CS certain') are the same as in Figure 4. The new abatement paths with lower temperature volatilities start out at the same level as their high volatility counterpart. The identical initial levels show that the anticipation of learning and the speed of learning do not affect optimal current climate polices. Abatement efforts are not lowered because it is optimal to "wait and see" (passive learning) or because higher carbon stocks speed up the learning process (active learning).[10] Present day climate policies are determined by the uncertainty in climate sensitivity alone. The optimal carbon tax is 21.1 percent higher with a temperature volatility of $\sigma_T^2 = .42$. Table 1 shows the optimal levels of abatement, the optimal social cost of carbon, investment and emissions for the year 2014. Only after considerable time the difference in confidence in the prior leads to the different paths approaching the level of pure temperature stochasticity at different speeds. We also observe an interesting interaction between climate and economy for emissions: The emission paths for $\sigma_T^2 = .2$ and $\sigma_T^2 = .7$ cross. Investment increases permanently, as the temperature stochasticity is irreducible. For high stochasticity $\sigma_T^2 = .7$ also the impact from subjective uncertainty last for the entire century. For $\sigma_T^2 = .2$ on the contrary, the increase in abatement wears off as the decision maker becomes more confident over time. Hence emissions increase faster, and eventually overtake emissions for the high stochasticity scenario.

Another important aspect of the learning dynamics is the correction of wrong expectations. In Figure 6 we show how a decision maker with a wrong initial prior adjusts abatement, and how the mean of her prior evolves. Here we use the low, empirically accurate temperature volatility of $\sigma_T^2 = 0.042$. Correcting the wrong belief takes long, even with low temperature volatility. Secondly, the decision maker insures herself against a "too low" expected climate sensitivity: The initial abatement rate under uncertainty is biased towards the optimal policy

[10] This contradicts the result by Leach (2007) in a setting with two uncertain parameters.

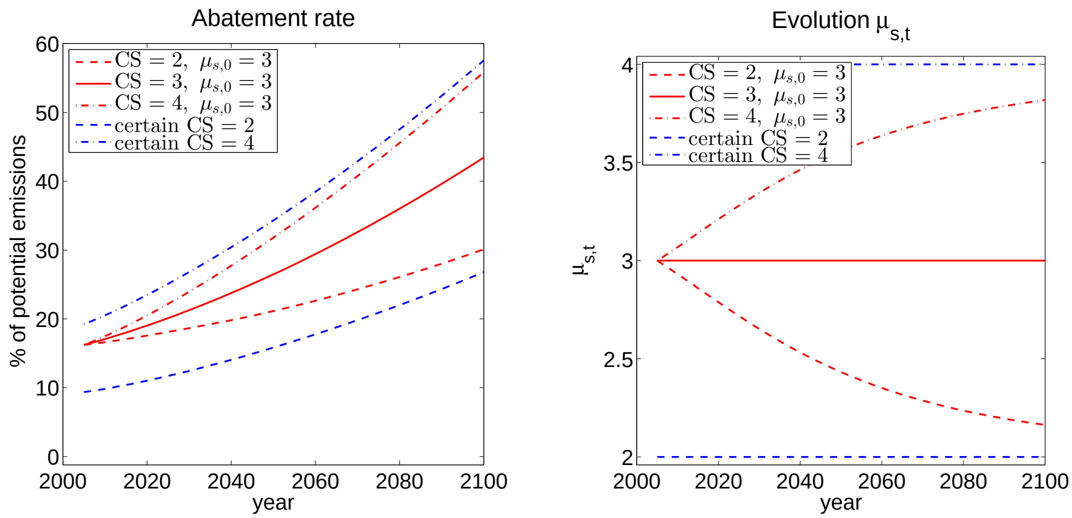

Figure 6: Abatement rate and dynamics of climate sensitivity prior mean ($\mu_{s,0} = 3$) for varying levels of true climate sensitivity (CS$= 2,\ 3,\ 4$), a prior variance of $\sigma^2_{s,0} = 3$ and temperature stochasticity $\sigma^2_T = 0.042$.

under high, certain climate sensitivity.[11]

4 Ambiguity aversion

Ambiguity aversion captures the attitude of a decision maker who prefers a world with well known probabilities to a world governed by subjective uncertainty. We model those two preferences by two different aggregator functions (Klibanoff et al., 2009). The social planner evaluates risk with her standard constant relative risk aversion (CRRA) preferences. The second aggregator function $f(z) = [(1-\eta)z]^{\frac{1-RAA}{1-\eta}}$ characterizes her additional aversion to subjective risk.[12] Given those

[11] This result contrasts with Kelly and Kolstad (1999), who however note that they face numerical difficulties.

