Environmental and Resource Economics (2024) 87:683–730
https://doi.org/10.1007/s10640-023-00832-z

Discounting the Future: On Climate Change, Ambiguity Aversion and Epstein–Zin Preferences

Stan Olijslagers[1] · **Sweder van Wijnbergen**[2]

Accepted: 20 December 2023 / Published online: 26 January 2024

Abstract
We show that ambiguity aversion and deviations from standard expected time separable utility have a major impact on estimates of the willingness to pay to avoid future climate change risk. We propose a relatively standard integrated climate/economy model but add stochastic climate disasters. The model yields closed form solutions up to solving an integral, and therefore does not suffer from the curse of dimensionality of most numerical climate/economy models. We analyze the impact of substitution preferences, risk aversion (known probabilities), and ambiguity aversion (unknown probabilities) on the social cost of carbon. Introducing ambiguity aversion leads to two offsetting effects on the social cost of carbon: a positive direct effect and a negative effect through discounting. Our numerical results show that for reasonable calibrations, the direct effect dominates the discount rate impact, so ambiguity aversion gives substantially higher estimates of the social cost of carbon.

Keywords Social cost of carbon · Ambiguity aversion · Epstein–Zin preferences · Climate change

PACS Q51 · Q54 · G12 · G13

1 Introduction

Climate change is one of the main risks the world will face in the upcoming decades or possibly even centuries. But although climate scientists agree on the fact that climate change will most likely have dramatic negative consequences for the environment and economic growth, there is still much uncertainty surrounding the extent and timing of future damages induced by climate change (cf IPCC (2021) for a recent assessment). Despite all the uncertainty about the timing and the exact structure and extent of the damages that

✉ Stan Olijslagers
s.w.j.olijslagers@cpb.nl

Sweder van Wijnbergen
s.j.g.vanwijnbergen@uva.nl

[1] CPB Netherlands Bureau for Economic Policy Analysis, The Hague, The Netherlands

[2] Tinbergen Institute, CEPR, University of Amsterdam, Amsterdam, The Netherlands

climate change will cause, we do know that if they are to be avoided policies need to be implemented today. This should place the issue of how to discount future uncertain cost of climate change back towards today to allow comparison to the costs of today's policy interventions, at the center-stage of the climate change debate. And hence the subject of this paper, on climate change, risk, ambiguity aversion and Epstein–Zin preferences, and what it all implies for the Social Cost of Carbon (SCC).

We focus on the fact that we often do not know the exact parameters of the climate model and cannot even assign probabilities to their possible values. There is a fast growing literature dealing with uncertainty and risk aversion but that literature starts from the assumption that we can in fact assign probabilities to specific possible realizations, i.e. that literature deals with risk, not with ambiguity, or, in the words of Knight (1921), fundamental uncertainty. Since so little is known about the exact distribution of uncertain climate shocks, we focus predominantly on ambiguity aversion under different assumptions about preferences, and its impact on the SCC. A key result we obtain is that ambiguity aversion leads to estimates of the social cost of carbon that are substantially higher than the estimates one gets with risk rather than ambiguity (and much higher than what comes out without uncertainty at all).

The impact of climate change on the economy is most commonly modeled using combined economy/climate models called Integrated Assessment Models (IAMs). IAMs integrate the knowledge of different domains into one model. In the case of climate change, IAMs combine an economic model with a climate model. Three well-known IAMs are DICE (Nordhaus, 2014), PAGE (Hope, 2006) and FUND (Tol, 2002).[1] These models are, among others, used as policy tools for cost-benefit analyses. They provide a conceptual framework to better understand the complex problem of climate change by combining different fields and allowing for feedback effects between those fields.

But IAMs also have major drawbacks. To quote Pindyck (2017): "*IAM-based analyses of climate policy create a perception of knowledge and precision that is illusory...*" His critique is that the models are (1) very sensitive to the choices of parameters and functional forms, especially the discount rate. Besides, we know very little about (2) climate sensitivity and (3) damage functions. Lastly, (4) IAMs don't incorporate tail risk. He recommends simplifying the problem by focusing on the catastrophic outcomes of climate change, instead of modeling the underlying causes. In line with that view we focus on disaster risk by modeling climate damages as disasters (Poisson shocks). And our focus on ambiguity aversion naturally follows from his observation that we know very little about the precise stochastics of climate disasters.

We model emissions, atmospheric carbon concentration and the temperature anomaly. In this setup we model climate risk as disaster risk instead of assuming that the damage that temperature increases generate occurs gradually. Climate disasters are events that occur rarely and take place abruptly (Goosse, 2015). To model this feature, we add a jump process to the endowment consumption stream to capture climate disaster risk. The arrival rate and the intensity of the disasters is assumed to be increasing in temperature.

A critical link between climate and the economy is conventionally modeled by postulating a *damage function*, that summarizes the reduction in output resulting from the occurrence of a climate disaster. Since there is so little known about the damage functions, we investigate the impact of both attitudes towards well defined measurable risks and

[1] The references do not contain the most recent versions of the IAMs.

ambiguity aversion towards unmeasurable uncertainty on the willingness to pay for avoiding climate risk. We thus assume that the agent does not know the exact probability distributions of the arrival rate of climate disasters and the size of the disasters: there is so called ambiguity about the characteristics of the jump risk component. And the agent is assumed to be averse to this ambiguity or Knightian uncertainty.

Finally we use the continuous time version of Epstein–Zin utility, which allows us to separate the elasticity of intertemporal substitution (EIS) from the degree of risk aversion γ. In the widely used power utility specification, risk aversion and elasticity of intertemporal substitution are captured by one parameter: under power utility, they are equal to each other's inverse. There is strong empirical evidence placing the relative degree of risk aversion in the range of 5–10 (Cochrane 2009). Using such estimates in combination with power utility then results in implied estimates for the EIS much lower than direct empirical estimates suggest. But especially for long term problems such as climate change intertemporal choices play an important role and restricting parameters such as the EIS is a severe limitation. Epstein–Zin preferences make it possible to separate risk aversion and the elasticity of intertemporal substitution. We can therefore disentangle risk aversion effects (known probabilities), ambiguity aversion effects (unknown probabilities) and substitution effects. Our main focus is on the interaction of other preference parameters such as risk aversion and the EIS with ambiguity aversion.

We explicitly consider the valuation of climate risk assuming the Business As Usual (BAU) scenario: we do not analyse optimal abatement policies in this paper. Optimal policy is integrated into the analysis in a companion paper (Olijslagers et al., 2023). The rationale for this is that an analysis of the environmental costs of current policies (not current plans...) is useful in the climate policy debate. The BAU scenario is also the default scenario to calculate the social cost of carbon in Nordhaus (2014).[2] The social cost of carbon using a baseline scenario can be interpreted as the monetized welfare loss of emitting one additional unit of carbon today, given the current global carbon abatement policy scenario under the assumption that no measures will be taken in the future either.

In the first part of the paper we provide analytical solutions. To make that possible we simplify the model by assuming that the economy is a pure exchange economy with exogenous stochastic endowments, since linking stochastic emissions to the stochastics of output and consumption processes precludes analytical solutions. Of course assuming an exogenous emissions stream is unrealistic, we therefore endogenize emissions in the second part of the paper where we conventionally link emissions to output and use numerical methods to solve the model. We also demonstrate that the analytical solutions, while derived under restrictive conditions on the nature of the emissions process, do provide valuable insights that help in interpreting the outcomes of the numerical analysis of the more complex but more realistic model of part two of the paper.

Similar to the literature, the SCC in our model is very sensitive to the choice of the input parameters. But because we have analytical solutions, we can easily explore the implications of parameter choices on ambiguity aversion in combination with risk aversion and the elasticity of intertemporal substitution (EIS). We show that including ambiguity into the analysis gives two offsetting effects. On the one hand, ambiguity aversion increases the willingness to pay to avoid climate risk, which pushes up the SCC. On the other hand, ambiguity aversion also increases the risk premium and therefore the discount rate, which

[2] Note that not incorporating optimal abatement policies implies that the social cost of carbon derived here in our model is not equal to the globally optimal Pigouvian carbon tax.

lowers the SCC. Our numerical example using best estimates of the various parameters shows that the first effect is dominating. The impact of higher risk premia is more than offset by the impact of ambiguity aversion on the certainty equivalence estimates that are subsequently discounted back to today. Introducing ambiguity aversion in our model yields a SCC that is between 65% and 83% higher depending on the structure of climate risk. The net impact of higher aversion to ambiguity is to substantially raise the SCC, in particular for the realistic case of stochastic emissions which we analyse in the numerical solutions section.

Moreover, we highlight that the social cost of carbon is also sensitive to choices about time discounting, either via the pure rate of time preference, risk aversion or the elasticity of intertemporal substitution, and that all these parameters interact with the cost of ambiguity aversion. But the overall conclusion remains: insufficient attention to ambiguity leads to substantial underestimation of the SCC.

The plan of the paper is as follows: After the introduction (Sect. 1) we discuss related literature in Sect. 2 and introduce the basic model in Sect. 3. In Sect. 3.1 we focus on the endowment process and in Sect. 3.2 on modeling climate change and its economic impact. Section 3.3 focuses on preference structure and on the consequences of assuming Epstein–Zin preferences. In Sect. 3.4 we outline our approach to ambiguity Aversion and what that implies in the current model (Sect. 3.5). In Sect. 4 we present our analytical results on discount rates, the social cost of carbon and ambiguity aversion. Section 5 switches to the use of numerical solution methods; we first calibrate our model (Sect. 5.1), and use the calibrated version to illustrate our analytical results quantitatively, still assuming deterministic emissions (Sect. 5.2). In Sect. 5.3 we analyse the full model with stochastic emissions. Section 5.4 considers an extension with ambiguity about the climate sensitivity. Section 6 concludes.

2 Related Literature

This paper is related to two strands of literature. First, several other papers have also focused on obtaining analytical results to improve the understanding of what drives the social cost of carbon. Second, the paper is related to the climate economics literature that takes into account uncertainty.

Most climate economy models are solved using numerical methods. However, since it has become clear that the choice of the input parameters has a large influence on the results, we think it is useful to know how exactly these parameters influence the outcomes and therefore opt for models that allow for analytical solutions in most of our paper. There are several papers that also focused on obtaining analytic solutions. Golosov et al. (2014) were the first to obtain closed form solutions in an IAM. However, this required quite strict assumptions such as logarithmic utility and full depreciation of capital every decade. Bretscher and Vinogradova (2018) develop a stylized production-based model where the current carbon concentration directly enters the damage function and obtain closed form solutions for the optimal abatement policy. Van den Bremer and Van der Ploeg (2021) consider a rich stochastic production-based model with Epstein–Zin preferences, convex damages, uncertainty in state variables, correlated risks and skewed distributions to capture climate feedbacks. Since the model is too complex to obtain exact analytic solutions, they obtain closed form approximate solutions using perturbation methods.

