# Pricing an Unknown Climate[a]

Svenn Jensen[b] Christian Traeger[c]

Preliminary Draft

**Abstract**

Anthropogenic climate change is subject to a multitude of highly uncertain feedback processes making the long-run impact of current emissions also highly uncertain. At present, we cannot reliably quantify the likelihood of differing global warming scenarios. Decision theory distinguishes between known, quantifiable risks and situations of ambiguity or deep uncertainty. A fully rational decision maker can respond differently to ambiguity and to risk, and real-world decision makers frequently do. We show how aversion to ambiguity affects optimal climate policy in an integrated assessment of climate change. We derive an analytic social cost of carbon formula for an ambiguity averse decision maker in a generic integrated assessment model. We also quantify the impact of recursive smooth ambiguity aversion for a stochastic dynamic programming implementation of DICE. Previous and paralleling approaches suggest substantial ambiguity premia on the optimal carbon tax. Our results show that the ambiguity premium is very small and optimal policy deriving from the standard Bayesian model is robust to ambiguity concerns under moderately large ambiguity aversion if climate policy is endogenous and policy maker's have rational foresight.

**JEL Codes:** Q54, Q00, D90, C63
**Keywords:** climate change, uncertainty, integrated assessment, DICE, ambiguity, ambiguity aversion, smooth ambiguity, dynamic programming.

[a] Several results of this paper were part of our earlier work "Optimally Climate Sensitive Policy", which we split into papers with a focus on risk and learning, and the present research on ambiguity.

[b] Oslo Business School, Oslo Metropolitan University; svenn.jensen@oslomet.no.

[c] Department of Economics, University of Oslo; ifo Institute for Economics Research; Frisch Centre; mail@traeger.eu.

# 1 Introduction

Carbon dioxide emissions drive our current economy and harm our future climate and well-being. The immediate benefits of our emissions are usually known (and priced). In contrast, future damages rely on a limited number of estimates, subjective guesstimates, and incomplete modeling. Integrating uncertainty into evaluations of climate change has been a major recent advancement. Unfortunately, the underlying probability distributions are themselves unknown or highly uncertain. A slowly growing literature takes "the next step" and incorporates ambiguity or Knightian uncertainty into climate change evaluation; they incorporate the lack of well quantified objective probabilities. These papers commonly find major ambiguity premia for the price of carbon, increasing the optimal carbon tax beyond the risk premium. The present paper argues that we have to carefully distinguish between ambiguity and ambiguity aversion. Agreeing with earlier papers, we have little doubt regarding the ubiquity of ambiguity in climate change evaluation. Yet, we argue that earlier results are less driven by ambiguity per se, but by choices of ambiguity aversion and, at times, the exogeneity of climate policy. We present a novel evaluation of the social cost of carbon under ambiguous climate feedbacks. Our social planner respects certain rationality constraints that give rise to the recursive smooth ambiguity model of Klibanoff et al. (2009). Moreover, we incorporate a climate policy that responds optimally to the resolution of uncertainty. We show that the employed interpretation of rationality and foresightedness shrinks the ambiguity premium to an almost negligible level.

A ton of carbon released today encounters several major uncertainties along its path to causing future economic damages. The main climate uncertainty is how atmospheric carbon affects temperatures, summarized by the climate sensitivity; the equilibrium warming from doubling the $CO_2$ concentration relative to pre-industrial levels. It depends on a set of highly uncertain feedback processes. There is substantial disagreement about its value, which is derived primarily from different climate models but also instrumental records and paleo-climatic research. The Intergovernmental Panel on Climate Change (IPCC) states that a doubling of $CO_2$ concentrations likely yields a warming between 2.5°C and a 4°C, with a current best guess of 3°C. 'Likely' is defined by the IPCC as two-thirds subjective probability. Significantly higher warming still carries substantial probability mass (Masson-Delmotte et al. 2021). Hence the same emission trajectory can result in drastically different economic realities, and there is substantial uncertainty governing any probabilistic prediction. Sir Nicholas Stern, a prominent figure in the climate change debate, argued already more than ten years ago that "[t]he difference between risk and uncertainty (. . . ) is a very important issue and a key topic for further research" (Stern 2008, p.3), a call that has been repeated over the years (Burke et al. 2016).

We analyze the response of optimal climate mitigation policy to climate ambiguity in an integrated assessment model of climate change. We derive an analytic social cost of carbon (SCC) formula for an ambiguity averse social planner in a generic climate-economy model. We quantify this SCC using the functional forms and their calibrations from the DICE integrated assessment model (Nordhaus 2017). Our analysis employs the smooth ambiguity model by Klibanoff et al. (2009). The social planer is uncertain about the right probability distribution for the climate sensitivity and also faces stochastic temperatures. The planner is more averse to the subjective uncertainty about the climate sensitivity than to the quantifiable temperature fluctuations.

Our analytic insights extend the analytic SCC formula in Jensen & Traeger (2021) by two terms: An 'ambiguity prudence' term and an 'ambiguity pessimism' term. These terms weigh the different possible welfare outcomes resulting from different possible climate sensitivities. The pessimism term gives more weight to bad outcomes, i.e., high realizations of the climate's sensitivity to emissions, which incentivizes more abatement. The prudence term is similar to the one in the precautionary savings literature, where a prudent decision maker saves more when facing uncertainty over future wealth. Here, a prudent decision maker operates in a more complex environment, hence her course of action isn't as definite. For decision maker's whose ambiguity attitude exhibits prudence (decreasing aversion in wealth and welfare) it generally raises the willingness to invest today into a better future. However, the non-linear relation translating climate uncertainty into welfare uncertainty can alter the effect's magnitude and sign.

We quantify our finding in a stochastic dynamic programming model where ambiguity and Bayesian learning govern the climate sensitivity prior. Today's decision maker rationally foresees future uncertainty and learning. We solve the infinite horizon problem and pin down the optimal carbon tax and the contribution of ambiguity aversion. We find that, along the optimal policy time path, ambiguity and ambiguity aversion hardly alter the optimal policy. For very high levels of atmospheric carbon (much higher than optimal) the impact of ambiguity aversion on the social cost of carbon is more pronounced but still not dominant.

Lange & Treich (2008) discuss ambiguity aversion in a stylized two-period climate change mitigation model. They show that even in such a two-period linear stock accumulation problem the qualitative effect of ambiguity on the optimal policy cannot be signed uniquely. Berger et al. (2016) apply recent robust preferences (Marinacci 2015) to a two-period climate-economy model. This robust preference specification is in a two-period setting mathematically equivalent to Klibanoff et al. (2009). Like Lange & Treich (2008), they find that the ambiguity effect is not easily signed. They analytically derive an 'ambiguity prudence' condition and quantify the impact of their ambiguity preferences. To that end they develop a version of the DICE model in which a catastrophe of fixed size either happens at a known date or never. Berger & Marinacci (2020) review the ambiguity preference specifications that were recently used in climate economics. In a stylized model they quantitatively compare the impact of the modeling choices on optimal mitigation. Millner et al. (2013) model smooth ambiguity aversion over climate sensitivity. In their model policy is exogenous and stochasticity derives from switching between two different deterministic trajectories generated with different assumptions on climate sensitivity. The authors find that ambiguity aversion causes a small welfare loss under the standard DICE damage function and large welfare loss when employing a more convex damage function. In our stochastic model where the system evolves smoothly and the policy maker chooses the mitigation policy endogenously we find a much smaller impact of ambiguity aversion on welfare. Closest to the present paper is the recent work by Barnett et al. (2020) who apply a continuous time version of the Hansen & Sargent (2001) robust preferences to a climate-economy model with uncertain climate sensitivity and damages. They quantify the impact of ambiguity aversion (as well as model misspecification) on the SCC, finding an economically significant increase. Their model differs in important respects from ours: Their preference specification is closely related to Klibanoff et al. (2005), while we use Klibanoff et al. (2009) (see below in section 2). They 'switch off' the learning dynamics while in our numeric analysis the social planner applies Bayesian updating. Finally, their risk- and consumption smoothing preferences are logarithmic while ours are the slightly more general constant relative risk aversion preferences. Lemoine & Traeger (2016), inspired by the present analysis, test whether the impact of ambiguity aversion is larger in a setting with tipping points that cause more abrupt changes in system dynamics and welfare.[1] They find that also in the case of tipping points ambiguity over the location of the tipping threshold implies only a moderate premium on the social cost of carbon. We further this literature by providing the first application of the smooth ambiguity model by Klibanoff et al. (2009) to analyze climate policy in a full-fledged integrated assessment model.