[12] RAA stands for: Constant coefficient of **R**elative **A**mbiguity **A**version. Traeger (2012b) defines the measure analogously to Arrow-Pratt relative risk aversion.

preferences, the Bellman equation now reads

$$V(k_t, M_t, t, T_t, \mu_{s,t}, \sigma_{s,t}) = \max_{c_t,\mu_t} \frac{c_t^{1-\eta}}{1-\eta} + \frac{\beta_t}{1-\eta} \times$$

$$\left\{ \int_\Theta \left((1-\eta)\, \mathrm{E}_{\psi(s)} \left[V(k_{t+1}, M_{t+1}, t+1, \tilde{T}_{t+1}, \mu_{s,t+1}, \sigma_{s,t+1}) \right] \right)^{\frac{1-RAA}{1-\eta}} \mathrm{d}\,\Pi(s) \right\}^{\frac{1-\eta}{1-RAA}} .$$

For a particular realization of climate sensitivity, temperature is stochastic and

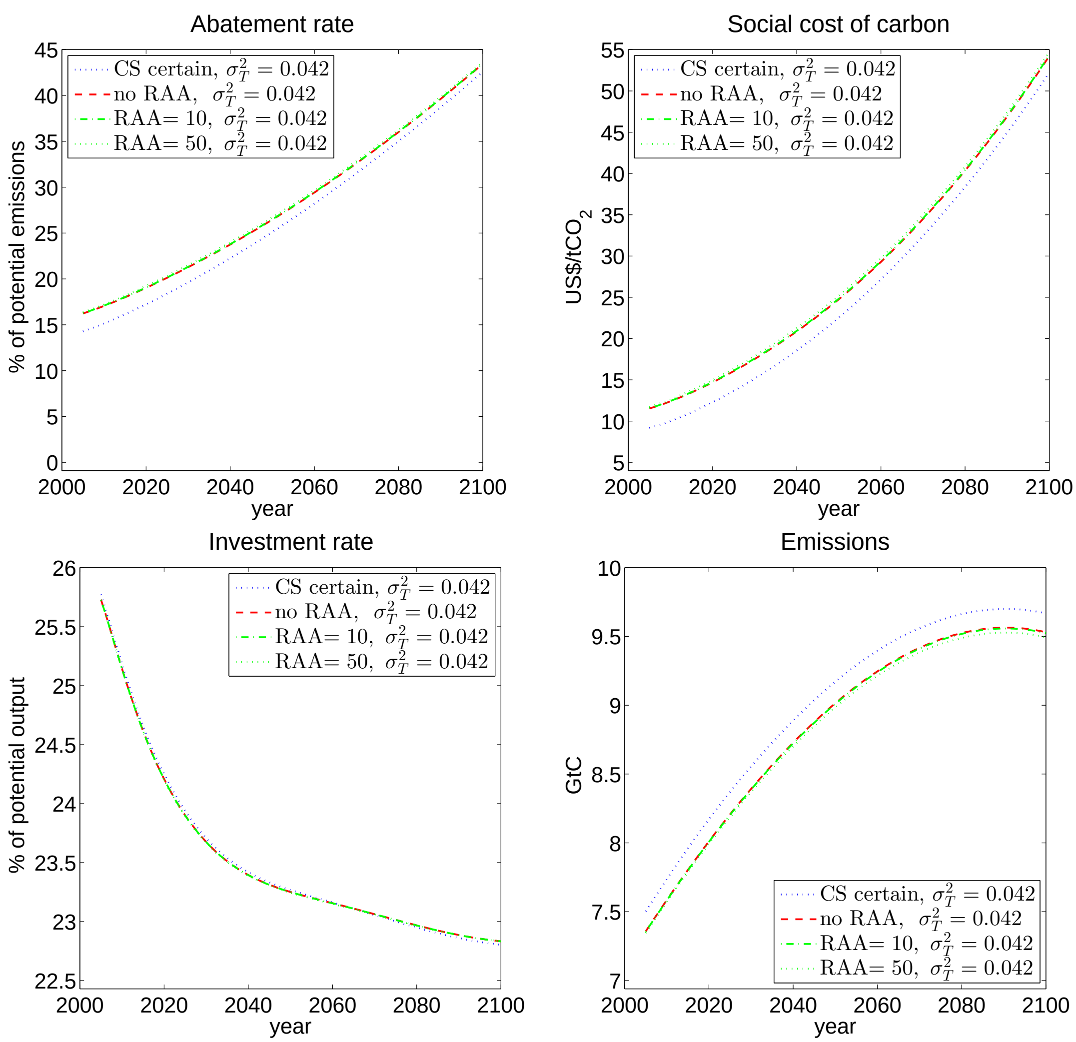

Figure 7: Abatement rate, social cost of carbon, investment rate and emissions for the current century with stochastic temperature ($\sigma_T^2 = .042$), uncertain climate sensitivity with initial prior variance $\sigma_{s,0}^2 = 3$ and three levels of ambiguity aversion: none, $RAA =$ 10 and $RAA = 100$.

normally distributed, $\mathcal{N}(\mu_{T,t}(s), \sigma_T^2)$. The expectation operator in the inner bracket takes expected future welfare with respect to this well-known stochasticity. In addition, the decision maker is subjectively uncertain about climate sensitivity over which she has the prior $\Pi(s) \sim \mathcal{N}(\mu_{s,t}, \sigma_{s,t}^2)$. The integral with respect to the prior Π expresses this second uncertainty integration. The ambiguity aversion function $f(z) = [(1-\eta)z]^{\frac{1-RAA}{1-\eta}}$ curves the argument of this second uncertainty aggregation additionally, expressing additional aversion because of the low confidence over the prior. Observe that for $RAA = \eta$ the additional aversion vanishes and the Bellman equation collapses to its standard form.