Lastly, Traeger (2023) extends the model of Golosov et al. (2014). Where in other analytic other models the atmospheric carbon concentration often directly enters the damage function (Golosov et al., 2014; Bretscher and Vinogradova, 2018), Traeger (2023) explicitly models the carbon cycle and the temperature anomaly while damages are induced by an increasing temperature. In an accompanying paper, Traeger (2021) focuses specifically on uncertainty. He finds that the SCC's risk premium is almost 50% in the baseline calibration. Our SCC formula mainly differs from Traeger (2021) because we do not assume a unitary elasticity of intertemporal substitution. When the $EIS = 1$, the effective discount rate equals the pure rate of time preference and uncertainty does have no impact on the discount rate. In our view the effect of preferences, risk and ambiguity on discount rates is important and this is one of the mechanisms that we focus on. Our assumption of an $EIS \neq 1$ is crucial to study this mechanism.

Uncertainty is not part of the best-known integrated assessment model, the DICE model (Nordhaus 2017). This model is still deterministic and the representative agent is assumed to have power utility. But several papers have recently studied the impact of risk and more complex preference structures on the social cost of carbon. Cai (2020) gives an overview of the climate economics literature that studies different types of uncertainty.

Barro (2015) extends his economic disaster risk model with environmental disasters. He shows that economic disaster risk affects optimal environmental investment through discount rates. He does not incorporate a climate model but rather assumes that the disaster probability is constant and that it can be reduced by environmental investment. Jensen and Traeger (2014) add long-run economic risk and Epstein–Zin preferences to an integrated assessment model and show that assumptions about economic risk and preferences are important for the social cost of carbon through the discounting channel.

In addition to economic risk, there are also several studies that look at the effect of climate risk on the social cost of carbon and climate policy. Lemoine and Traeger (2014) and Cai et al. (2016) both study different types of tipping points and show that these tipping points lead to a higher ex-ante social cost of carbon. Karydas and Xepapadeas (2019) consider a dynamic asset pricing framework with both macroeconomic disasters and climate change related disasters and analyze the implications for portfolio allocation. Our approach differs from these papers by focusing on the effect of ambiguity aversion on the social cost of carbon.

Bansal et al. (2019) also develop a theoretical model with climate risk in the form of climate disasters, similar to our specification of climate damages. The predictions of their theoretical model are then used to test whether climate risk is already priced into asset markets. They show that long-run temperature fluctuations have a positive risk premium in equity markets. Additionally, they use information embedded in asset valuations to obtain a semi-parametric estimate of the welfare cost of carbon emissions. In contrary, we obtain estimates of the social cost of carbon using different utility specifications and using ranges of utility parameters from the literature.

Cai and Lontzek (2019) and Hambel et al. (2021) are two important contributions to the literature on uncertainty and the SCC. Both studies include Epstein–Zin preferences and calculate the SCC for a large range of combinations of risk aversion and the elasticity of intertemporal substitution. Cai and Lontzek (2019) conclude that taking into account economic risks shows that the future SCC is highly uncertain and that adding tipping points lead to a significantly higher social cost of carbon. The use of the stochastic discount factor in the valuation of climate damages is extensively discussed. Hambel et al. (2021) study uncertainty in the carbon concentration, temperate distribution and GDP and additionally consider different damage specifications and several parameter combinations. They conclude that the effect of the degree of risk aversion is modest with most damage specifications, but becomes important when damages are assumed to be large. And the interaction between different sources of risk is highlighted. Our results complement the findings from these numerical models by providing additional insights due to the closed form solutions. We for example show that risk aversion has two offsetting effects, which can be decomposed using the SCC formula. This gives a better understanding of the mechanisms that are mentioned in these papers.

Closest to our paper is the literature that considers the effect of ambiguity or deep uncertainty on the social cost of carbon and optimal climate policy. Berger and Bosetti (2020) conclude that policymakers are generally ambiguity averse using a field experiment at the COP21 conference in Paris. Brock and Hansen (2017) discuss the importance and the challenges of including different types of uncertainty in climate models. Rudik (2020) studies uncertainty, learning and concerns for misspecification in an integrated assessment model. He finds that damage learning can deliver large welfare gains. Millner et al. (2013) study the effect of ambiguity in the climate sensitivity on optimal abatement policies and conclude that ambiguity aversion could have a large effect on optimal abatement.

In contrary to our findings, Lemoine and Traeger (2016) find only a marginal impact of ambiguity aversion on the social cost of carbon. They consider ambiguity with respect to which climate regime will prevail from an uncertain arrival date onwards. The different outcome is probably caused by two differences: 1) we allow for Epstein–Zin preferences instead of power utility and 2) in our paper ambiguity directly prevails about the size of climate damages, while in Lemoine and Traeger (2016) it affects damages in a more indirect way, through the climate system.

Barnett et al. (2020) propose a model framework with three components of uncertainty in a climate economy model, namely risk, ambiguity and model misspecification. Damages enter the utility function directly. In their stylized numerical example, to simplify computations, the model misspecification channel is shut down. The analysis shows that uncertainty can be of first order importance in a climate change setting. Barnett et al. (2022) build further on the analysis of Barnett et al. (2020). They consider a much less stylized setting with three types of continuous shocks (climate shocks, damage shocks and technology shocks). Additionally an uncertain Poisson event is introduced that changes the damage function steepness. It is shown how different forms of uncertainty contribute to the SCC in the presence of ambiguity and model misspecification.

Our approach differs at several points from these two contributions. First, both papers use the ambiguity approach developed in Hansen and Miao (2018), which extends the smooth ambiguity model Klibanoff et al. (2005) by allowing for model misspecification. Instead, we use the ambiguity specification of Chen and Epstein (2002), which is a recursive extension of the Gilboa-Schmeidler max-min utility. Second, our damages are modeled as climate disasters instead of continuous damages with a possible Poisson event that changes the coefficient of the continuous damages. Lastly, because of the assumption of a unitary EIS, these papers do not focus on the effect of uncertainty on the discount rate. Using our analytical solution and non-unitary EIS, we are able to disentangle a direct effect of ambiguity aversion on the SCC given discount rates and an indirect effect through its impact on the discount rate.

3 The Model

We extend a standard endowment economy by assuming that the stochastic endowment stream is subject to climate disasters, where the probability of a climate disaster depends on the temperature level. An endowment economy is in our view a suitable starting point given our focus on the social cost of carbon and the way it depends on uncertainty and ambiguity for given policies. In particular we analyse the SCC in Nordhaus' Business As Usual scenario. In a companion paper (Olijslagers et al., 2023) we endogenize abatement policy and analyse the price of carbon under optimal abatement policies and different objective functions.

3.1 The Economy

The aggregate endowment process follows a geometric Brownian motion with an additional jump component that represents climate disasters:[3]

$$dC_t = \mu C_t dt + \sigma C_t dZ_t + J_t C_{t-} dN_t. \tag{1}$$

In equilibrium, aggregate consumption must equal the aggregate endowment and therefore we also refer to the process as the aggregate consumption process. The growth rate μ and the volatility σ are constant. Z_t is a standard Brownian motion that captures 'normal' uncertainty, i.e. non-climate uncertainty. N_t is a Poisson process which represents climate disasters. The arrival rate of a climate disaster equals λ_t, which we assume to be a function the temperature level T_t. Note that this way of modeling climate disasters implies that disasters are proportional to consumption.

When a climate disaster strikes at time t, the size of the disaster is controlled by the random variable J_t. The distribution of the size of disasters is assumed to be the same for any t. We assume that J_t has the density $f(x) = \eta(1+x)^{\eta-1}$ where $-1 < x < 0$. J_t represents the percentage loss of aggregate consumption after a disaster. The expected disaster size then equals $E[J_t] = \frac{-1}{\eta+1}$ and the moments $E\left[(1+J_t)^n\right] = \frac{\eta}{\eta+n}$ can be easily calculated. In line with the subject of climate disasters, jumps can only be negative.

3.2 The Climate Model

The arrival rate of disasters is assumed to be temperature dependent. We assume that damages are linearly increasing in temperature: $\lambda_t = \lambda_T T_t$. We make this simplifying assumption because we want to focus on other non-linearities in the SCC. However, all our derivations remain valid for convex specifications of the arrival rate. We discuss this assumption in more detail in the calibration section.

In the first part of the paper we make a number of simplifying assumptions to allow for analytic solution of the model. The main solvability requirement is that the state variables of the climate submodel are deterministic, and this in particular affects the way we model emissions. Carbon emissions are the product of the carbon intensity of aggregate output and aggregate output itself. We will introduce emissions in this way in the numerical part of the paper, but doing so precludes analytical solution. So our main simplification in the

[3] C_{t-} denotes aggregate endowment just before a jump ($C_{t-} = \lim_{h\downarrow 0} C_{t-h}$).

analytical part of the paper is the assumption that aggregate emissions are driven by an independent deterministic process, an unavoidable simplification if one is to obtain analytical solutions. In Sect. 5.3 we use numerical methods and introduce stochastic emissions correlated with output. Making emissions stochastic clearly adds realism and enriches the results, but interpreting the numerical results benefits substantially from the additional insights obtained from the analytical results obtained earlier.

So assume for the first part of the paper that emissions E_t are exogenous. E_t is growing at a non-stochastic rate $g_{E,t}$. The growth rate itself moves gradually towards the long-run equilibrium $g_{E,\infty}$ at a rate δ_E. By assuming a high initial growth rate but a negative long run rate ($g_{E,\infty} < 0$), we have growing emissions today; but the growth rate starts declining immediately and eventually turns negative because of $g_{E,\infty} < 0$, so emissions will go to zero eventually. This is a plausible assumption since there is a point where the stock of fossil fuels will be depleted. All this leads to the following process for emissions:

$$\begin{aligned} dE_t &= g_{E,t}E_t dt, \\ dg_{E,t} &= \delta_E(g_{E,\infty} - g_{E,t})dt. \end{aligned} \tag{2}$$

We calibrate this process to match the baseline scenario in Nordhaus (2017).

We use the climate model (carbon cycle and temperature model) discussed in Mattauch et al. (2018), which they call the IPCC AR5 impulse-response model. This model is in line with recent insights from the climate literature and is also used in IPCC (2013). Specifically, this climate model incorporates the fact that thermal inertia plays a smaller role than commonly assumed in the climate modules in economic models. Climate modules commonly used in economic models tend to overstate the time it takes for the earth to warm in response to carbon emissions (cf Dietz et al. 2021).