# 2 Climate Change and Ambiguity Evaluation

Climate feedbacks showcase the complexity of global warming and the resulting depth of our uncertainty. These feedbacks responds to the initial warming (radiative forcing) caused by our carbon dioxide emissions. (i) The initial warming increases the atmospheric content of water vapor, which itself is a strong greenhouse gas currently contributing the largest share to the $\sim 18°$C warming of the natural greenhouse effect. (ii) Higher temperatures melt glaciers and ice caps, reducing the reflectivity of our planet's surface, further increasing its uptake of solar radiation. (iii) Soils that are currently subject to permafrost will thaw and release methane, which is another powerful greenhouse gas, potentially starting a powerful thawing-warming cycle. (iv...) Other feedbacks affect the heat exchange in the atmosphere or carbon cycle processes, and yet others may still be unknown. As a consequence of such feedbacks and their interactions, the actual warming resulting from a given greenhouse gas concentration is deeply uncertain. Figure 1 shows an overview over confidence intervals compiled by the IPCC in its Fifth Assessment Report (AR5) (Stocker et al. 2013). How to aggregate vastly different distributions into a single probabilistic estimate is not obvious. AR5

[1] The present paper is a continuation of work initially reported in Jensen & Traeger (2013).

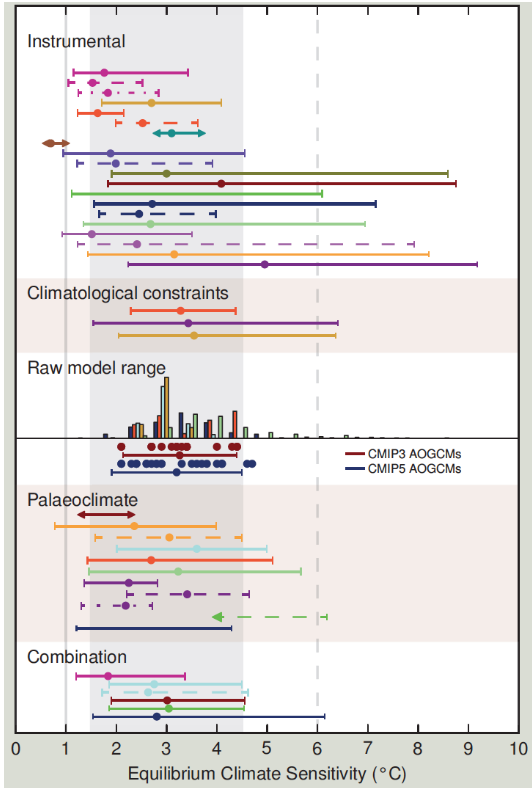


Figure 1: Subjective estimates of the equilibrium climate sensitivity from different sources. Figure from the Fifth Assessment Report by the IPCC (Stocker et al. 2013).

stated that a doubling of $CO_2$ concentrations 'likely'[2] yields a warming between 1.5°C and a 4.5°C, but refused to give a 'best guess'. The Sixth Assessment Report (Masson-Delmotte et al. 2021) changed this likely interval to 2.5°C to 4°C with a best estimate of $3°C$.

Already Keynes (1921) suggests that the usual probabilistic model is ill-equipped to deal with corresponding situations where we lack confidence into our probabilistic description of the world. Decision theorists demonstrate in experiments that, indeed, many decision makers do not adhere to our standard economic models of objective or subjective expected utility when confronted with unknown or highly subjective probability distributions. Today, decision theory mostly summarizes this type of uncertainty under 'ambiguity'; phrases such as

[2]They assign the term a subjective 66 percent probability.

'deep uncertainty' or 'hard uncertainty' remain in circulation in the more applied literature. We refer to, e.g., Machina & Siniscalchi (2014) for a discussion of the literature. This literature offers a large variety of alternative approaches to evaluate situations lacking a trusted and unique probabilistic description of the future. We briefly review three of the more common classes of models. A first class abandons the use of probabilities, replacing them by so-called capacities. Capacities lack the additive nature of probabilities and an event can happen, not happen, and then there is room for something else, which makes their interpretation and application to the climate change evaluation rather intricate. As second class introduces sets of probability distributions rather than singletons. This approach seems more suitable for policy analysis. Indeed, this specification has been chosen by several of the earlier applications in the integrated assessment of climate change. The most popular models evaluate outcomes merely based on worst (or best) expected outcomes, which holds in particular for all of the climate change application in this category. The common evaluation based on the worst expected outcome can also be recovered as a limiting case of a third approach. Here, the set of possible probability distributions is governed by a subjective second order distribution that weighs the different possible objective (stochastic) descriptions of the world. The most famous model in this class is Klibanoff et al.'s (2005) model of smooth ambiguity aversion. Here, ambiguity aversion governs the additional aversion with respect to subjective second order beliefs. The approach allows the authors to characterize ambiguity aversion in the same way as usual risk aversion using Arrow-Pratt measures. Letting (relative) smooth ambiguity aversion approach infinity brings us to the limit case of only evaluating the worst possible expected outcome (as in the second class of models).

Most ambiguity models are formulated in an atemporal setting. Applying these models in a dynamic context gives rise to a variety of inconsistencies that make them rather unattractive for social planning. In our view, the major exception is Klibanoff et al. (2009) recursive smooth ambiguity model, a dynamic extension of the atemporal smooth ambiguity model. By recursively evaluating the decision tree, the model achieves time consistency and is compatible with simple Bayesian updating. The model also falls into Traeger's (2011) normative axiomatic foundation focusing on decision makers that satisfy the von Neumann

& Morgenstern (1944) axioms, once we distinguish between objective and subjective (or low confidence) priors. As a result, we adopt Klibanoff et al.'s (2009) recursive smooth ambiguity framework for our evaluation. Another advantage of the framework is its distinction between ambiguity and ambiguity aversion.

Our evaluation of climate change will maximize welfare under constant relative (smooth) ambiguity aversion. In the limit of infinite ambiguity aversion, this model converges to Gilboa & Schmeidler's (1989) maximin expected utility; here, a decision maker only considers the worst possible expected outcome under the set of possible probability distributions. On the level of subjective priors, this approach is analogous to Arrow & Hurwicz's (1972) maximin criterion judging a set of possible outcomes merely by its worst outcome.[3] Instead of simple outcomes as in Arrow & Hurwicz's (1972), the maximin expected utility model in consider a set of possible distributions and expected outcomes. Just as we can obtain maximin utility in the limit of constant relative risk aversion going to infinite, we can obtain maximin expected utility by letting aversion with respect to the subjective prior over possible stochastic models go to infinity. Hansen & Sargent (2001) give assumptions under which this limit of maximin expected utility is equivalent to robust control. Robust control, maximin expected utility, and the transformation between the two models is used repeatedly in the literature on ambiguity or deep uncertainty in climate change (Brock & Xepapadeas 2021, Barnett et al. 2020, Rudik 2020, Olijslagers & van Wijnbergen 2019, Xepapadeas & Yannacopoulos 2018, Berger et al. 2016). A similar approach is also known as "playing a game against nature", where nature always takes the worst possible move, i.e., draws the worst possible realization. The first observation is that, from the perspective of (smooth) ambiguity aversion, robust control or maximin expected utility correspond to infinite smooth ambiguity aversion. Our second observation is that Gilboa & Schmeidler's (1989) maximin expected utility model, just like the simple smooth ambiguity model of Klibanoff et al. (2005) is atemporal. As Klibanoff et al.'s (2009) explain, a naive combination of smooth ambiguity to a temporal setting leads to inconsistencies. As a result, they develop a more sophisticated model recursively evaluating over time and ambiguity. This recursive evaluation

[3]Arrow & Hurwicz (1972) show that under certain assumptions on ignorance, a decision should only be based on the worst or the best possible outcome.

reflects that present decision maker's incorporates how future decision maker's evaluate and respond to uncertainty. It turns out that this recursive consideration of future decisions will lower the effective premium of ambiguity aversion. The robust control approach has, to the best of our understanding, no such recursive interpretation.