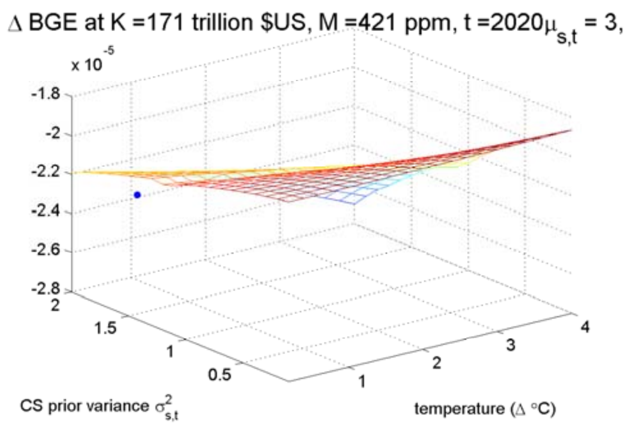

Figure 8: Difference in balanced growth equivalent $\frac{\bar{c}^{RAA10} - \bar{c}^{learn}}{\bar{c}^{RAA10}}$ between expected utility maximizer and ambiguity averse decision maker with $RAA = 10$. Plotted over climate sensitivity prior variance and temperature for the year 2020, a carbon stock of 896 GtC (421 ppm CO_2), a capital stock of 171 US trillion dollars and a climate sensitivity prior mean of $\mu_{s,0} = 3$.

In the case of subjectively uncertain climate sensitivity the decision-maker's ability to change or avoid subjective uncertainty is limited. She can increase her emissions in order to learn faster. However, the learning comes at the cost of being even worse-off in the situation where climate sensitivity turns out to be high. We find that overall ambiguity aversion has virtually no effect on optimal policies (Figure 7, Table 1). The ambiguity averse social planner acts identically to one who evaluates risk and subjective uncertainty equally. Also, no loss in welfare is experienced even with strong aversion to subjectivity. Figure 8 compares the balanced growth equivalent for an ambiguity averse decision maker to a standard expected utility maximizer. The percentage difference in per capita consumption that makes them equally well off is in the order of magnitude of 10^{-5}, or numerically zero.