Define by M_t the carbon concentration with respect to pre-industrial emissions M_{pre}. In our model, M_t is the sum of four artificial carbon boxes: $M_t = \sum_{i=0}^{3} M_{i,t}$. This specification can capture that the decay of carbon has multiple time scales and that a fraction of emissions will stay in the atmosphere forever. The dynamics of carbon box i are given by:

$$dM_{i,t} = v_i\Big(E_t - \delta_{M,i}M_{i,t}\Big)dt. \tag{3}$$

v_i is the fraction of emissions that ends up in carbon box i, which implies that $\sum_{i=0}^{3} v_i = 1$. $\delta_{M,i}$ controls the decay rate of carbon in box i. We assume that all carbon that ends up in box 0 will permanently stay in the atmosphere, such that $\delta_{M,0} = 0$. The other three boxes have a positive decay rate: $\delta_{M,i} > 0, i = \{1, 2, 3\}$.

The next step is to model the impact of carbon concentration on temperature. This requires modeling what is called radiative forcing: the difference between energy absorbed by the earth from sunlight and the energy that is radiated back to space. A higher atmospheric carbon concentration strengthens the greenhouse effect and therefore leads to higher radiative forcing. The relation between atmospheric carbon concentration and radiative forcing is logarithmic:

$$F_{M,t} = \alpha \frac{\upsilon}{\log(2)} \log\left(\frac{M_t + M_{pre}}{M_{pre}}\right). \tag{4}$$

α equals the climate sensitivity: the long-run change in temperature due to a doubling of the carbon concentration compared to the pre-industrial level. υ is a parameter that is also part of the temperature module and this parameter will be discussed later.

Finally we also include non-carbon related (exogenous) forcing $F_{E,t}$, which follows:

$$dF_{E,t} = \delta_F(F_{E,\infty} - F_{E,t})dt. \tag{5}$$

Total radiative forcing is the sum of carbon-related radiative forcing and exogenous forcing: $F_t = F_{M,t} + F_{E,t}$.

The final step moves from F_t to the actual surface temperature T_t. T_t is the difference between the actual temperature compared to the pre-industrial temperature level. The change in surface temperature is a delayed response to radiative forcing. Call the heat capacity of the surface and the upper layers of the ocean τ while τ_{oc} equals the heat capacity of the deeper layers of the ocean. The parameter κ captures the speed of temperature transfer between the upper layers and the deep layers of the ocean. The dynamics of temperature are then given by:

$$\begin{aligned} dT_t &= \frac{1}{\tau}\Big(F_t - \upsilon T_t - \kappa(T_t - T_t^{oc})\Big)dt, \\ dT_t^{oc} &= \frac{\kappa}{\tau_{oc}}(T_t - T_t^{oc})dt. \end{aligned} \tag{6}$$

From this equation, one can derive a long run equilibrium temperature for a given level of radiative forcing F_t:

$$T_t^{eq} = \frac{F_t}{\upsilon} \tag{7}$$

The parameter υ controls the equilibrium temperature response to a given level of forcing. Note that Eq. (4) tells us that when $M_t = 2M_{pre}$, we get that $F_t = \alpha\upsilon + F_{E,t}$ and $T_t^{eq} = \alpha + \frac{F_{E,t}}{\upsilon}$. Therefore the parameter α can indeed be interpreted as the equilibrium temperature response to doubling of the carbon concentration.

Using Eq. (7), we can rewrite the first line of Eq. (6) as:

$$dT_t = \frac{1}{\tau}\Big(\upsilon(T_t^{eq} - T_t) - \kappa(T_t - T_t^{oc})\Big). \tag{8}$$

Written this way the equation is more intuitive, since it captures the fact that the temperature moves towards its equilibrium level at a rate proportional to $T_t^{eq} - T_t$. The second part shows that the oceans are delaying this convergence. It takes time for T_t^{oc} to adjust towards T_t and this will also delay the convergence of T_t towards the equilibrium level T_t^{eq}. As specified earlier, the arrival rate of climate disasters is a linear function of temperature T_t.

3.3 Preference Specification

The representative agent maximizes utility of consumption over an infinite planning horizon. Because of the different roles played by intertemporal substitution and risk aversion in determining risk premia, the safe of interest, and therefore also the social cost of carbon, we use Epstein–Zin (EZ) preferences (Epstein and Zin 1989); EZ preferences allow us to vary the elasticity of intertemporal substitution (EIS) and the coefficient of relative risk aversion independently. This is important since both the elasticity of intertemporal substitution (through the discount rate) and risk aversion interact with ambiguity aversion in our framework. We use the continuous time version of of Epstein–Zin utility, a special case of stochastic differential utility introduced by Duffie and Epstein (1992b).

The agent's utility or value function is:

$$V_t = E_t\Big[\int_t^\infty f(C_s, V_s)ds\Big]$$

where

$$f(C,V) = \frac{\beta}{1-1/\epsilon}\frac{C^{1-1/\epsilon} - \big((1-\gamma)V\big)^{\frac{1}{\zeta}}}{\big((1-\gamma)V\big)^{\frac{1}{\zeta}-1}} \quad \text{for } \epsilon \neq 1 \tag{9}$$

with $\zeta = \dfrac{1-\gamma}{1-1/\epsilon}$.

γ denotes risk-aversion, ϵ is the elasticity of intertemporal substitution and β equals the time preference parameter. We will focus on the more general case where $\epsilon \neq 1$. For the case $\epsilon = 1$ one can take the limit $\epsilon \to 1$ or follow the same derivation but with $f(C,V) = \beta(1-\gamma)V\Big(\log C - \frac{1}{1-\gamma}\log\big((1-\gamma)V\big)\Big)$. Finally if $\gamma = \frac{1}{\epsilon}$, the utility specification reduces to standard power utility.

3.4 Ambiguity

There is much uncertainty regarding the arrival rate and magnitude of climate disasters. Pindyck (2017) already stresses that we know very little about the damage functions. And where consumption growth and volatility can be estimated accurately from historical data, the estimation of the climate disaster parameters will be much harder since climate disasters do not happen that often. It is fair to state that we simply do not know the exact distribution of climate damages. We should therefore account for the possibility that the 'best estimate' model is not the true model: there is ambiguity. We assume that the representative agent is ambiguity averse.

It is important to stress the difference between risk and ambiguity. When we are talking about risk, an agent knows the probabilities and possible outcomes of all events. When the agent has to deal with ambiguity, the probabilities attached to particular events are unknown. The distinction between risk and ambiguity is already extensively discussed in Knight (1921), which is why ambiguity is often referred to as Knightian uncertainty. Ellsberg (1961) shows using the Ellsberg Paradox that people are ambiguity averse, i.e. they prefer known probabilities over unknown probabilities. We follow Chen and Epstein (2002) by introducing uncertainty over probability measures by introducing a set of 'plausible' models $\mathcal{P}$; subjective beliefs over this set can be summarized by a subjective probability μ.

We use the *recursive multiple priors utility* developed in continuous time by Chen and Epstein (2002) to model ambiguity aversion. This method selects a discrete set of models that are relatively similar. The size of the set depends on the degree of ambiguity aversion. In the language of Hansen and Sargent (2001), the decision maker has to make a robust decision given the set of reasonable models. In the literature two approaches to that decision problem are used, the *smooth ambiguity* approach proposed by Klibanoff et al. (2005) and the *Maximin* approach advocated among others by Gilboa and Schmeidler (1989) and used in Chen and Epstein (2002).

To apply the Maximin approach of Chen and Epstein (2002) to modeling ambiguity aversion we begin by defining the 'best estimate' model or reference model as the agent's most reliable model, with probability measure $\mathbb{P}$. But the agent takes into account that his reference model may not be the true model and specifies a set of models $\mathcal{P}^{\theta}$ that he also considers possible. The alternative models have measure $\mathbb{Q}^{a,b}$; the jump arrival rate becomes $\lambda_t^{\mathbb{Q}^{a,b}} = a\lambda_t$ and the jump size parameter becomes $\eta_t^{\mathbb{Q}^{a,b}} = b\eta$. Remember that the expected jump size equals $\frac{-1}{\eta+1}$, i.e. a low b leads to a more negative jump size. Given the set of models $\mathcal{P}^{\theta}$, Chen and Epstein (2002) then assume that the agent optimizes assuming the worst case, in line with the axiomatic Minimax approach advocated by Gilboa and Schmeidler (1989). In this approach the subjective probability μ over the set $\mathcal{P}$ of all probability measures $\mathbb{Q}^{a,b}$ has no impact on the final outcome: the MiniMax approach advocated by Gilboa and Schmeidler (1989) and applied in Chen and Epstein (2002) optimizes assuming only the worst measure of the set $\mathcal{P}$ (hence the moniker Minimax).

A more general and also often used way to model ambiguity and ambiguity aversion is the the *smooth ambiguity* model proposed by (Klibanoff et al. 2005). Klibanoff et al. (2005) also start by considering a set of possible measures $\mathcal{P}$, with an associated subjective probability μ over the set $\mathcal{P}$. In particular, assume that the agent does not know the true values of λ and η. In this approach the agent first constructs a prior probability distribution μ that reflects his beliefs on λ and η. To incorporate ambiguity aversion, he then transforms this distribution using a mapping ϕ to put more weight on the events that give him low utility and less weight on the events that give high utility. This transformation ϕ summarizes the representative consumer's attitude towards uncertainty over what constitutes the "right" prior. The concavity of ϕ indicates the degree of ambiguity aversion, in the same way concavity of a standard utility function defined over different realizations (given a measure mapping those realizations to probabilities) indicates risk aversion. Ambiguity aversion adds as it were another layer to our utility specification. For a given set $\mathcal{P}$ the Maximin approach advocated by Gilboa and Schmeidler (1989) and Chen and Epstein (2002) is in fact the limiting case of the smooth ambiguity approach of Klibanoff et al. (2005) as ambiguity aversion goes from zero to infinity. We opt for the Maximin approach, which of course has the main benefits of being simple to apply and easy to interpret.

We assume that all models with a distance smaller than θ are in the set $\mathcal{P}$ of admissible models $\mathbb{Q}^{a,b}$, so the size of the set of models depends on the ambiguity aversion parameter θ; and θ can be interpreted as a measure of the extent of ambiguity. We measure distance between the reference model $\mathbb{P}$ and an alternative model $\mathbb{Q}^{a,b}$ using the concept of *relative entropy*, a common metric for the distance between two probability measures (see for example Hansen and Sargent (2008)). Relative entropy thus gives information about how similar two probability measures are. To obtain our distance measure, we scale relative entropy by the arrival rate λ_t.[4] Without this scaling, the optimal a^* and b^* would be time-varying. This would imply that the decision maker is continuously updating a^* and b^*. A constant a^* and b^* are both more intuitive and more tractable.

The distance between the reference and alternative model depends on the parameters a and b and can therefore be written as $d(a, b)$. The distance measure satisfies $d(a, b) \geq 0 \; \forall (a, b)$ and $d(1, 1) = 0$: the distance of the reference model to itself is by definition equal to 0. If θ is large, the agent is very ambiguity averse and thus considers a large set of models. The preferences of the agent then become:

[4] Liu et al. (2004) and Maenhout (2004) also use a normalisation factor to scale their distance measure in order to get tractable results. In an earlier version, which is available on request from the authors, we used an unscaled entropy measure. This adds complexity but has a negligible impact on the results.