# 3 The General Model

This section introduces a simple and general integrated assessment model of climate change. It couples a growing world economy to the climate through greenhouse gas emissions and economic impacts of climate change. Figure 2 presents a schematic of our model (the full set of equations can be found in Appendix B). The right hand side characterizes a standard Ramsey growth model. World output is a function $F(K_t, T_t, E_t, t)$ of endogenous capital $K_t$, atmospheric temperature $T_t$, carbon dioxide emissions $E_t$, and a set of exogenous processes including labor and technology levels that depend on time $t$. The temperature increase $T_t$ (measured in degree Celsius above 1900 levels) reduces the economy's productivity. Output is spent either on consumption $C_t$, capital investment, or emission reductions. We treat the economy's carbon dioxide emissions as an input into the production process; it is a reduced representation of an energy production sector whose carbon dioxide emissions trade off with capital and labor.Emissions build up the stock of atmospheric carbon $M_t$. Atmospheric carbon together with other (exogenous) greenhouse gases trap our planet's outgoing infrared radiation. The resulting shift in the planet's energy balance causes an initial warming and induces further feedback processes.

The climate sensitivity $s$ ('CS' in Figure 2) measures the medium to long-run temperature increase resulting from a doubling of pre-industrial atmospheric carbon dioxide concentration. The approximate "climate change law" specifies that every doubling of atmospheric carbon dioxide (or fractions thereof) implies a $s$ degree Celsius temperature increase (or fraction thereof). It takes decades to centuries to reach a new equilibrium temperature. As we discuss in the next section, climate sensitivity is an unknown parameter governed by a prior and Bayesian updating. The updating relies of the observation of temperature and

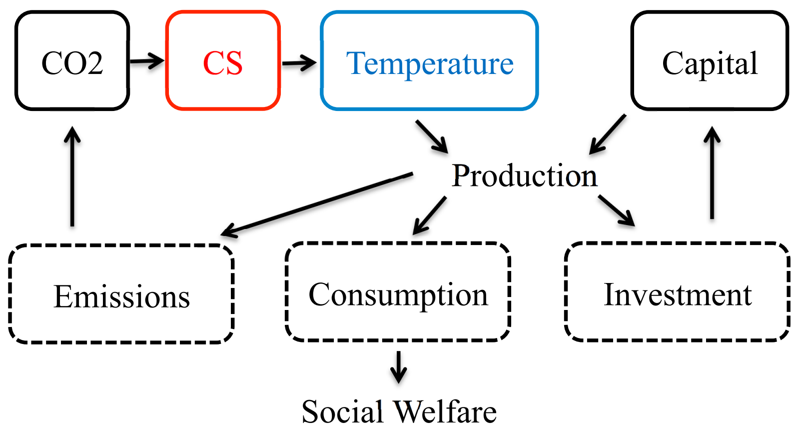


Figure 2: The main relations in the climate-economy model. Dashed rectangles represent control variables. Solid rectangles depict the main state variables of the system. Climate sensitivity ('CS') is uncertain. The model with learning represents climate sensitivity by a Bayesian prior (2 state variables). Temperature is stochastic.

carbon concentrations. Temperature evolution is also subject to a annual volatility modeled as Gaussian white noise. The change in temperature feeds back into economic production causing damages. The decision maker chooses emissions and consumption (and implicitly investment) in every period. We discuss her objective function in section 3.2 and solve the model using stochastic dynamic programming. Our analytic results rely only on this general model structure. For our numeric results, we pick the functional forms and quantitative assumptions of the DICE model summarized in Appendix B.

## 3.1 Temperature stochasticity and Bayesian learning

A model of the (slow) structural learning observed in climate change requires a model of Bayesian uncertainty over the medium to long-run temperature that will prevail after the feedback processes adjusted. It also requires accounting for temperature stochasticity, which shields the learning. For a given climate sensitivity $s$, temperature $T_t$ evolves according to

$$\tilde{T}_{t+1} = \chi_t(M_{t+1})s + \xi_t(T_t) + \tilde{\epsilon}_t \ . \tag{1}$$

The factor $\chi_t$ captures the forcing from atmospheric $CO_2$ and other greenhouse gases. The term $\xi_t$ reflects that atmospheric warming is a slow process and approximately governed by an AR(1) type process with a moving target.

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

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

The variance $\sigma_T^2$ is exogenous. Empirical estimates suggest annual volatility in global mean temperature in $\sigma_T^2 = 0.042$.[4] For analytic purposes we will also use considerably larger values.

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

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

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, $\sigma_{s,0}^2 = 3$ is a rounded-up empirical approximation to the set of distributions found in Stocker et al. (2013)[5]

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 assumed to be 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.

[4]Kelly & 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 $\sigma_T^2 = 0.042$.

[5]A normal distribution with mean 3 and variance 3 has the one-standard-deviation bands [1.27,4.73], which also mimics the IPCC's "likely" range [1.5,4.5].

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 show the following updating rules for the expected value

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

and the variance

$$\sigma^2_{s,t+1} = \frac{\sigma_T^2 \sigma_{s,t}^2}{\chi_t^2 \sigma_{s,t}^2 + \sigma_T^2} \; . \tag{3}$$

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 $\chi_t$, 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 (3) grows in $1/\sigma_T^2$ and in $\chi_t$. With uncertainty about climate sensitivity, temperature realizations are themselves uncertain (not only stochastic) and governed by the predictive distribution $\tilde{T}_{t+1} \sim \mathcal{N}(\xi_t + \chi_t \mu_{s,t}, \chi_t^2 \sigma^2_{s,t} + \sigma_T^2)$. We can conveniently use this distribution (see Appendix B) to evaluate the uncertainty in the optimization.

## 3.2 Welfare and Aversion to Climate Change Ambiguity

The decision maker's evaluation of future climate change depends on her uncertainty preferences. As discussed in section 2, we follow Klibanoff et al.'s (2009) axiomatization of smooth ambiguity aversion because it satisfies normatively attractive axioms including time consistency. These preferences capture a decision maker who prefers a world with well-known probabilities to a world governed by subjective uncertainty.

The social planner evaluates the risk resulting from stochastic movements of temperature $T$ with an increasing and concave utility function $u(c_t)$. As in the common intertemporally additive expected utility model, this utility function describes her risk aversion as well as her desire to smooth consumption over time. Another increasing and concave function $f(z)$ captures the *additional* aversion towards the subjective uncertainty. In our application, utility is a population-weighted constant relative risk aversion function $L_t u(c_t) = L_t \frac{c_t^{1-\eta}}{1-\eta}$ from per capita consumption $c_t = \frac{C_t}{L_t}$ with the Arrow-Pratt coefficient of relative risk aversion $\eta$. Similarly, we use a constant relative ambiguity aversion function characterized by the coefficient of relative ambiguity aversion RAA.[6] In our application, the concavity of the magnitude of RAA (or concavity of $f$) captures additional aversion when evaluating the future uncertainty resulting from the epistemological uncertainty governing the climate sensitivity parameter. The following Bellman equation pins down the social planner's value function

$$V(K_t, M_t, t, T_t, \mu_{s,t}, \sigma_{s,t}) = \max_{c_t, \mu_t} \; L_t u(c_t) + \tag{4}$$
$$\exp[-\delta_u] \quad f^{-1}\left\{ \int_S f\Big( \mathbb{E}_{\epsilon_t}\left[ V(K_{t+1}, M_{t+1}, t+1, T_{t+1}, \mu_{s,t+1}, \sigma_{s,t+1})\right] \Big) \, \mathrm{d}\, \Pi(s,t) \right\} \;,$$

which is optimized subject to the dynamic equations summarized in Appendix B, and the informational updating equations (2) and (3). The Bellman equation states that the maximized social welfare today is equal to the sum of the instantaneous social welfare. and the future maximized social welfare. By choosing the consumption level $c_t$, the decision maker balances immediate consumption against future physical capital stocks, and with her abatement decision $\mu_t$ she trades off immediate consumption and lower future carbon concentration.