5 Conclusions

The relation between greenhouse gas emissions and global equilibrium temperature is governed by major scientific uncertainty. We derive the optimal mitigation policy in the face of climatic uncertainty, and compare it to the deterministic policy assuming the expected climate response. We thereby disentangle effects of Bayesian uncertainty, learning, and stochasticity. Uncertainty about climate sensitivity increases the optimal abatement rate in the present by 15%, and the optimal carbon tax by approximately 25%. The optimal investments in physical capital remain constant. Over time, learning reduces the uncertainty, but the present day policy is not affected by anticipating this learning process or by the speed of learning, a major focus of earlier analyses.

In particular, passive learning does not imply a wait and see policy, and also active learning has no notable effect on policies: the decision maker should not optimally raise current emissions to learn faster about the climate response. If the true climate sensitivity parameter happens to coincide with the expected value of the climate sensitivity prior, then optimal abatement slowly converges back to the deterministic scenario.

The stochasticity of temperature not only determines the speed of learning, but also directly affects the expected damages. We find that the implied damage stochasticity has no effect on optimal abatement, but induces precautionary savings: the decision maker should slightly increase the wealth level to reduce her vulnerability to the stochastic temperature shocks. If global risk sharing occurs, the precautionary savings effect is minimal. If such risk sharing is absent, then precautionary savings can play a slightly larger role. Higher temperature stochasticity also reduces the speed of learning. Keeping expectations fixed, this second effect of temperature stochasticity increases optimal mitigation over the coming decades. The joint impact of uncertainty, stochasticity, and learning always increases the abatement rate and decreases absolute emission levels.

The scientific literature disagrees on the climate sensitivity prior and its expected value. We choose a normal distribution, which is particularly convenient for the dynamic analysis. While other forms likely resemble current estimates more closely, any given prior and expectation lacks confidence by a major part of the scientific community. We therefore acknowledge the low confidence of the Bayesian prior by making it ambiguous. We show that ambiguity about the Bayesian prior in combination with ambiguity aversion, i.e., aversion to the lack of confidence, has virtually no effect on the optimal climate policy.

The currently prevailing uncertainty with respect to the climate's response to emission increases the optimal carbon tax about 25%. Neither the anticipation of learning, nor the lack of confidence into the climate sensitivity prior change this result. Temperature stochasticity has no effect on the present abatement policy, but implies that a small amount of precautionary savings can increase welfare.

A Details on the climate enriched economy model

The following model emulates DICE-2007. The three most notable differences are the annual time step (DICE-2007 features ten year time periods), the infinite time horizon, and the replacement of the ocean feedbacks by exogenous processes. This simplification is neccessary because the ocean carbon sink and ocean temperature would each require an own state variable in a recursive framework, which is computationally too costly. Instead we calibrate a decay rate for atmospheric carbon and a temperature difference between atmosphere and ocean which closely match the behavior of DICE's original carbon cycle. For a detailed description of the procedure, see Traeger (2012a), who also shows how to reformulate the decision problem when expressing capital stock and consumption in efficient labor units. All parameters are characterized and quantified in Table B on page 31.

Global average temperatures respond with a delay to the forcing from atmospheric carbon stocks M_t (above preindustrial level M_{pre}) and other non-CO2 forcing. Restating Equation (1) with climate sensitivity as a uncertain parameter

$$\tilde{T}_{t+1} = (1-\sigma)T_t + \sigma \tilde{\mathrm{s}} \left[\frac{\ln \frac{M_t}{M_{pre}}}{\ln 2} + \frac{EF_{t+1}}{\lambda} \right] - \sigma_{ocean} \Delta T_t + \tilde{\epsilon}_t \ .$$

The ocean temperature difference ΔT_t replicates the relation between oceanic and atmospheric temperatures in DICE. It follows the simple quadratic equation

$$\Delta T_t = \max\{0.7 + 0.02. * t - 0.00007. * t.^2 \ , \ 0\} \ .$$

Exogenous forcing EF_t from non-CO2 greenhouse gases, aerosols and other processes is assumed to follow the process

$$EF_t = EF_0 + 0.01(EF_{100} - EF_0) \times \max\{t, 100\} \ .$$

Note that it starts out slightly negatively. Carbon in the atmosphere accumulates according to

$$M_{t+1} = M_{pre} + (M_t - M_{pre})\,(1 - \delta_M(t)) + E_t \quad \text{with}$$
$$\delta_{M,t} = \delta_{M,\infty} + (\delta_{M,0} - \delta_{M,\infty}) \exp[-\delta_M^* \ t] \ .$$