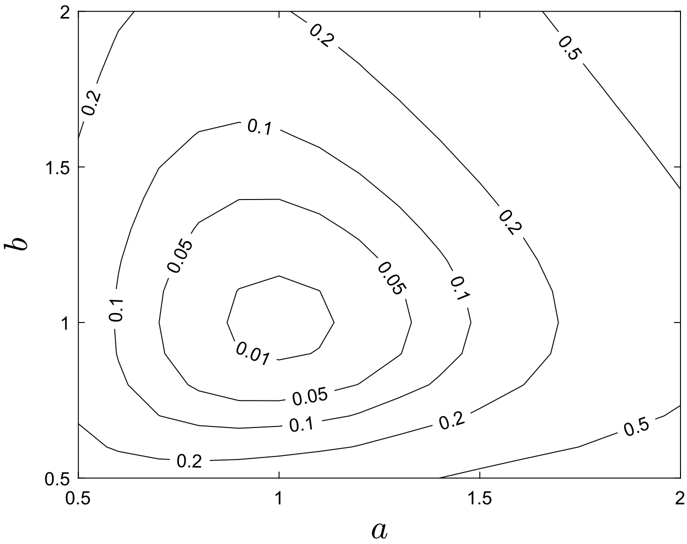

Fig. 1 Distance measure for different values of a and b

$$
\begin{aligned}
& V_t = \min_{\mathbb{Q}^{a,b} \in \mathcal{P}^\theta} V_t^{\mathbb{Q}^{a,b}} \\
& \text{where } V_t^{\mathbb{Q}^{a,b}} = E_t^{\mathbb{Q}^{a,b}}\left[\int_t^\infty f(C_s, V_s^{\mathbb{Q}^{a,b}})ds\right] \\
& \text{and } \mathcal{P}^\theta = \{\mathbb{Q}^{a,b} : d(a,b) \le \theta \ \forall t\}.
\end{aligned}
\tag{10}
$$

Here $V_t^{\mathbb{Q}^{a,b}}$ is the value function assuming $\mathbb{Q}^{a,b}$ is the true probability measure. $\theta = 0$ implies that $\mathcal{P}^\theta = \{\mathbb{P}\}$ and the agent only considers one measure, namely the reference measure. Thus there is no ambiguity aversion when $\theta = 0$. Where the risk aversion parameter γ can be seen a parameter that is relevant for any risky bet, the parameter θ captures intrinsic ambiguity aversion (one person might be more ambiguity averse than another), but it is also source dependent. If there is a lot of information and data available about a process, θ will be smaller and the set of admissible models will be smaller compared to a process about which not much is known. But at the same time θ also captures aversion to ambiguity similar to risk aversion.

In "Appendix A" we derive that the distance measure equals:

$$
d(a,b) = (1-a) + a\left(\log(ab) + \frac{1}{b} - 1\right). \tag{11}
$$

It is easy to verify that d(1,1) = 0, the distance between the reference distribution and itself is zero. When one or both of the two variables a and b deviate from the reference model, $d(a, b)$ increases. Every contour in Fig. 1 gives a set of combinations (a, b) that yields the same distance. If for example $\theta = 0.1$, then all (a, b) combinations within that contour line are included in the set of admissible models. The worst case probability measure will be the probability measure for which either a is large (high arrival rate) and/or b is small, since the expectation of the jump size under the alternative measure is inversely related to b: $E^{\mathbb{Q}^{a,b}}[J_t] = \frac{-1}{b\eta+1}$.

From the current setup, it is hard to argue what a reasonable value for ambiguity aversion θ would be. In order to give more guidance about reasonable values for θ, we use the concept of *detection error probabilities* introduced by Anderson et al. (2003).[5] Consider the following thought experiment. Assume that the representative agent would be able to observe the process of consumption over the next N years, and after observing the process the agent has to choose which of the two models (the reference model or the worst-case model) is most likely. There are two types of errors in this case. The agent could choose the reference model while the process was actually generated by the worst-case model and he could also make the opposite error. The detection error probability is defined as the average of the probability of the two errors. "Appendix B" describes how the detection error probability is calculated.

The detection error probability depends on N, since when the agent observes the process for a longer period, the probability of a mistake will be smaller. The detection error probability also depends on the ambiguity aversion parameter: when θ is small, the reference and worst-case model are similar to each other and the probability of a mistake is large. On the other hand, when the agent is extremely ambiguity averse (or there is a lot of ambiguity) the reference and worst-case models are very different and the detection error probability becomes small. The representative agent wants to make the set of models sufficiently large to make a robust decision, but on the other hand does not want to take into account implausible models. The detection error probability gives guidance about whether the set of admissible models is too small or too large. Since the detection error also depends on the other parameters of the model, we come back to the issue of calibrating the ambiguity aversion parameter in the calibration section.

3.5 Optimal *a* and *b*

As discussed before, the agent has the following utility function: $V_t = \min_{\mathbb{Q}^{a,b} \in \mathcal{P}^\theta} V_t^{\mathbb{Q}^{a,b}}$, where $\mathcal{P}^\theta$ is the set of all probability measures that satisfy the distance constraint. Since every probability measure $\mathbb{Q}^{a,b}$ is uniquely defined by the parameters a and b, minimizing over $\mathbb{Q}^{a,b}$ is equivalent to minimizing over the parameters a and b. In "Appendix C" we show that this minimization problem boils down to minimizing the following expression:

$$\min_{(a,b)} a\lambda_t \frac{-1}{b\eta + 1 - \gamma} \quad s.t. \;\; d(a,b) \leq \theta. \tag{12}$$

We start with discussing this minimization problem assuming that the agent would be risk-neutral ($\gamma = 0$) but ambiguity averse ($\theta > 0$). At every time unit, the expected loss equals the probability of a disaster times the expected disaster size. Assuming that $\mathbb{Q}^{a,b}$ is the true measure, the arrival rate becomes $a\lambda_t$ and the expected disaster size equals $\frac{-1}{b\eta+1}$. The expected loss of a climate disaster therefore equals $a\lambda_t \frac{-1}{b\eta+1}$. The agent then chooses the combination of a and b that gives the most negative expected loss while still satisfying the distance restriction $d(a,b) \leq \theta$. A higher θ allows a larger range of values for a and b that satisfy the distance constraint.

In our specification the agent is both risk averse and ambiguity averse. Instead of minimizing the expected loss, the agent minimizes the certainty equivalent of a climate

[5] See for example Maenhout (2006) for another application of detection error probabilities.

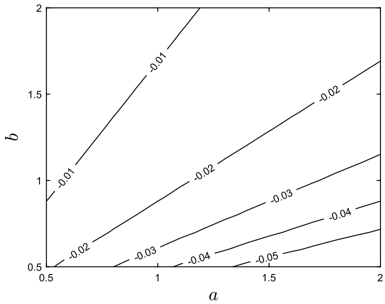

(a) Contour plot of the objective function of the constrained minimization problem for different values of a and b.

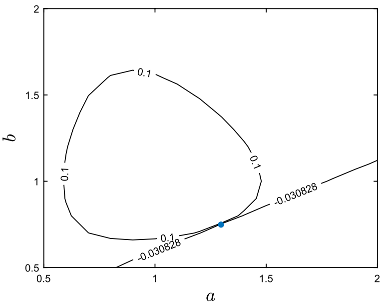

(b) Illustration of selection of optimal (a, b). The oval area shows all admissible values for a and b that are within the ambiguity budget of 0.1. The straight line is the objective function.

Fig. 2 Selection of the optimal a and b

disaster: $a\lambda_t \frac{-1}{b\eta+1-\gamma}$. The certainty equivalent is more negative than the expected loss since it contains a correction for risk aversion. The optimal parameters a^* and b^* are thus a function of ambiguity aversion θ, the jump size parameter η but also of risk aversion γ.

Figure 2 illustrates the optimization problem. Given an ambiguity budget θ, one can determine the feasible set of (a, b). Figure 1 shows the feasible sets for several budgets. A contour plot of the objective function for several (a, b) combinations is given in Fig. 2a. Clearly combinations in the bottom right corner (high a, low b) give the lowest value of the objective function. The optimization will thus lead to $a^* > 1$ and $b^* < 1$ since b^* and the disaster size are inversely related. The goal is to minimize this function, given the distance constraint. Figure 2b shows how the optimal combination (a^*, b^*) is determined. The point where objective function touches the feasible region is the optimal solution. From now on we use the following notation for the optimal arrival rate and jump parameter: $\lambda_t^* = a^*\lambda_t$ and $\eta^* = b^*\eta$.

4 Discounting and the Social Cost of Carbon: Analytical Solutions

We are now ready to address the key questions raised in the introduction, how to discount future carbon damages and what that implies for the Social Cost of Carbon (SCC). We focus first on the appropriate discount rates.

4.1 On Discounting

Consider first the risk-free rate and the risk premium; we then derive the growth-adjusted consumption discount rate, the rate at which future consumption streams (or their decline) need to be discounted towards today, which is used to discount future damages when calculating the SCC.

In "Appendix D" we derive the expressions for the interest rate, the risk premium and the consumption rate of interest from no arbitrage conditions for the valuation of respectively a safe asset B_t, aggregate wealth and a synthetic asset paying out a proportion of aggregate consumption at a specified future time. We label latter the CDR_t; it equals the return on wealth $r + rp$ minus a correction for the growth in consumption.

Consider first the expression for the safe rate of interest, derived from the no arbitrage condition of a safe asset B_t (cf "Appendix D" for the derivation details):

$$\begin{aligned} r_t = \beta + \frac{\mu}{\epsilon} - \left(1 + \frac{1}{\epsilon}\right)\frac{\gamma}{2}\sigma^2 - \left(\gamma - \frac{1}{\epsilon}\right)a^* \lambda_t \frac{-1}{b^*\eta + 1 - \gamma} \\ - a^* \lambda_t \left(\frac{b^*\eta}{b^*\eta - \gamma} - 1\right). \end{aligned} \tag{13}$$

We return to this expression below in the discussion. The relevant risk premium is the excess return on a claim on consumption, or, more precisely, a stock S_t paying out continuous dividends C_t. The value of the stock can also be interpreted as aggregate wealth, since total wealth of the representative agent is equal to the total claim on future consumption. Requiring once again the familiar no arbitrage condition gives the expression for the risk premium (again cf "Appendix D"):

$$rp_t = \gamma\sigma^2 + a^* \lambda_t \left(\frac{-1}{b^*\eta + 1} - \frac{b^*\eta}{b^*\eta + 1 - \gamma} + \frac{b^*\eta}{b^*\eta - \gamma}\right). \tag{14}$$

Without climate risk (the Poisson terms), the risk premium boils down to the well known expression: $\gamma\sigma^2$.