Next period's temperature realization depends, apart from the atmospheric carbon concentration and current temperature, on the climate sensitivity $s$ and

[6] *RAA* stands for: Constant coefficient of **R**elative **A**mbiguity **A**version. For $\eta > 1$ the utility and value function are negative. As a result, the power functional form of the constant relative ambiguity aversion aggregator has to aggregate over the negative of utility resulting in a slightly more sophisticated definion of RAA. See Traeger (2010) for a detailed and axiomatic discussion of defining aversion in such a setting.

the realization of the iid shock $\epsilon_t$ (see equation 1). For a given value of climate sensitivity, next period temperature $T_{t+1}$ is objectively stochastic and normally distributed with variance $\sigma_T^2$. The expectation operator in the inner bracket of the Bellman equation (4) takes expectations with respect to this well-known stochasticity. Yet, the decision maker is also subjectively uncertain about the true climate sensitivity, captured by the prior $\Pi(s,t) \sim \mathcal{N}(\mu_{s,t}, \sigma_{s,t}^2)$. The integral over the space of possible climate sensitivity realizations $S$ with respect to this prior aggregates over this subjective uncertainty.[7] Here, the increasing and concave transformation $f$ expresses additional aversion with respect to future payoffs that are subject to the decision maker's epistemological uncertainty.

Next period's expected climate sensitivity also depends on both the present epistemological prior $\Pi(s,t)$ and the realization of the shock $\epsilon_t$ to global average surface temperature (see equation 2). It depends on the objective shock because of the Bayesian updating of climate sensitivity uncertainty based on the temperature observation. Again, the uncertainty deriving from the subjective prior is evaluated using the additional aversion embedded in the concave transformation $f$, and the uncertainty deriving from the objective shock $\epsilon_t$ is evaluated ambiguity neutrally.

# 4 The Social Cost of Carbon with Ambiguity Aversion

An optimal mitigation policy equates the costs from mitigating a ton of $CO_2$ with the present value of the future social damages from emitting a ton of $CO_2$. These damages are known as the social cost of carbon (SCC). Jensen & Traeger (2021) derive the optimal social cost of carbon under uncertainty for standard CRRA preferences from the first order conditions of the Bellman equation *without* ambiguity aversion. We reproduce their intuitive presentation here and then contrast it with the SCC under ambiguity aversion. We denote by $\frac{\partial M_\tau}{\partial E_0}$ the change of the carbon stock in period $\tau$ as a result of an additional ton of car-

[7] In our case $S$ is a continuous paramter space, whereas it is a finite space in Klibanoff et al.'s (2009) axiomatization.

bon emitted today. The carbon stock change in period $\tau$ affects the temperature in subsequent periods $t > \tau$ as $\frac{\partial T_t}{\partial M_\tau}$ through direct radiative forcing and feedbacks. The resulting change in future atmospheric temperature impacts the output $Y_t = F_t(K_t, T_t, E_t, t)$, causing a marginal damage of $-\frac{\partial F_t}{\partial T_t}$. The output loss reduces period $t$ welfare proportional to marginal welfare $u_t'(c_t)$. Summing the discounted welfare loss over an infinite time horizon and translating it into present day consumption units results in the analytic expression for the SCC for our integrated assessment model

$$\mathrm{SCC}_0 = -\frac{1}{u_0'(c_0)}\, \mathbb{E}_0 \sum_{t=1}^{\infty} \sum_{\tau=1}^{t} u_t'(c_t)\, \frac{\partial F_t}{\partial T_t}\, \frac{\partial T_t}{\partial M_\tau}\, \frac{\partial M_\tau}{\partial E_0}.$$

The optimal carbon tax is the SCC evaluated along the optimal trajectory.

Starting from the recursive Bellman equation, Appendix XX shows that ambiguity aversion adds two novel terms into the social cost of carbon formula: a prudence term ($\mathcal{R}$) and a pessimism term ($\mathcal{P}$). The resulting formula is

$$\mathrm{SCC}_0 = -\frac{1}{u'(c_0)}\ \mathbb{E}_0^s\, \mathbb{E}_0^T \left[ \sum_{t=1}^{\infty} \sum_{\tau=1}^{t} \prod_{j=1}^{t} \mathcal{R}_j\, \mathbb{E}_j^s\, \mathcal{P}_j\, \mathbb{E}_j^T\ \ u_t'\ (c_t)\ \frac{\partial F_t}{\partial T_t} \frac{\partial T_t}{\partial M_\tau} \frac{\partial M_\tau}{\partial E_0} \right] .$$

Expectations $\mathbb{E}_j^s$ are over the climate sensitivity prior conditional on information in period $j$, and expectations $\mathbb{E}_j^T$ govern the climate conditional on a given climate sensitivity and information in period $j$. The novel pessimism term is

$$\mathcal{P}_t = \frac{f'(\cdot)}{\mathbb{E}_t^s\left(f'(\cdot)\right)}\ ,$$

where $f'(\cdot) = f'\left(\mathbb{E}_{\Omega_s}\left[V_{t+1}(\cdot)\right]\right)$. The pessimism term acts as a weighting bias. Assuming ambiguity aversion ($f'' < 0$), events with a higher welfare $V_{t+1}$ contribute a lower marginal weight $f'(\cdot)$. Thus, the pessimism term gives more weight to bad outcomes, which are the high realizations of the climate's sensitivity to emissions. More weight on the high sensitivity scenarios increases the incentives to abate and the optimal carbon tax.

The novel prudence term is

$$\mathcal{R}_t = \frac{\mathbb{E}_t^s\left(f'(\cdot)\right)}{f'\left(f^{-1}(\cdot)\right)},$$

where again $f'(\cdot) = f'\left(\mathbb{E}_{\Omega_s}\left[V_{t+1}(\cdot)\right]\right)$. The prudence term is less symmetric in its dependence on the value function than the pessimism term. As a result, it plays a more important role how the value function $V_{t+1}$ itself responds to uncertainty. Assume for the moment that the value function itself would be subject to mean-zero shocks. Then the prudence term plays an analogous role to the prudence term known in the standard precautionary savings model. We call the decision maker ambiguity prudent if she exhibits decreasing ambiguity aversion in welfare, i.e., in $V_{t+1}$ (rather than consumption or wealth as in the precautionary savings case). Under mean-zero shocks to welfare, ambiguity prudence increases the decision maker's willingness to invest today into a better future; she does so because a higher welfare level in the future will reduce her suffering when she is hit by the uncertainty.[8]

# 5 Numeric Results

## 5.1 Numeric Implementation

The numeric model is s stochastic version of DICE with Bayesian learning about climate sensitivity based on Traeger (2014) and Jensen & Traeger (2021). It simplifies DICE's carbon cycle and the ocean cooling. This adjustment saves us three computationally costly state variables. Traeger (2014) explains the details of the implementation and shows that our calibration compares at least as favorably to the scientific climate change models (AOGCMs) as the original DICE model. We model climate sensitivity uncertainty and -learning by Bayesian updating with a conjugate normal prior for climate sensitivity and normal stochastic temperature shocks. The mean and variance of climate sensitivity are the informational states that fully capture uncertainty. After observing the temperature in a given year, the decision maker infers the implied climate sensitivity and updates her prior about it accordingly.

[8]The prudence term increases the SCC for mean-zero shocks over the next period welfare, if, and only if, absolute ambiguity aversion $-\frac{f''}{f'}$ decreases in welfare (**?**). This condition is always met for our isoelastic preference specification that we will adopt in the quantitative assessment. For a more detailed discussion of prudence and pessimism terms in integrated assessment, in the context of standard risk with Epstein-Zin-Weil preferences, see Jensen & Traeger (2014).

We solve the dynamic programming equation (4) by function iteration, using the collocation method to approximate the value function. As basis functions we choose Chebychev polynomials with $35,200$ Chebychev nodes and coefficients. The normal distributions for temperature stochasticity and the climate sensitivity prior are approximated by Gauss-Legendre quadrature.[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 & Fackler (2002) to generate and evaluate the Chebychev polynomials, and we let the solver KNITRO to carry out the optimization.