The stock of CO_2 (M_t) exceeding preindustrial levels (M_{pre}) decays exponentially at the rate $\delta_M(M,t)$. This decay rate falls exogenously over time to replicate the carbon cycle in DICE-2007, mimicking that the ocean reservoirs reduce their uptake rate as they fill up (see Traeger, 2012a). The variable E_t characterizes yearly CO_2 emissions, consisting of industrial emissions and emissions from land use change an forestry B_t

$$E_t = (1-\mu_t)\,\sigma_t A_t L_t k_t^\kappa + B_t \ .$$

Emissions from land use change and forestry fall exponentially over time

$$B_t = B_0 \exp[g_B\ t] \ .$$

Industrial emissions are proportional to gross production $A_t L_t k_t^\kappa$. They can be reduced by abatement μ_t. As in the DICE model, the carbon intensity of production falls at an exogenous rate of decarbonization σ_t

$$\sigma_t = \sigma_{t-1} \exp[g_{\sigma,t}] \qquad \text{with} \qquad g_{\sigma,t} = g_{\sigma,0} \exp[-\delta_\sigma\ t] \ .$$

The economy accumulates capital according to

$$k_{t+1} = [(1-\delta_k)\,k_t + y_t - c_t] \exp[-(g_{A,t} + g_{L,t})] \ ,$$

where δ_K denotes the depreciation rate, y_t denotes production net of abatement costs and climate damage, and c_t denotes aggregate global consumption of produced commodities (both in per effective labor units, i.e. $y_y = \frac{Y_t}{A_t L_t}$). Population grows exogenously

$$L_{t+1} = \exp[g_{L,t}] L_t \quad \text{with} \quad g_{L,t} = \frac{g_L^*}{\frac{L_\infty}{L_\infty - L_0} \exp[g_L^*\ t] - 1} \ .$$

Here L_0 denotes the initial and L_∞ the asymptotic population. The parameter g_L^* characterizes the convergence from initial to asymptotic population. Technology grows exogenously

$$A_{t+1} = A_t \exp\left[g_{A,t}\right] \quad \text{with} \quad g_{A,t} = g_{A,0} * \exp\left[-\delta_A t\right] \ .$$

Net global GDP is obtained from the gross product as follows

$$y_t = \frac{1 - \Lambda(\mu_t)}{1 + D(T_t)} k_t^\kappa$$

where production is expressed in per effective labor units and

$$\Lambda(\mu_t) = \Psi_t \mu_t^{a_2}$$

characterizes abatement costs as percent of GDP depending on the emission control rate $\mu_t \in [0, 1]$. The coefficient of the abatement cost function Ψ_t follows

$$\Psi_t = \frac{\sigma_t}{a_2} a_0 \left(1 - \frac{(1 - \exp[g_\Psi\, t])}{a_1}\right)$$

with a_0 denoting the initial cost of the backstop, a_1 denoting the ratio of initial over final backstop, and a_2 denoting the cost exponent. The rate g_Ψ describes the convergence from the initial to the final cost of the backstop.

Climate damage as percent of world GDP depends on the temperature difference T_t of current to preindustrial temperatures and is characterized by

$$D(T_t) = b_1 T_t^{b_2} \ .$$

Nordhaus (2008) estimates $b_1 = 0.0028$ and $b_2 = 2$, implying a quadratic damage function with a loss of 0.28% of global GDP at a 1 degree Celsius warming.

B Updating rules for climate sensitivity prior and predictive distribution

This appendix derives the updating rules for the climate sensitivity prior and the predictive distribution for temperature. Let $l_t(x_{t+1}|s) = \mathcal{N}(\mu_{x,t+1}, \sigma_T^2|s, x_t, h_t)$ denote the likelihood function in period t. Then[13]

$$\Pi(s|\hat{T}_1, ..., \hat{T}_{t+1}) = \frac{l_t(x_{t+1}|s)\Pi(s|\hat{T}_1, ..., \hat{T}_t)}{\int_{-\infty}^{\infty} l_t(x_{t+1}|s)\Pi(s|\hat{T}_1, ..., \hat{T}_t) ds} \ .$$

[13]This simplified updating equation only using the latest prior and the latest observation is a consequence of our convenient choice of the conjugate prior.