Finally we use the results for the safe rate of interest and the risk premium in the derivation of the expression for the growth-adjusted Consumption Discount Rate CDR_t. The consumption discount rate CDR_t is the relevant discount rate for discounting climate damages when calculating the social cost of carbon, because damages in our setup are proportional to consumption. In "Appendix D" we derive the expression for this discount rate using the results we obtained so far for r_t and rp_t. The no-arbitrage condition is applied to a synthetic asset H_t paying out a proportion of aggregate consumption at a unique time $s > t$:

$$\begin{aligned} CDR_t &= \underbrace{r_t}_{I} + \underbrace{rp_t}_{II} - \underbrace{\left(\mu + a^* \lambda_t \frac{-1}{b^*\eta + 1}\right)}_{III} \\ &= \beta + (1/\epsilon - 1)\left(\mu - \frac{\gamma}{2}\sigma^2 + a^* \lambda_t \frac{-1}{b^*\eta + 1 - \gamma}\right). \end{aligned} \tag{15}$$

CDR_t consists of three terms, labeled *I*, *II* and *III*. Part *I* is the risk-free rate. But future economic growth is uncertain, so we need to add a risk premium since damages are a fraction of the economy and thus have an impact on consumption: part *II*. Lastly, the discount rate should be corrected for the growth of the aggregate consumption process (part *III*). Future damages are larger because the future economy is larger, which is why we need to correct the discount rate for future growth. On average, consumption grows at a rate $\mu + a^* \lambda_t \frac{-1}{b^*\eta+1} < \mu$: the average growth rate is smaller than μ since climate disasters are expected to have a negative impact on consumption.

In the simplest case, without any risk at all, the risk premium is zero and the interest rate then reduces to the well-known Ramsey rule (Ramsey, 1928):

$$(\sigma, \lambda_T) = (0, 0) \;\; => \;\; r_t = \beta + \frac{\mu}{\epsilon}, \tag{15a}$$

which implies a growth corrected discount rate $r_{n,t}$ for the case of $(\sigma, \lambda_T) = (0, 0)$ equal to:

$$r_{n,t} = \beta + (1/\epsilon - 1)\mu. \tag{15b}$$

Clearly a higher value for ϵ implies a lower growth corrected discount rate: a higher willingness to substitute over time implies less discounting of the future. Adding diffusion risk ($\sigma > 0$, $\lambda_T = 0$) leads to well known results: this will both affect the safe interest rate, which falls due to a flight to safety effect, and the risk premium, which now becomes $\gamma\sigma^2$:

$$(\sigma > 0, \lambda_T = 0) => r_t = \beta + \frac{\mu}{\epsilon} - (1 + 1/\epsilon)\frac{\gamma}{2}\sigma^2, \tag{15c}$$

$$rp_t = \gamma\sigma^2. \tag{15d}$$

Adding the risk premium to the risk-free rate and again correcting for the growth rate gives the growth-adjusted discount rate, still assuming $\sigma > 0$, $\lambda_T = 0$:

$$r_{n,t} = \beta + (1/\epsilon - 1)\left(\mu - \frac{\gamma}{2}\sigma^2\right). \tag{15e}$$

So the impact on the safe rate and on the risk premium are in opposite directions, as is well known from the literature. For $\epsilon = 1$ the two effects cancel out, for $\epsilon > 1$ the risk premium impact dominates and the overall discount rate increases with risk. For $\epsilon < 1$ the opposite result is obtained and discount rates will actually go down with higher risk as the flight to safety effect dominates the impact on the risk premium. While ϵ determines the relative importance of the risk-free rate and risk premium effects, risk aversion γ determines their magnitude. A high degree of risk aversion amplifies the effect of risk on the discount rate. Of course when the agent is risk neutral ($\gamma = 0$), risk has no effect on the discount rate.

We now introduce climate uncertainty in addition to diffusion risk, and ambiguity aversion, the main topic of this paper. Adding climate disaster risk to diffusion risk implies: $\sigma > 0$ and $\lambda_T > 0$. To set a benchmark we first analyse the case where there is no ambiguity aversion ($\theta = 0$). This corresponds to $a^* = 1, b^* = 1$, the optimal and the reference case actually coincide when $\theta = 0$. Equation (15) then shows that adding climate disaster risk has an effect on both the safe interest rate and the risk premium very much like changes in σ have. The climate risk term is premultiplied by $(1/\epsilon - 1)$ in Eq. (15), so when $\epsilon < 1$ the risk free rate effect dominates and adding disasters leads to a lower discount rate. But when $\epsilon > 1$, the risk premium effect dominates and adding climate disasters actually leads to higher discount rates. Finally when $\epsilon = 1$, the two effects cancel.

In our no-ambiguity-aversion benchmark case $a^* = 1, b^* = 1$, the climate related term in Eq. (15) then becomes:

$$\lambda_t \frac{-1}{\eta + 1 - \gamma}. \tag{16}$$

$\frac{-1}{\eta+1-\gamma}$ equals the certainty equivalent of the climate shock. When $\gamma = 0$, the certainty equivalent is equal to the expected value $E_t[J_t] = \frac{-1}{\eta+1}$. The term scales with the arrival rate λ_t: more frequent disasters have a larger effect on discount rates. Finally a higher γ leads to a smaller certainty equivalent (i.e. a larger negative shock), since η is substantially larger than 1.

Now introduce ambiguity aversion. The climate term in Eq. (15) now equals:

$$a^* \lambda_t \frac{-1}{b^* \eta + 1 - \gamma}. \tag{17}$$

Including ambiguity aversion leads to a larger worst case arrival rate: $a^* > 1 => a^* \lambda_t > \lambda_t$ so one can see from Eq. (17) that ambiguity aversion leads to a larger worst case arrival rate. Thus ambiguity aversion amplifies the impact that the arrival rate of climate disasters has on discounting. Also, we can see from (17) that ambiguity aversion implies a more negative certainty equivalent term since $b^* < 1$; so once again we find that ambiguity aversion leads to a larger impact of climate risk on discounting. Therefore we can unambiguously conclude that there is AA amplification: ambiguity aversion amplifies the effect of climate risk on discounting, both through its impact on the worst case arrival state and on the worst case certainty equivalent conditional on arrival.

Whether AA amplification leads to a higher or lower discount rate depends on the value of ϵ, much like in the earlier discussion on the impact of (climate) risk on interest rates in the absence of ambiguity.

$\epsilon = 1$: the impact of AA amplification on the safe rate and on the risk premium cancel each other out and the discount rate simply becomes β irrespective of climate risk (or for that matter any other risk).

$\epsilon < 1$: the flight to safety effect of magnification dominates the impact of a higher risk premium, so AA amplification actually leads to a lower discount rate.

$\epsilon > 1$: we get the presumably more intuitive outcome, with $\epsilon > 1$ the risk premium effect dominates and AA magnification actually leads to a higher discount rate than obtained without AA magnification.

4.2 The Social Cost of Carbon

With the machinery developed sofar and using the value function from Eq. (9) we can take the next step and calculate the Social Cost of Carbon (SCC). We define the SCC as the marginal cost in terms of reduced welfare of increasing carbon emissions by one ton carbon scaled by the marginal welfare effect of one additional unit of consumption. This gives us the social cost of carbon in terms of the price of time t consumption units terms (conventionally referred to as 'in dollar terms'). In "Appendix E" we derive the following expression for the SCC based on this definition:

$$SCC_t = C_t \int_0^\infty \underbrace{\exp\left\{ - \int_t^{t+u} CDR_s ds \right\}}_{I} \underbrace{\int_t^{t+u} a^* \lambda_T \frac{\partial T_s}{\partial M_t} ds}_{II} \underbrace{\frac{1}{b^* \eta + 1 - \gamma}}_{III} du. \tag{18}$$

Equation (18) shows first of all that the social cost of carbon is proportional to C_t, the aggregate consumption level: when the current aggregate consumption level C_t doubles, the SCC doubles as well. For a given consumption level, the SCC depends on three terms, labeled *I*, *II* and *III* respectively in Eq. 18. The social cost of carbon. the marginal welfare loss due to emitting an additional unit of carbon today, is the discounted sum of all current and future damages done by emitting one ton of carbon today. The outer integral indicates that all future marginal damages are included in the SCC. Future damages are discounted with the (cumulative) consumption discount rate (term *I*). Term *II* is the change in the expected number of disasters between time t and time $t + u$ due to an additional unit of emissions today. This change in expected number of disasters is a function of the derivative of future temperature levels with respect to current carbon emissions, because the arrival rate is temperate dependent. The marginal changes in the arrival rate are integrated, because the expected number of disasters between time t and time $t+u$ equals the integral over the time-varying arrival rates.[6] Term *III* captures the damages when a disaster actually takes place. It can be interpreted as a certainty equivalent: the expected value is adjusted for risk and ambiguity preferences.

Consider first the impact of risk aversion as measured by γ. Term *III* is clearly increasing in risk aversion. But risk aversion also has an effect on the discount rate CDR_t. As discussed before, increasing risk aversion increases the discount rate when $\epsilon > 1$. So when $\epsilon > 1$ the discounting effect works in opposite direction of the effect on the certainty equivalent: for $\epsilon > 1$ the impact of γ on the SCC is therefore ambiguous in general and will depend on the specific parameter values chosen (cf the numerical analysis in Sect. 5).

Consider next the impact of ϵ. The elasticity of intertemporal substitution ϵ only plays a role in the discount rate. When ϵ increases, the willingness to substitute over time increases which leads to lower discount rates. So a higher ϵ unambiguously leads to a higher SCC.

The ambiguity aversion parameter θ does not directly show up in the formula for the SCC, but its effect works through the choice of a^* and b^*. When ambiguity aversion is present, i.e. $\theta > 0$, $a^* > 1$ (higher worst-case arrival rate) and $b^* < 1$ (more negative worst-case jump size). With $\theta > 0$, the increase in the probability of a disaster happening (term *II*) is larger because the worst case arrival rate of disasters $a^* \lambda_t$ is higher. And term *III* in expression 18, the certainty equivalent damage term conditional on a disaster happening, is also higher. So through these two channels ambiguity aversion leads to a higher social cost of carbon.

But ambiguity aversion also affects discount rates and the sign again depends on the elasticity of intertemporal subtitution ϵ. When $\epsilon < 1$, ambiguity aversion additionally leads to a lower discount rate and thus an even higher SCC. When $\epsilon = 1$, the discount rate is simply β and ambiguity has no effect on the discount rate so in that case the SCC increases with ambiguity also. Lastly, when $\epsilon > 1$, increasing θ leads to higher discount rates. Therefore increasing ambiguity aversion then has two offsetting effects in this case and the net sign of the impact of θ on the SCC is in principle ambiguous. Since ambiguity aversion always leads to a higher SCC when $\epsilon \leq 1$, we mostly focus on the case of $\epsilon > 1$ in the Numerical Solutions Sect. 5 below.