## 5.2 Ambiguity aversion along optimal policy trajectory

Figure 3 shows the impact of ambiguity aversion on socially optimal climate policy (abatement, emission levels and the social cost of carbon) and investment. We plot time paths until 2050. The policy maker has an initial climate sensitivity prior of $\mathcal{N}(3, 3)$ and expects stochastic temperature shocks $\mathcal{N}(0, 0.11^2)$. The time paths are "expected draws": Each period the actual shock is zero, and the decision maker hence observes the climate sensitivity she (correctly) expects ($\hat{s}_t = 3$). Hence her expectation doesn't change, only her confidence in the value grows over time. That way, the only difference between the different scenarios derives from the difference in policy, not different realizations of shocks and hence beliefs. The scenarios depicted are known climate sensitivity ($s = 3$) with stochastic temperature (solid black lines), uncertain climate sensitivity with constant relative risk aversion preferences (dashed red), and uncertain climate sensitivity with relative ambiguity aversion RAA= 20 (dashed-dotted blue) and RAA= 80 (dashed blue).

We see that the optimal policy is surprisingly unaffected by the presence of ambiguity aversion, irrespective of its level. There is no sizable "ambiguity premium". We 'zoom in' on this result in Figure 4. The figure shows the social cost of carbon in the year 2015 for different levels of ambiguity aversion. In the left panel, we see the same scenario as before: the DICE specification with low (but empirically accurate) temperature stochasticity. We see that the social cost of carbon increases roughly linearly with the level of ambiguity aversion. But a

[9]The results are robust to increasing Gauss-Legendre nodes and Chebychev nodes.

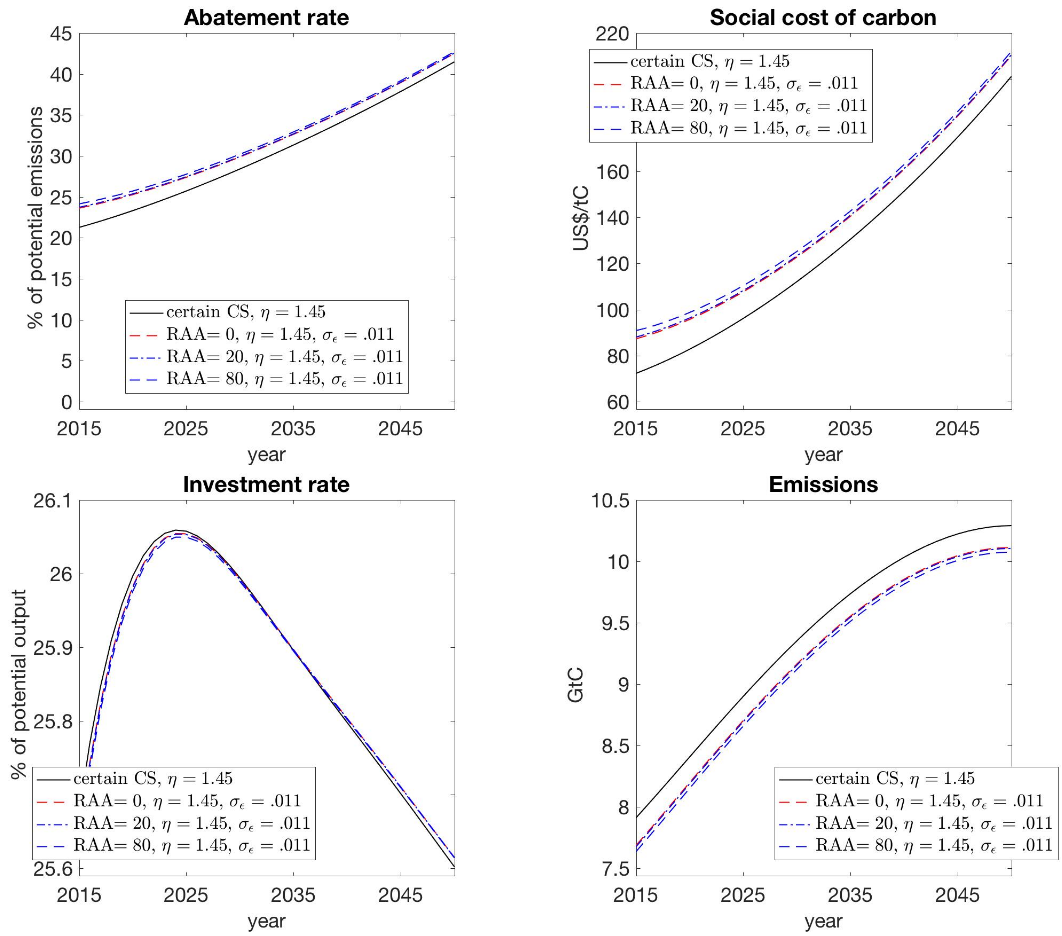


Figure 3: Abatement rate, social cost of carbon, investment rate and emissions for the current century with stochastic temperature ($\sigma_\epsilon = .011$), uncertain climate sensitivity with initial prior variance $\sigma^2_{s,0} = 3$ and three levels of ambiguity aversion: none, RRA= 20 and RRA=80. Stochastic paths generated by drawing expected value each period.

look at the scale reveals that the overall effect is small. An increase from RAA= 0 (no additional ambiguity aversion) to RAA= 80 increases the SCC by less than USD 4. The right hand panel shows the same figure for cubic damages but lower levels of relative ambiguity aversion. Why ambiguity aversion has so little impact on optimal policy is not immediately obvious and requires further investigation.

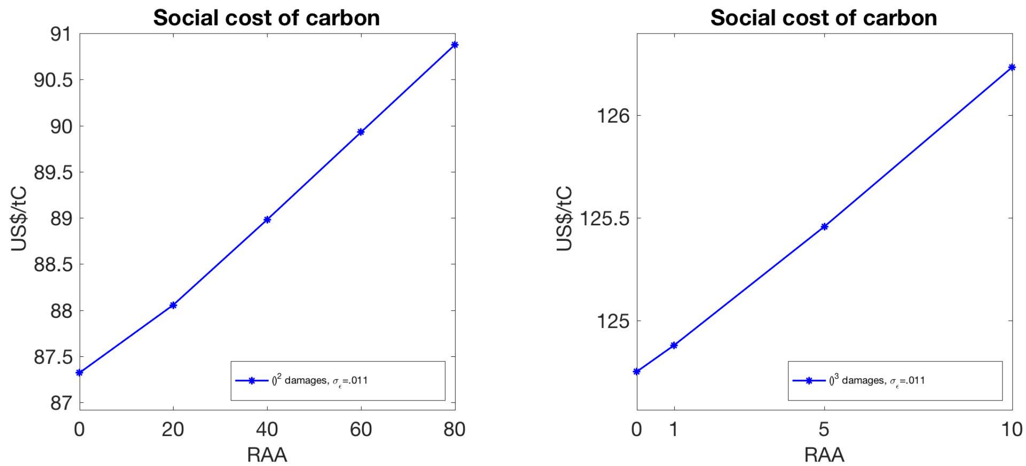


Figure 4: Social cost of carbon in 2015 over different levels of ambiguity aversion. The left panel assumes quadratic damages (as before), and the right panes assume cubic damages. Note the different scales of the axes. Climate sensitivity is uncertain with prior $\mathcal{N}(\mu_s = 3, \sigma_s^2 = 3)$.

## 5.3 Ambiguity aversion off the optimal policy path

Why do we observe no real impact of ambiguity and ambiguity aversion? Potentially, the decision maker's ability to affect the subjective uncertainty is limited. She can increase her emissions in order to learn faster and reduce the uncertainty. However, the additional learning is incremental, and it comes at the cost of being even worse-off in if the climate sensitivity turns out to be high. Alternatively, could it be that the subjective uncertainty doesn't hurt the decision maker much? To explore this second idea, we compare the numeric value functions for decision makers with and without ambiguity aversion for the same scenario. If they are identical over the state space, the ambiguity doesn't harm the averse decision maker.