We the sign $\propto$ to denote proportionality and suppress the normalization constants of the distributions, finding

$$
\begin{aligned}
l_t(x|s)\ \Pi(s|\hat{T}_1,...,\hat{T}_t) &\propto \exp\left(-\frac{(x-\mu_{x,t+1}(s))^2}{2\sigma_T^2}\right)\exp\left(-\frac{(s-\mu_{s,t})^2}{2\sigma_{s,t}^2}\right)\\
&\propto \exp\left(-\frac{(x-(s\chi_t+\xi_t))^2}{2\sigma_T^2}-\frac{(s-\mu_{s,t})^2}{2\sigma_{s,t}^2}\right)\\
&\propto \exp\left(-\frac{x^2-2x(s\chi_t+\xi_t)+(s\chi_t+\xi_t)^2}{2\sigma_T^2}-\frac{s^2-2s\mu_{s,t}+\mu_{s,t}^2}{2\sigma_{s,t}^2}\right)\\
&\propto \exp\left(-\frac{x^2-2xs\chi_t-2x\xi_t+s^2\chi_t^2+2s\chi_t\xi_t+\xi_t^2}{2\sigma_T^2}-\frac{s^2-2s\mu_{s,t}+\mu_{s,t}^2}{2\sigma_{s,t}^2}\right)\\
&\propto \exp\left(-\frac{1}{2}\left[s^2\left(\frac{\chi_t^2}{\sigma_T^2}+\frac{1}{\sigma_{s,t}^2}\right)-2s\left(\frac{(x-\xi_t)\chi_t}{\sigma_T^2}+\frac{\mu_{s,t}}{\sigma_{s,t}^2}\right)+\frac{x^2-2x\xi_t+\xi_t^2}{\sigma_T^2}+\frac{\mu_{s,t}^2}{\sigma_{s,t}^2}\right]\right)\\
&\propto \exp\left(-\frac{1}{2}\left[s^2\left(\frac{\chi_t^2}{\sigma_T^2}+\frac{1}{\sigma_{s,t}^2}\right)-2s\left(\frac{(x-\xi_t)\chi_t}{\sigma_T^2}+\frac{\mu_{s,t}}{\sigma_{s,t}^2}\right)+\frac{(x-\xi_t)^2}{\sigma_T^2}+\frac{\mu_{s,t}^2}{\sigma_{s,t}^2}\right]\right)\\
&\propto \underbrace{\exp\left(-\frac{1}{2}\left(\frac{\chi_t^2}{\sigma_T^2}+\frac{1}{\sigma_{s,t}^2}\right)\left(s-\frac{\frac{(x-\xi_t)\chi_t}{\sigma_T^2}+\frac{\mu_{s,t}}{\sigma_{s,t}^2}}{\frac{\chi_t^2}{\sigma_T^2}+\frac{1}{\sigma_{s,t}^2}}\right)^2\right)}_{\equiv\bar{\Pi}}\\
&\qquad\cdot\exp\left(-\frac{1}{2}\left[-\frac{\left(\frac{(x-\xi_t)\chi_t}{\sigma_T^2}+\frac{\mu_{s,t}}{\sigma_{s,t}^2}\right)^2}{\frac{\chi_t^2}{\sigma_T^2}+\frac{1}{\sigma_{s,t}^2}}+\frac{(x-\xi_t)^2}{\sigma_T^2}+\frac{\mu_{s,t}^2}{\sigma_{s,t}^2}\right]\right)\\
&\propto \bar{\Pi}\cdot\exp\left(\frac{1}{2}\frac{\cancel{\left(\frac{(x-\xi_t)\chi_t}{\sigma_T^2}\right)^2}+2\frac{(x-\xi_t)\chi_t}{\sigma_T^2}\frac{\mu_{s,t}}{\sigma_{s,t}^2}+\cancel{\left(\frac{\mu_{s,t}}{\sigma_{s,t}^2}\right)^2}-\cancel{\frac{(x-\xi_t)^2}{\sigma_T^2}\frac{\chi_t^2}{\sigma_T^2}}-\frac{\mu_{s,t}^2}{\sigma_{s,t}^2}\frac{\chi_t^2}{\sigma_T^2}-\frac{(x-\xi_t)^2}{\sigma_T^2}\frac{1}{\sigma_{s,t}^2}-\cancel{\frac{\mu_{s,t}^2}{\sigma_{s,t}^2}\frac{1}{\sigma_{s,t}^2}}}{\frac{\chi_t^2}{\sigma_T^2}+\frac{1}{\sigma_{s,t}^2}}\right)\\
&\propto \bar{\Pi}\cdot\exp\left(-\frac{1}{2\sigma_T^2\sigma_{s,t}^2}\frac{(x-\xi_t)^2-2(x-\xi_t)\chi_t\mu_{s,t}+\mu_{s,t}^2\chi_t^2}{\frac{\chi_t^2}{\sigma_T^2}+\frac{1}{\sigma_{s,t}^2}}\right)\\
&\propto \bar{\Pi}\cdot\exp\left(-\frac{1}{2}\frac{(x-\xi_t-\chi_t\mu_{s,t})^2}{\chi_t^2\sigma_{s,t}^2+\sigma_T^2}\right)\ .
\end{aligned}
$$