Summarizing, when considering the effect of ambiguity aversion on the social cost of carbon we can identify two effects. First, including ambiguity aversion leads to a higher arrival rate and a larger certainty equivalent of expected damages conditional on arrival, which unambiguously pushes the social cost of carbon up. We call this effect the *direct* effect of ambiguity aversion. Second, there is a more *indirect* general equilibrium effect through the impact of ambiguity aversion on discount rates. The discount rate that should be used to discount future climate disasters is the consumption discount rate. When $\epsilon \leq 1$, the discount rate goes down as ambiguity aversion increases but when $\epsilon > 1$ and the elasticity of substitution is larger than 1, ambiguity aversion leads to a higher consumption discount rate. This is an intuitive result: if the representative agent is very ambiguity averse about climate disasters, he would rather like to consume today than to postpone consumption since the future consumption level is uncertain. Ambiguity aversion therefore

[6] Mathematically: $\Lambda(t, t+u) = \int_t^{t+u} \lambda_s ds$, where $\Lambda(t, t+u)$ is the expected number of disasters between t and $t+u$.

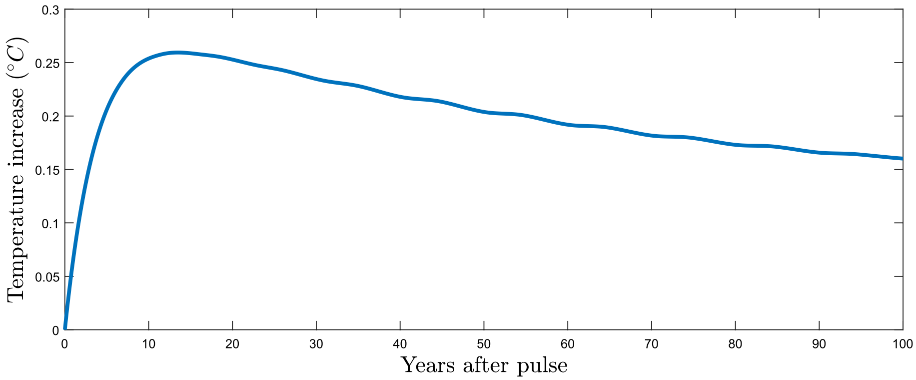

Fig. 3 Dynamic temperature response to an instantaneous 100GtC emissions pulse

increases the consumption discount rate when $\epsilon > 1$. For $\epsilon > 1$ it is ultimately a numerical issue which of the two effects dominates. We will highlight both effects separately in the numerical section and show that for our calibration the first effect dominates. Thus in our numerical analysis more ambiguity aversion leads to a higher SCC for all values of ϵ.

5 Climate Change and the Social Cost of Carbon: Numerical Results

We now move to a quantitative assessment of the theoretical results derived so far. We do this in three parts after first discussing our calibration choices (Sect. 5.1). In the first part we numerically evaluate our analytical formula and look at the magnitude of the direct and discounting effects (Sect. 5.2). In Sect. 5.3 we extend the model with the more realistic endogenous stochastic emissions. Last, in Sect. 5.4 we consider the case where also the climate senstivity is ambiguous.

There are two reasons for numerically solving the model with exogenous emissions before consider the model version with stochastic emissions. First, analyzing both variants allows us to show what adding endogeneity and thus stochastics to the emission process adds to the results on the SCC. But there is a second reason for also analyzing the exogenous emissions case: the analytical solution and its quantitative version allow us one more insight. Ambiguity aversion has an impact both on the discount rate and on the certainty equivalents being discounted but the two effects work against each other. Without an analytical solution we do not have a separate expression for the discount rate. The full model version therefore does not allow for separate assessment of the impact of ambiguity aversion on the discount rate and on the certainty equivalent damages being discounted. Numerically we can only assess the net impact.

5.1 Calibration

"Appendix F" gives the full details of the calibration of the climate model. Parameters for the growth rate of emissions and the initial level are chosen to match the baseline scenario of the DICE-2016 calibration (Nordhaus, 2017). The parameters of the carbon cycle and temperature model are taken from Mattauch et al. (2018). In addition, and different from Mattauch et al. (2018), we also include a base level of non-carbon related radiative forcing

Table 1 Parameters for the economic model

Par	Description	Value
C_t	Initial consumption level (PPP, in trillion 2015$)	83.07
λ_T	Arrival rate parameter	0.02 / 0.04
η	Disaster size parameter	30.25 / 61.5
$E[J]$	Expected disaster size	− 0.032 / − 0.016
γ	Risk aversion	5
θ	Ambiguity aversion parameter	0.1
a^*	Optimal ambiguity parameter	1.27 / 1.30
b^*	Optimal ambiguity parameter	0.74 / 0.75
ϵ	Elasticity of substitution	1.5
CDR_0	Consumption discount rate	1.5%

and calibrate it to match exogenous forcing in DICE-2016. This calibration leads to a temperature increase of 3.87 °C in 2100.

What in the end matters for the social cost of carbon is the derivative of temperature in the future with respect to additional emissions today. To consider the performance of the climate model, we therefore consider an instantaneous carbon pulse of 100GtC and look at the temperature increase. The result is given in Fig. 3. The temperature response peaks after approximately 12 years. After 100 years, the increase is around 0.16 degrees Celsius. By comparing this to Figure 1 in Dietz et al. (2021), it is clear that this temperature response is much better in line with climate science models than most climate modules in economic models.

As a robustness check, we also calculate the SCC using a simpler climate model in which temperature increases are assumed to be a linear function of cumulative emissions. This simple model turns out to be quite a good approximation of more complex climate models and is for example used in Dietz and Venmans (2019). The linear model would imply a flat line in Fig. 3. We take a Transient Climate Response to cumulative emissions of 1.75, which is the best estimate value from IPCC (2021). This implies that emitting 1000 GtC (3667 $GtCO_2$) would lead to to a global temperature increase of 1.75 °C. The results are given in “Appendix G”. The social cost of carbon is a bit higher using this alternative climate model. But the relative increase after adding risk aversion or ambiguity aversion does hardly change.

The calibration of the economic parameters is given in Table 1. Since we consider an exogenous endowment economy, output and consumption are the same thing in our model. That leaves the question open whether we should calibrate the endowment to output or to consumption data. The focus of the paper is on the social cost of carbon. What ultimately matters for the social cost of carbon is consumption, since utility depends on consumption and not on output. To make our results more comparable to other models, we therefore calibrate endowment to consumption data. The next choice to be made is whether one should aggregate output or consumption data using market exchange rates or using purchasing power parities (PPP). In line with the DICE-2016 model we use purchasing power parity exchange rates. Consumption data is not directly available in PPP. To obtain a proxy for world consumption in PPP we first obtain output data in PPP. Then we determine the world consumption ratio using market exchange rates. Our proxy for world consumption in PPP is then output in PPP multiplied by the world consumption ratio. Real world GDP (PPP) in

2015 equals 114.137 trillion 2015$ (IMF World Economic Outlook October 2016). World consumption in 2015 using market exchange rates equals 55.167 (in trillion 2010 $), while world GDP using market exchange rates equals 75.803 (in trillion 2010 $) (Worldbank Database). This yields a consumption-output ratio of 72.78%. Applying this ratio to World GDP (PPP) then gives 83.065 (in trillion 2015 $) for aggregate consumption in PPP terms.

The next step is to calibrate the climate disaster distribution, and in particular the parameters λ_T and η.[7] Our setup does allow for an arrival rate that is convex in temperature, but we do not consider this extension for two reasons. First, we want to keep things simple because there are enough other non-linearities in the SCC that we study. Second, it would also give another free parameter to calibrate.

Karydas and Xepapadeas (2019) also study a model with climate disasters and assume, based on natural disaster data, that for every degree warming the arrival rate increases by 6%. The disaster size is calibrated to 1.6%. This implies that the expected growth loss due to climate change would be $6\% \times 1.6\% = 0.096\%$ per degree global warming. We decide to choose $\lambda_T = 4\%$, which is smaller than Karydas and Xepapadeas (2019). We choose a smaller value for the arrival rate coefficient than Karydas and Xepapadeas (2019), because that places our calibration more in the middle of other estimates in the literature that we will discuss below. And we pick $\eta = 61.5$ which yields $E_t[J_t] = -1.6\%$, in line with Karydas and Xepapadeas (2019). Additionally, we consider a variant with less frequent but on average larger disasters: $\lambda_T = 2\%$, and a disaster size parameter $\eta = 30.25$ which gives $E_t[J_t] = -3.2\%$. While both calibrations have on average the same impact, their impact on risk premia is very different.

Let us compare these numbers to other studies in the literature. Several studies assume that climate damages have a level impact on the economy. In our setup, climate disasters have in expectation a growth impact of $4\% \times 1.6\% = 0.064\%$ per degree warming. Hambel et al. (2021) consider both level and a growth impacts of climate damages. They calculate three different growth impact variants. Their first specification is chosen to be similar to the damage function of the DICE model (Nordhaus 2017) and gives a loss of 0.026% per degree warming. Their second specification $0.0075T_t^{3.25}$ is non-linear and should match the damage function proposed by Weitzman (2012). At two degrees warming the growth loss is 0.07%, close to our specification. At higher temperature the losses quickly increase because of the high convexity. The last specification of Hambel et al. (2021) is again linear and is based on estimates from Dell et al. (2012), but the loss of 0.14% per degree is around 5 times higher than the loss of the DICE-case. Lastly, Bansal et al. (2019) develop a stylized model with climate disasters. Their arrival rate increases linearly with 1% per degree Celcius, and their disaster size is assume to be 5%. This gives an expected growth loss of 0.05% per degree. Our damage calibration fits well within the calibrations of these other models.

We now turn to the calibration of risk aversion and ambiguity aversion. We set risk aversion equal to 5. This level of risk aversion can be seen as conservative if we compare it to common values in the asset pricing literature.[8]

[7] Another option would be to assume that the damage size is a function of temperature instead of the arrival rate. However, making the damage size state-dependent would complicate the analysis in two ways. First, the damage size is a very non-linear function of the parameter η. There is not a simple equivalent to the linear specification that we currently have for the arrival rate. Second, this would also make relative entropy state-dependent. We therefore leave such an extension for further research.

[8] A coefficient of relative risk aversion between 5 and 10 is common in the asset pricing literature according to Cochrane (2009).