For most of the state space, the value functions indeed are almost identical. Exceptions are states of the world where the expected climate sensitivity is high and highly uncertain, the levels of atmospheric carbon are high, and the physical capital stock is vast. Figure 5 shows the relative value function difference for decision makers with and without ambiguity aversion (RAA= 80) over sections of the state space (the other state space variables are kept at the values they take along the optimal time path). The scenario plotted features quadratic damages

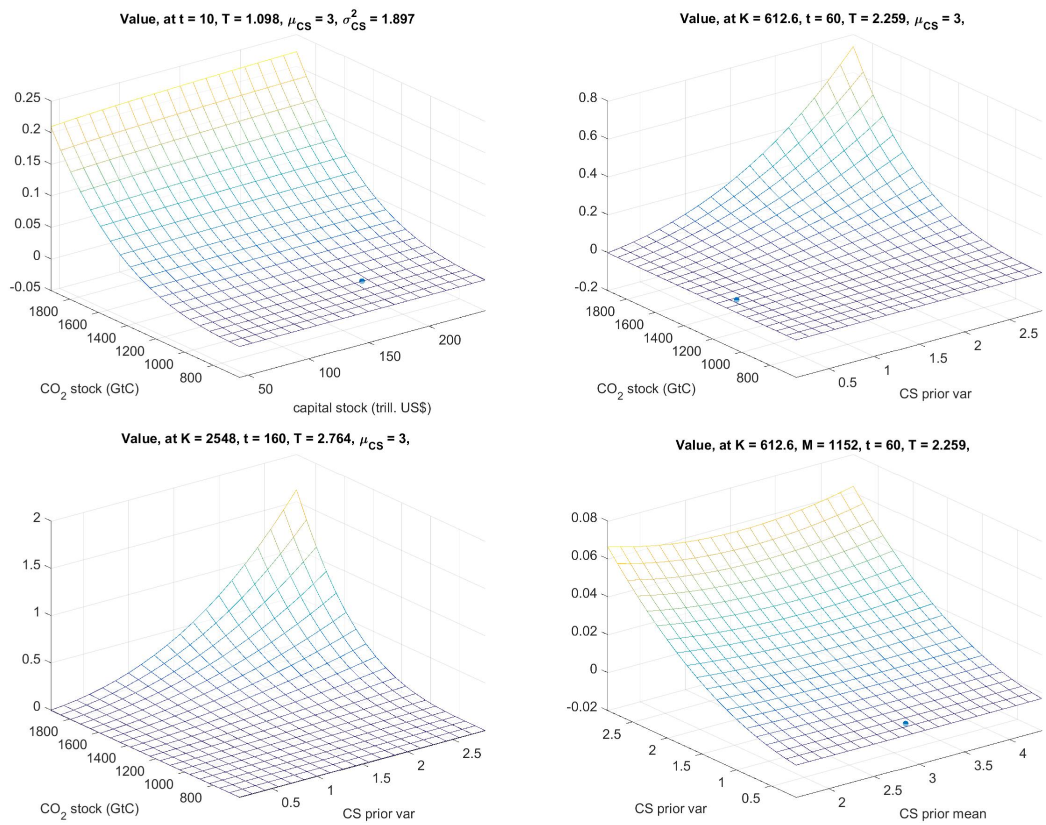


Figure 5: Relative difference in value function in percent for decision makers with and without ambiguity aversion RAA= 80 for quadratic damages and temperature stochasticity $\sigma_\epsilon = .11$. Plotted along optimal time path. Four state space variables are kept fixed, two are plotted.

and low temperature stochasticity ($\mathcal{N}(0, 0.11^2)$), as in the left panel of figure 4. The scale suggest a difference of less than 0.2 percent.Figure 7 shows the same difference for cubic damages with ambiguity aversion RAA= 10.

Figure 6 plots the relative difference in the social cost of carbon corresponding to the value function difference in Figure 5.[10] In the upper left panel we see a relative difference of up to 33 percent. The blue dot indicates the optimal time path trajectory, i.e. for the given state-space values for time, temperature and the climate sensitivity prior, the carbon and capital stocks and the social cost of carbon difference take those values. The lower left panel shows an ambiguity

[10]The social cost of carbon can be calculated from first derivates of the value function by carbon, capital and temperature, see Jensen & Traeger (2021).

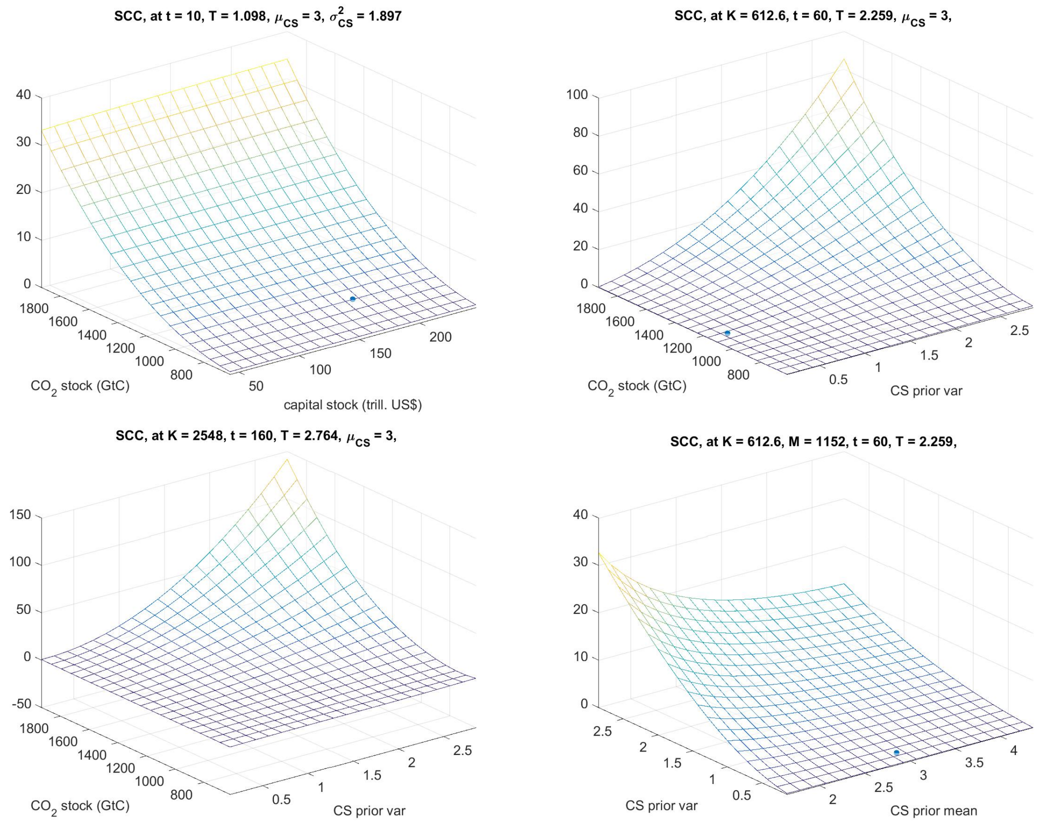


Figure 6: Relative difference in social cost of carbon in percent for decision makers with and without ambiguity aversion RAA= 80 for quadratic damages and temperature stochasticity $\sigma_\epsilon = .11$. Plotted along optimal time path. Four state space variables are kept fixed, two are plotted.

premium of almost 150 percent fir high carbon stocks and prior variance.

Figure 8 shows the same difference for cubic damages with ambiguity aversion RAA= 10.

# 6 Conclusions

Uncertainties governing the level of future climate change are large and subjective. We analyze how a social planner should or could account for the underlying ambiguity when calculating the optimal carbon tax. Our evaluation framework relies on two assumptions. First, future planners respond to the resolution of uncertainty. Second, the planner employs a dynamically consistent framework.