The following predictive distribution P_{t+1} governs the temperature realization in period $t+1$ incorporating stochasticity and parameter uncertainty

$$P_{t+1}(x) = \int_{-\infty}^{\infty} l_t(x_{t+1}|s)\Pi(s|\hat{T}_1, ..., \hat{T}_t)ds \propto \exp\left(-\frac{1}{2}\frac{(x-\xi_t-\chi_t\mu_{s,t})^2}{\chi_t^2\sigma_{s,t}^2+\sigma_T^2}\right) .$$

It is the normal distribution $\mathcal{N}(\chi_t\mu_{s,t}, \chi_t^2\sigma_{s,t}^2+\sigma_T^2)$. We find the posterior

$$\Pi(s|\hat{T}_1, ..., \hat{T}_{t+1}) = \frac{l_t(x_{t+1}|s)\Pi(s|\hat{T}_1, ..., \hat{T}_t)}{\int_{-\infty}^{\infty} l_t(x_{t+1}|s)\Pi(s|\hat{T}_1, ..., \hat{T}_t)ds}$$

$$\propto \exp\left(-\frac{1}{2}\left(\frac{\chi_t^2}{\sigma_T^2}+\frac{1}{\sigma_{s,t}^2}\right)\left(s-\frac{\frac{(\hat{T}_{t+1}-\xi_t)\chi_t}{\sigma_T^2}+\frac{\mu_{s,t}}{\sigma_{s,t}^2}}{\frac{\chi_t^2}{\sigma_T^2}+\frac{1}{\sigma_{s,t}^2}}\right)^2\right) .$$

Thus, if $\Pi(s|\hat{T}_1, ..., \hat{T}_t)$ is distributed normally with expected value $\mu_{s,t}$ and variance $\sigma_{s,t}$, then the posterior in the subsequent period $\Pi(s|\hat{T}_1, ..., \hat{T}_{t+1})$ is also distributed normally with expected value

$$\mu_{s,t+1} = \frac{\frac{\chi_t^2}{\sigma_T^2}\frac{\hat{T}_{t+1}-\xi_t}{\chi_t}+\frac{1}{\sigma_{s,t}^2}\mu_{s,t}}{\frac{\chi_t^2}{\sigma_T^2}+\frac{1}{\sigma_{s,t}^2}} = \frac{\chi_t^2\sigma_{s,t}^2\frac{\hat{T}_{t+1}-\xi_t}{\chi_t}+\sigma_T^2\mu_{s,t}}{\chi_t^2\sigma_{s,t}^2+\sigma_T^2}$$

and variance

$$\sigma_{s,t+1} = \left(\frac{\chi_t^2}{\sigma_T^2}+\frac{1}{\sigma_{s,t}^2}\right)^{-1} = \frac{\sigma_T^2\sigma_{s,t}^2}{\chi_t^2\sigma_{s,t}^2+\sigma_T^2} .$$

References

Anthoff, D. and R. Tol (2009). The impact of climate change on the balanced growth equivalent: An application of fund. *Environmental and Resource Economics 43*, 351–367. 10.1007/s10640-009-9269-5.

Fischer, C. and M. Springborn (2011). Emissions targets and the real business cycle: Intensity targets versus caps or taxes. *Journal of Environmental Economics and Management 62*(3), 352 – 366.