Fig. 4 Detection error probabilities as a function of θ

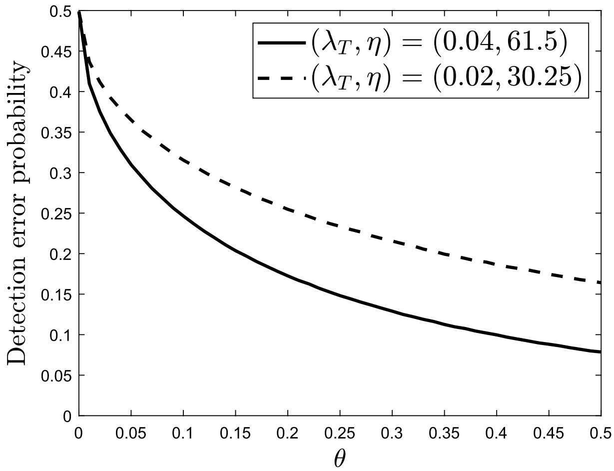

The level of ambiguity aversion is harder to calibrate. To get a feeling for reasonable values of ambiguity aversion, we use detection error probabilities. This method can be seen as a thought experiment: First, choose a value for the ambiguity aversion parameter. The ambiguity aversion parameter θ pins down the arrival rate and the expected jump size in the worst-case scenario. A higher θ leads to a higher worst-case arrival rate and a more negative worst-case expected jump size. Then assume that a decision maker observes the realisations of the stochastic process for N years, but he does not know whether the process is generated by the true or worst case model. He has to choose which of the two models is most likely. We can repeat this experiment many times. The detection error probability is then the probability of choosing the wrong model (so choosing the reference model $\mathbb{P}$ when the data was generated by the worst-case model $\mathbb{Q}^{a,b}$ and vice-versa).

For a given level of θ, the detection error probability gives information about how similar the worst-case and the reference model are. It helps in determining whether a value of θ is high or low. In the case without ambiguity aversion, both models are the same and the processes become indistinguishable. Choosing the most likely model becomes a random guess and the detection error probability equals 50%. A detection error probability that is close to 50% therefore indicates that the level of ambiguity aversion is small, because the worst case model is still indistinguishable form the reference model. When θ is very high, the two models are very different and the probability of making a mistake is close to zero. This indicates that the worst-case model is extreme and the ambiguity aversion parameter is very high.

We calculate the detection error probability assuming that the consumption process can be observed over a period of 100 years. A longer time horizon would have given lower detection error probabilities. The ambiguity aversion parameter θ is varied between 0 and 0.5. The results are given in Fig. 4. Detection error probabilities are decreasing in θ and are higher for a lower λ_T. This is intuitive, since a lower λ_T implies that there are less disasters over the observed time period and the probability of choosing the wrong model is therefore larger. We choose to set $\theta = 0.1$ in the base calibration, which gives a detection error probability of 24.6% for $(\lambda_T, \eta) = (0.02, 30.25)$ and 31.6% for $(\lambda_T, \eta) = (0.04, 61.5)$ (cf Fig. 4). This level of ambiguity aversion balances the trade-off between wanting to make a robust decision, but not taking into account too extreme models. The detection error probabilities for $\theta = 0.1$ are sufficiently far away from 50%, which implies that the reference model and the worst case model are not too similar. On the other hand, the detection error probabilities are also not close to 0, which would indicate an extreme amount of ambiguity aversion. However, since this parameter remains hard to calibrate, we do vary θ in robustness checks.

For the calibration $(\lambda_T, \eta) = (0.04, 61.5)$ the resulting optimal parameters with $\theta = 0.1$ are: $a^* = 1.30$ and $b^* = 0.75$. The arrival rate under the worst-case probability measure is 30% higher compared to the reference model. And the expected jump size becomes $\frac{-1}{b^*\eta+1} = -2.12\%$ compared to -1.6% in the reference model. The optimal parameters for the case $(\lambda_T, \eta) = (0.02, 30.25)$ are quite similar: $a^* = 1.27$ and $b^* = 0.74$.

The parameters that still have to be calibrated affect the social cost of carbon only indirectly, via the discount rate. Equation (15) shows that one can separate the expression for the Consumption Discount Rate (the relevant discount rate for the social cost of carbon) CDR_t in a time-independent part CDR_0 and a part that does depend on time as:

$$\begin{aligned} CDR_t &= CDR_0 + (1/\epsilon - 1)a^*\lambda_t \frac{-1}{b^*\eta + 1 - \gamma} \\ CDR_0 &= \beta + (1/\epsilon - 1)\Big(\mu - \frac{\gamma}{2}\sigma^2\Big). \end{aligned} \tag{19}$$

CDR_0 is the consumption discount rate in the absence of climate disasters. First, the value of the elasticity of intertemporal substitution ϵ determines whether additional risk increases or decreases the discount rate. There is no consensus in the profession on whether ϵ is greater or smaller than one.

In many applications the EIS is assumed to be smaller than one for an indirect reason: econometric evidence uniformly suggests that the rate of risk aversion γ is substantially larger than one and with the commonly used assumption of power utility that results in $\epsilon < 1$.[9] In the asset pricing literature where Epstein–Zin preferences are used, it is often estimated that the $EIS > 1$ (Van Binsbergen et al. 2012; Vissing-Jørgensen and Attanasio 2003). This is because a high EIS is necessary to obtain low enough equilibrium risk-free interest rates in an asset pricing model that are similar to observed risk-free interest rates. In integrated assessment models, $1/EIS$ is sometimes interpreted as a measure of intergenerational inequality aversion. The idea is that when EIS is small, the representative agents has a large preference for consumption smoothing. This can then be interpreted as a preference for equality between generations. Using this interpretation, an $EIS < 1$ is often assumed. Rezai and Van der Ploeg (2016) for example assume that intergenerational inequality aversion $IIA = 2$, which gives and EIS of 0.5. And Venmans and Groom (2021) estimate $IIA > 1$ using hypothetical decision tasks related to environmental inequalities across space and time.

We follow the asset pricing approach, because we focus on discounting and we want our parameter estimates to be in line with empirical evidence in this literature. We therefore choose $\epsilon = 1.5$, which is a common value in the literature on Epstein–Zin preferences (see for example Epstein and Zin (1989)). Additionally, we also consider a special case with $\epsilon = 1$. The growth rate μ, the volatility σ and the pure rate of time preference β only affect the social cost of carbon via CDR_0. The calibration of β has been widely discussed in the climate change literature. Additionally, we could calibrate σ from observed consumption volatility. However, as Mehra and Prescott (1985) point out, the model in that case

9 Because under the assumption of power utility $\epsilon = \frac{1}{\gamma}$.

Table 2 Social cost of carbon as function of risk aversion and ambiguity aversion

Social cost of carbon	$(\lambda_T, \eta) = (0.04, 61.5)$	$(\lambda_T, \eta) = (0.02, 30.25)$
$\epsilon = 1, \gamma = 0, \theta = 0$	175	175
$\epsilon = 1, \gamma = 5, \theta = 0$	191	209
$\epsilon = 1, \gamma = 5, \theta = 0.1$	338	381
$\epsilon = 1.5, \gamma = 0, \theta = 0$	363	363
$\epsilon = 1.5, \gamma = 5, \theta = 0$	360	392
$\epsilon = 1.5, \gamma = 5, \theta = 0.1$	599	664

would generate a way too low risk premium (the equity premium puzzle). A way to circumvent this is to calibrate σ to the volatility of stock prices, but this solution is also not very satisfactory. There have been several (partial) solutions proposed to the equity premium puzzle, for example including economic disaster risk. Solving the equity premium puzzle goes beyond the scope of this paper. Since both β and σ only affect the SCC via CDR_0, we choose to directly calibrate the consumption discount rate in the absence of climate risk. In our base calibration, we choose $CDR_0 = 1.5\%$, but we show our results for values of CDR_0 between 0.5% and 2.5%. The parameter combinations $(\beta, \mu, \sigma) = (2.25\%, 2.5\%, 3\%)$ and $(\beta, \mu, \sigma) = (1.5\%, 2.5\%, 10\%)$ for example yield a consumption discount rate $CDR_0 = 1.5\%$. Note that the actual consumption discount rate CDR_t is higher because of the impact of climate disasters on discounting.

5.2 The Social Cost of Carbon: The Analytical Model Quantified

Our base calibration yields a social cost of carbon of \$599 per ton of carbon (\$163 per ton CO_2) with $(\lambda_T, \eta) = (0.04, 61.5)$ and \$664 per ton carbon (\$181 per ton CO_2) with $(\lambda_T, \eta) = (0.02, 30.25)$.[10] Comparing the two cases shows that it matters whether the disasters are frequent but small (large η) or more infrequent but larger (smaller η). The two sets of assumptions yield the same expected disaster shock, but in the low frequency/large-shock case risk aversion and ambiguity aversion play a larger role and the social cost of carbon is correspondingly higher. In the following sections, we discuss different variants to decompose discounting, risk and ambiguity effects. The numerical outcomes are summarized in Table 2.

Risk aversion and ambiguity aversion with a unitary EIS

In the literature authors make widely varying assumptions on whether the EIS is smaller, larger or equal to 1. We start with a simple case where we assume that the EIS $\epsilon = 1$. This implies that the consumption discount rate simply equals the pure rate of time preference $\beta = 2.25\%$. Risk and ambiguity have therefore no impact on discounting and will only affect the direct valuation of damages. By definition, the SCC is the same for both calibrations when risk aversion γ and ambiguity aversion θ are both 0. In that case the expected value of both calibrations is the same and since risk is not priced under those assumptions, the SCC is the same for both calibrations. The case without risk aversion and ambiguity aversion gives an SCC of 175\$ per ton carbon. Adding risk aversion increases

[10] We express the social cost of carbon in the rest of this paper in dollars per ton carbon. To convert in dollars per ton CO_2, divide by 3.67.

the SCC with 9% for the small but frequent disaster case and with 19% for the larger disaster case. Adding ambiguity aversion leads to a much higher increase of 77% and 82% in respectively the small and large disaster case. The intuition behind the difference is that risk aversion and ambiguity aversion have a larger effect with less frequent but larger disasters. Therefore the relative increase in the SCC due to ambiguity aversion is larger with the $\lambda_T = 0.02$ setup than it is with $\lambda_T = 0.04$ setup. This example also suggests that ambiguity aversion has a substantially larger effect on the SCC than risk aversion.

Risk aversion and ambiguity aversion when EIS > 1

To clarify the impact of the EIS we now increase the elasticity of intertemporal substitution to 1.5. A higher willingness to substitute consumption over time yields a lower consumption discount rate. This lower discount rate implies that even in the case without risk aversion and ambiguity aversion, the SCC is twice as high compared to the situation with $EIS = 1$. This directly shows the importance of the discount rate in the SCC. Given the long horizon, small changes in the discount rate have large effects on the SCC. At the end of Sect. 5.2 we will show more results when varying discount rates.