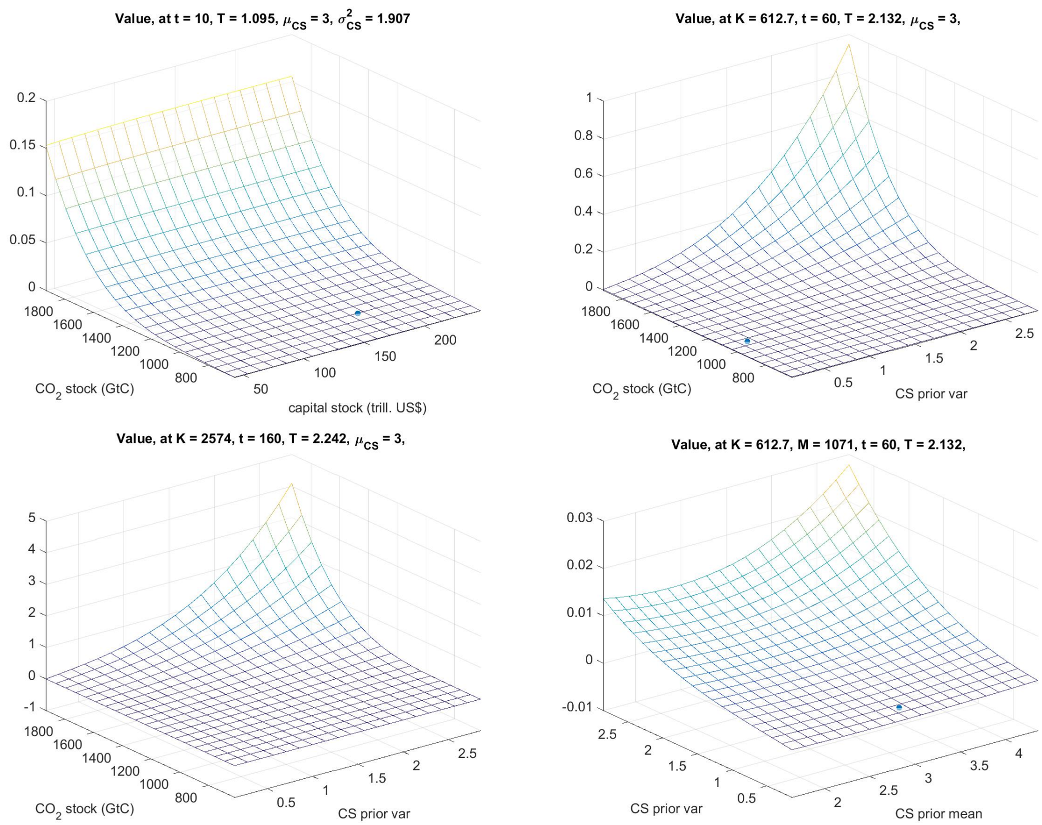


Figure 7: Relative difference in value function in percent for decision makers with and without ambiguity aversion RAA= 10 for cubic damages and temperature stochasticity $\sigma_\epsilon = .11$. Plotted along optimal time path. Four state space variables are kept fixed, two are plotted.

In addition, our evaluation assumes moderately strong ambiguity aversion as compared to a more extreme form of ambiguity aversion where the planner only weighs the worst possible (expected) outcome.

In this framework, We show how smooth ambiguity attitude implies two modifications of the optimal carbon pricing formula. First, ambiguity aversion implies pessimism weighting. Welfare in the high climate sensitivity scenarios receives more attention than welfare in the scenarios with a low climate sensitivity. As a result, the pessimism weighting effect increases the incentive to mitigate today. Second, ambiguity prudence implies a tendency to save for the future through mitigation and standard savings in order to be "less affected" by the future deep uncertainty. This effect is similar to the implications of prudence in the precau-

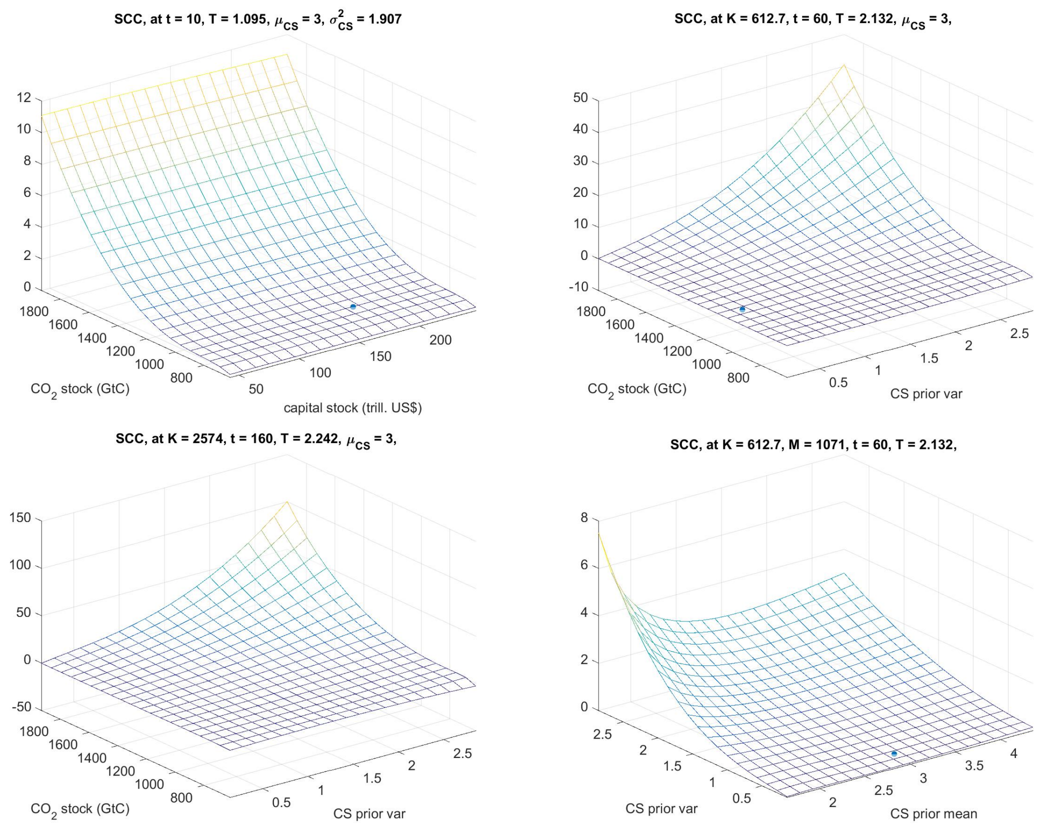


Figure 8: Relative difference in social cost of carbon in percent for decision makers with and without ambiguity aversion RAA= 10 for cubic damages and temperature stochasticity $\sigma_\epsilon = .11$. Plotted along optimal time path. Four state space variables are kept fixed, two are plotted.

tionary savings literature. Yet, in the more complex climate change context, the non-linearity between climate uncertainty and welfare shocks alters the nature of the common prudence effect.

Quantitatively, we find that the resulting carbon tax is surprisingly close to the simpler Bayesian learning model where the decision maker compounds all uncertainties into a simple best guess probability distribution. The results is encouraging in that even under deep uncertainty we should simply let the best-guess compilations of our probabilistic scientific guesstimates guide our actions. Ambiguity and ambiguity aversion add a minor premium to the social cost of carbon, which is always positive in our quantitative assessment. The results suggest that we should not shy away from using state of the art stochastic dynamic programming models for the assessment of climate change even if probability distributions governing the uncertainty are no more than the current scientific best guess.

Our model relies on optimal response to uncertainty resolution. We show that ambiguity premia can be substantially larger in cases where we deviate substantially from the expected optimal path. In particular, this finding explains why studies that merely evaluate exogenous scenarios without adaptive policy can find much larger ambiguity premia. If the climate's sensitivity to emissions turns out very high, and we do not expect future policy makers to respond to this information, then current policy makers should set a higher carbon tax. Yet, it might be questionable to suggest that present policy makers should tax carbon emissions even higher because we do not expect future policy makers with more knowledge that climate change turns out really bad further increasing the taxes.

Two clarifications remain. First, our model does not analyze fat-tails. Fat tails can deliver even larger uncertainty premia and possibly also increase the ambiguity premium. Second, variants of an alternative approach in the literature have found substantial risk premia even though these models endogenize policy. These models essentially employ a robust control approach. From a decision theoretic perspective, the common robust control approach relates closely to our limiting case of infinite ambiguity aversion. In such a model, it becomes crucial to manage the worst possible model of the world. Our model permit worlds that are far too bad to usefully employ an approach with infinite ambiguity aversion. Instead, we would have have to decide on how to limit the worlds we permit, reducing the underlying ambiguity. Such an approach follows a different philosophy and we merely emphasize that from our decision-theoretic perspective, the approaches rely on extreme ambiguity attitude rather than ambiguity per se.