Hoel, M. and L. Karp (2001). Taxes and quotas for a stock pollutant with multiplicative uncertainty. *Journal of Public Economics 82*(1), 91 – 114.

IPCC (2013). Climate change 2013 - the physical science basis - contribution of working group 1 to the ipcc - summary for policymakers. Technical report, IPCC.

Karp, L. and J. Zhang (2006). Regulation with anticipated learning about environmental damages. *Journal of Environmental Economics and Management 51*, 259–270.

Kelly, D. L. (2005). Price and quantity regulation in general equilibrium. *Journal of Economic Theory 125*(1), 36 – 60.

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

Kelly, D. L. and Z. Tan (2013). Learning and climate feedbacks: Optimal climate insurance and fat tails. Working paper, University of Miami.

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

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

Lemoine, D. and C. Traeger (2013). Watch your step: Optimal policy in a tipping climate. *American Economic Journal: Economic Policy forthcoming*.

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

Millner, A., S. Dietz, and G. Heal (2013). Scientific ambiguity and climate policy. *Environmental and Resource Economics 55*(1), 21–46.

Miranda, M. J. and P. L. Fackler (2002). *Applied Computational Economics and Finance*. MIT Press.

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

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

Traeger, C. P. (2011). Subjective risk, confidence, and ambiguity. CUDARE Working Paper 1103R2, UC Berkeley: Department of Agricultural and Resource Economics.

Traeger, C. P. (2012a). A 4-stated dice: Quantitatively addressing uncertainty effects in climate change. CUDARE Working Paper 1130, University of California, Berkeley.

Traeger, C. P. (2012b). Why uncertainty matters - discounting under intertemporal risk aversion and ambiguity. Cudare Working Paper 1092, Department of Agricultural and Resource Economics, University of California Berkeley.

Table 2: Parameters of the model

Economic Parameters		
η	2	intertemporal consumption smoothing and risk aversion
RAA	$10, 100$	coefficient of relative ambiguity aversion
b_1	0.00284	damage coefficient
b_2	2	damage exponent
δ_u	1.5%	pure rate of time preference
L_0	6514	in millions, population in 2005
L_∞	8600	in millions, asymptotic population
g_L^*	0.035	rate of convergence to asymptotic population
K_0	137	in trillion 2005-USD, initial global capital stock
δ_K	10%	depreciation rate of capital
κ	0.3	capital elasticity in production
A_0	0.0058	initial labor productivity; corresponds to total factor productivity of 0.02722 used in DICE
$g_{A,0}$	1.31%	initial growth rate of labor productivity; corresponds to total factor productivity of 0.92% used in DICE
δ_A	0.1%	rate of decline of productivity growth rate
σ_0	0.1342	CO_2 emission per unit of GDP in 2005
$g_{\sigma,0}$	-0.73%	initial rate of decarbonization
δ_σ	0.3%	rate of decline of the rate of decarbonization
a_0	1.17	cost of backstop 2005
a_1	2	ratio of initial over final backstop cost
a_2	2.8	cost exponent
g_Ψ	-0.5%	rate of convergence from initial to final backstop cost
Climatic Parameters		
T_0	0.76	in °C, temperature increase of preindustrial in 2005
σ_T^2	$0.042, 0.2, 0.7$	temperature stochasticity
M_{pre}	596.4	in GtC, preindustiral stock of CO2 in the atmosphere
M_0	808.9	in GtC, stock of atmospheric CO_2 in 2005
$\delta_{M,0}$	1.7%	initial rate of decay of CO2 in atmosphere
$\delta_{M,\infty}$	0.25%	asymptotic rate of decay of CO2 in atmosphere
δ_M^*	3%	rate of convergence to asymptotic decay rate of CO2
B_0	1.1	in GtC, initial CO2 emissions from LUCF
g_B	-1%	growth rate of CO2 emisison from LUCF
$\mu_{s,0}$	3	climate sensitivity prior mean in $t = 0$
$\sigma_{s,0}^2$	3	climate sensitivity prior variance in $t = 0$
EF_0	-0.06	external forcing in year 2000
EF_{100}	.3	external forcing in year 2100 and beyond
σ_{forc}	3.2%	warming delay, heat capacity atmosphere
σ_{ocean}	0.7%	warming delay, ocean related