# A Details on the climate enriched economy model

The following model emulates DICE The three most notable differences are the annual time step, 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 (2014), who also shows how to reformulate the decision problem when expressing capital stock and consumption in efficient labor units.

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 $\Delta T_t$ replicates the relation between oceanic and atmospheric temperatures in DICE. It follows the simple quadratic equation

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

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 2014). 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 $\mu_t$. As in the DICE model, the carbon intensity of production falls at an exogenous rate of decarbonization $\sigma_t$

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

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 $\delta_K$ 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 $L_\infty$ 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 $\Psi_t$ 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 $g_\Psi$ 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[11]

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

[11] 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

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

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

Arrow, K. & Hurwicz, L. (1972), An optimality criterion for decision-making under ignorance, *in* C. Carter & J. Ford, eds, 'Uncertainty and expectations in economics: Essays in honour of G.L.S. Shackl', Basil Blackwell, Oxford, pp. 1–11.

Barnett, M., Brock, W. & Hansen, L. P. (2020), 'Pricing Uncertainty Induced by Climate Change', *The Review of Financial Studies* **33**(3), 1024–1066.
**URL:** *https://doi.org/10.1093/rfs/hhz144*

Berger, L., Emmerling, J. & Tavoni, M. (2016), 'Managing catastrophic climate risks under model uncertainty aversion', *Management Science* **63**(3), 749 – 765.

Berger, L. & Marinacci, M. (2020), 'Model uncertainty in climate change economics: A review and proposed framework for future research', *Environmental and Resource Economics* **77**, 475 – 501.

Brock, W. & Xepapadeas, A. (2021), 'Regional climate policy under deep uncertainty: robust control and distributional concerns', *Environment and Development Economics* **26**(3), 211â€"238.

Burke, M., Craxton, M., Kolstad, C. D., Onda, C., Allcott, H., Baker, E., Barrage, L., Carson, R., Gillingham, K., Graff-Zivin, J., Greenstone, M., Hallegatte, S., Hanemann, W. M., Heal, G., Hsiang, S., Jones, B., Kelly, D. L., Kopp, R., Kotchen, M., Mendelsohn, R., Meng, K., Metcalf, G., Moreno-Cruz, J., Pindyck, R., Rose, S., Rudik, I., Stock, J. & Tol, R. S. J. (2016), 'Opportunities for advances in climate change economics', *Science* **352**(6283), 292–293.
**URL:** *https://www.science.org/doi/abs/10.1126/science.aad9634*

Gilboa, I. & Schmeidler, D. (1989), 'Maxmin expected utility with non-unique prior', *Journal of Mathematical Economics* **18**(2), 141–53.

Hansen, L. P. & Sargent, T. J. (2001), 'Robust control and model uncertainty', *American Economic Review* **91**(2), 60–66.

Jensen, S. & Traeger, C. P. (2013), Optimally climate sensitive policy under uncertainty and learning, Working paper, University of California, Berkeley.

Jensen, S. & Traeger, C. P. (2014), 'Optimal climate change mitigation under long-term growth uncertainty: Stochastic integrated assessment and analytic findings', *European Economic Review* **69**, 104–125.

Jensen, S. & Traeger, C. P. (2021), Pricing Climate Risk, CESifo Working Paper Series 9196, CESifo Group Munich.

Kelly, D. L. & Kolstad, C. D. (1999), 'Bayesian learning, growth, and pollution', *Journal of Economic Dynamics and Control* **23**(4), 491–518.
**URL:** *http://ideas.repec.org/a/eee/dyncon/v23y1999i4p491-518.html*

Keynes, J. M. (1921), *A Treatise on Probability*, Macmillan, London.

Klibanoff, P., Marinacci, M. & Mukerji, S. (2005), 'A smooth model of decision making under ambiguity', *Econometrica* **73**(6), 1849–1892.

Klibanoff, P., Marinacci, M. & Mukerji, S. (2009), 'Recursive smooth ambiguity preferences', *Journal of Economic Theory* **144**(3), 930 – 976.
**URL:** *http://www.sciencedirect.com/science/article/pii/S0022053108001749*

Lange, A. & Treich, N. (2008), 'Uncertainty, learning and ambiguity in economic models on climate policy: some classical results and new directions', *Climatic Change* **89**(1), 7–21.
**URL:** *http://dx.doi.org/10.1007/s10584-008-9401-5*

Leach, A. J. (2007), 'The climate change learning curve', *Journal of Economic Dynamics and Control* **31**(5), 1728–1752.
**URL:** *http://ideas.repec.org/a/eee/dyncon/v31y2007i5p1728-1752.html*

Lemoine, D. M. & Traeger, C. P. (2016), 'Ambiguous tipping points', *Jounral of Economic Behaviour & Organization* **132B**, 5–18.

Machina, M. J. & Siniscalchi, M. (2014), Ambiguity and ambiguity aversion, *in* M. J. Machina & K. Viscusi, eds, 'Handbook of the Economics of Risk and Uncertainty', Vol. 1, ELSEVIER SCIENCE BV, chapter 13, pp. 729–807.

Marinacci, M. (2015), 'Model uncertainty', *Journal of the European Economic Association* **13**(6), 1022–1100.

Masson-Delmotte, V., Zhai, P., Pirani, A., Connors, S., Pean, C., Berger, S., Caud, N., Chen, Y., Goldfarb, L., Gomis, M., Huang, M., Leitzell, K., Lonnoy, E., Matthews, J., Maycock, T., Waterfield, T., Yelekci, O., Yu, R. & Zhou, B., eds (2021), *Climate Change 2021: The Physical Science Basis. Contribution of Working Group I to the Sixth Assessment Report of the Intergovernmental Panel on Climate Change*, Cambridge University Press, Cambridge, United Kingdom and New York, NY, USA.

Millner, A., Dietz, S. & Heal, G. (2013), 'Scientific ambiguity and climate policy', *Environmental and Resource Economics* **55**(1), 21–46.
**URL:** *http://dx.doi.org/10.1007/s10640-012-9612-0*

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

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

Nordhaus, W. D. (2017), 'Revisiting the social cost of carbon', *Proceedings of the National Academy of Sciences* **114**(7), 1518–1523.

Olijslagers, S. & van Wijnbergen, S. (2019), Discounting the Future: on Climate Change, Ambiguity Aversion and Epstein-Zin Preferences, Tinbergen Institute Discussion Papers 19-030/VI, Tinbergen Institute.
**URL:** *https://ideas.repec.org/p/tin/wpaper/20190030.html*

Rudik, I. (2020), 'Optimal climate policy when damages are unknown', *American Economic Journal: Economic Policy* **12**(2), 340–73.
**URL:** *https://www.aeaweb.org/articles?id=10.1257/pol.20160541*

Stern, N. (2008), 'The economics of climate change', *American Economic Review P&P* **98**(2), 1 – 37.

Stocker, T., Qin, D., Plattner, G.-K., Tignor, M., Allen, S., Boschung, J., Nauels, A., Xia, Y., Bex, V. & Midgley, P., eds (2013), *Technical Summary*, Cambridge University Press, Cambridge, United Kingdom and New York, NY, USA, book section TS.
**URL:** *www.climatechange2013.org*

Traeger, C. P. (2010), 'Capturing intrinsic risk attitude', *CUDARE Working Paper No. 1102, 2010, Updated RUD 2016* .
**URL:** *https://drive.google.com/file/d/0B*$_M e6IvRR1faSTlnMGt4eXctMGs/view$

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

Traeger, C. P. (2014), 'A 4-stated dice: Quantitatively addressing uncertainty effects in climate change', *Environmental and Resource Economics* **59**(1), 1–37.
**URL:** *http://dx.doi.org/10.1007/s10640-014-9776-x*

von Neumann, J. & Morgenstern, O. (1944), *Theory of Games and Economic Behaviour*, Princeton University Press, Princeton.

Xepapadeas, A. & Yannacopoulos, A. (2018), Spatially Structured Deep Uncertainty, Robust Control, and Climate Change Policies, DEOS Working Papers 1807, Athens University of Economics and Business.
**URL:** *https://ideas.repec.org/p/aue/wpaper/1807.html*