We are mostly interested in how risk and ambiguity change the SCC. Introducing risk aversion with $EIS = 1.5$ has a negligible effect on the SCC for the frequent disasters with low disaster size. Higher risk aversion again lowers the certainty equivalent but it now also increases the discount rate. The net effect is even a slight decrease of the SCC from \$363 to \$360. This changes when damages are more infrequent but larger. In the alternative calibration with $(\lambda_T, \eta) = (0.02, 30.25)$, risk aversion does increase the social cost of carbon with 8%, from \$363 to \$392. Either way the impact of risk aversion on the SCC is much smaller when compared to the previous case with $EIS = 1$. This is because risk aversion now also affects the discount rate.

Introducing ambiguity aversion again leads to a significantly higher value of the social cost of carbon. When θ goes from 0 to 0.1, the SCC increases by 66% and 69% in respectively the small and large disaster case. Again the increase is smaller compared to the situation where $EIS = 1$ because of the discounting effect. But in the case of ambiguity aversion the direct effect clearly dominates the indirect discounting effect.

We should point out that our claim of a relatively small impact of risk aversion and a much larger impact of Ambiguity Aversion on the SCC is influenced by our use of the MaxiMin approach to assess the impact of Ambiguity Aversion. An important difference between the impact of ambiguity aversion and risk aversion is that under risk aversion the physical probability density function and the associated expected values do not change, although of course the risk neutral distribution function that matters for valuation of damages does change. But with higher ambiguity aversion, the physical probability density function and the associated expected losses change also. A small impact of risk aversion on the SCC is in line with much of the literature, but under our use of the MaxiMin approach to ambiguity aversion the likelihood and severity of bad outcomes go up while introducing risk aversion leaves the (physical) distribution and associated expected values unchanged.[11]

Risk aversion and the discounting effect

In this section we dive deeper into the direct and discounting effects when varying the level of risk aversion. We can disentangle the direct effect of risk aversion on the certainty equivalent and the indirect effect of risk aversion on the discount rate in the social cost of

[11] Note that in the smooth preferences approach associated with Klibanoff et al. (2005), using MaxiMin corresponds to assuming an infinite coefficient of ambiguity aversion. Lowering that coefficient will also lower the impact of ambiguity aversion on the SCC.

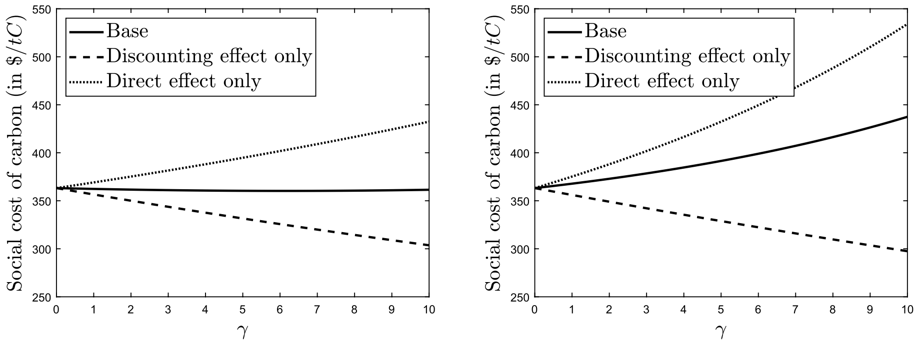

Fig. 5 Social cost of carbon as a function of risk aversion γ. This figure shows the social cost of carbon as a function of the risk aversion parameter γ. The total effect of risk aversion on the SCC is given by the solid line *(base)*. We additionally distinguish two special cases. In the *discounting effect only* case (dashed line) we assume that increasing γ does lead to an increase in the discount rate, but does not change the certainty equivalent in the SCC formula. In the *direct effect only* case (dotted line) we look at the opposite case, where increasing γ is assumed to have an effect on the certainty equivalent, but not on the consumption discount rate CDR_t

carbon. Figure 5 shows that even though the total effect of risk aversion on the social cost of carbon is small, the separate effects are not. Increasing risk aversion from 0 to 5 leads to an increase in the SCC from \$363 to \$392 in the case of infrequent but larger disasters. But only considering the direct effect would lead to an SCC of \$432, while only looking at the discounting effect lowers the SCC to \$329. Both effects largely cancel out and the resulting total effect is therefore small.

The size of the discounting effect is mainly dependent on two parameters. First, it depends on the elasticity of intertemporal substitution ϵ. The discounting effect is only negative if $\epsilon > 1$. When $\epsilon < 1$, additional risk, or more risk aversion would lower discount rates and both the indirect discounting effect and the direct effect of ambiguity aversion would have the same sign. From an asset pricing perspective, this leads to counter-intuitive effects. It would imply that the valuation of future consumption would increase when the volatility of consumption increases. For $\epsilon = 1$, the consumption discount rate CDR_t simply equals β and risk, risk aversion and ambiguity aversion do not affect discount rates.

Second, economic volatility plays an important role. The economic volatility does not change the distribution of climate disasters and therefore doesn't change the direct effect. But if economic volatility is large, then increasing risk aversion does have a larger effect on the discount rate. And the discounting effect will therefore be larger.

These results show that the discount effect is not negligible and this effect can actually dominate the direct effect of risk aversion. This can even give the counter-intuitive result that increasing risk aversion leads to a lower social cost of carbon.

Ambiguity aversion and the discounting effect

Figure 6 shows for each of the two sets of assumptions on the disaster risk parameters the social cost of carbon for different values of θ. We saw already from Eq. (19) that ambiguity aversion affects both the arrival rate and the certainty equivalent of climate disasters, but also the discount rate. In our calibration with $\epsilon = 1.5 > 1$, more ambiguity aversion leads to a higher discount rate which means the direct effect via the arrival rate and the certainly equivalent and the indirect effect via the discount rate have the opposite effect on the

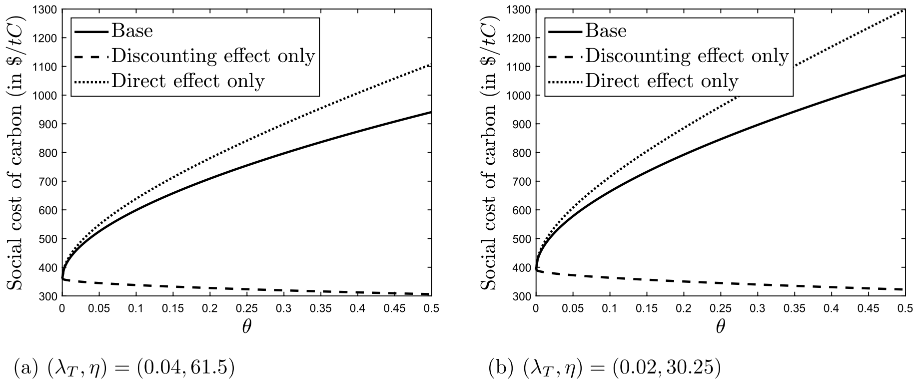

Fig. 6 Social cost of carbon as a function of θ. This figure shows the social cost of carbon as a function of the ambiguity aversion parameter θ. The total effect of ambiguity aversion on the SCC is given by the solid line *(base*). We additionally distinguish two special cases. In the *discounting effect only* case (dashed line) we assume that increasing θ does lead to an increase in the discount rate, but does not change the arrival rate and the certainty equivalent in the SCC formula. In the *direct effect only* case (dotted line) we look at the opposite case, where increasing θ is assumed to have an effect on the arrival rate and the certainty equivalent, but not on the consumption discount rate CDR_t

Fig. 7 Social cost of carbon as a function of CDR_0

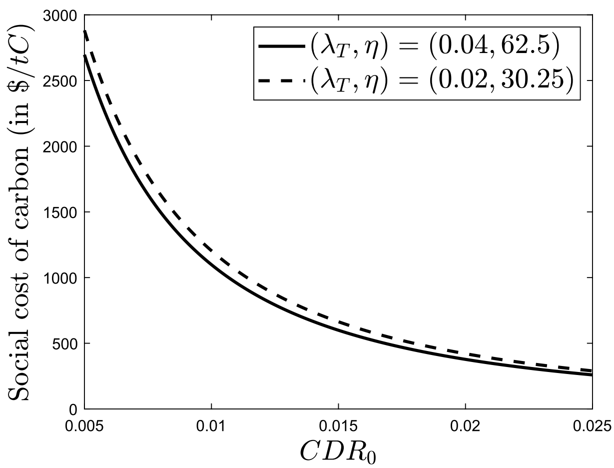

SCC. We show the two effects separately and combined in Fig. 6. There we consider the indirect *discounting only* effect, in which we assume ambiguity aversion only affects the discount rate CDR_t (the dashed line); and the direct effect where we leave the consumption discount rate CDR_t unchanged, but take into account the direct effect of ambiguity aversion on the arrival rate and certainty equivalent of the climate disasters (dotted line) in Fig. 6. The two effects are combined in the case labeled "Base" (solid line). Figure 6 clearly indicates that ambiguity aversion increases the discount rate (remember we assume $\epsilon > 1$ in this set of simulations); but we also see that the direct effect on the SCC dominates, the solid line slopes upward. We conclude that even for $\epsilon > 1$ ambiguity aversion leads to a higher social cost of carbon, and in our calibration actually substantially so.

In contrast to risk aversion, the discounting effect seems to be smaller when increasing ambiguity aversion. This happens because we have assumed that ambiguity aversion only

Table 3 Social cost of carbon as function of risk aversion and ambiguity aversion with stochastic emissions correlated to output

Social cost of carbon	$(\lambda_T, \eta) = (0.04, 61.5)$	$(\lambda_T, \eta) = (0.02, 30.25)$
$\epsilon = 1.5, \gamma = 0, \theta = 0$	352	352
$\epsilon = 1.5, \gamma = 5, \theta = 0$	368	399
$\epsilon = 1.5, \gamma = 5, \theta = 0.1$	609	673

affects climate disasters. It therefore only has a limited effect on the discount rate. Risk aversion on the other hand also changes the way economic risk is valued and it therefore has a larger effect on the discount rate.

Discount rates and the SCC

Figure 7 shows the dependence of the SCC on the time-independent part of the consumption discount rate CDR_0, the core discount rate. Note that the actual discount rate that is used to discount future damages (CDR_t) is higher than CDR_0 due to the effect of climate disasters itself on discounting. When core discount rates are close to zero, the social cost of carbon becomes very high. With $CDR_0 = 0.5\%$, the SCC is even above \$2000, around four times higher than in the base calibration. On the other hand, setting $CDR_0 = 2.5\%$ gives a social cost of carbon that is less than half the value in the base calibration. This figure highlights the importance of the discount rate when analyzing climate change and in particular its impact on the social cost of carbon.

5.3 The Social Cost of Carbon with Stochastic Emissions

So far we have made the obviously counterfactual assumption that emissions are a non-stochastic process, since assuming otherwise would preclude analytical solutions. In this section we remedy this shortcoming by modeling emissions as an explicitly stochastic process correlated to the process generating output. The short answer to the question what this brings about is that the main results are still true in this more realistic case. But stochastic emission processes add to risk and uncertainty, with as implication that the social cost of carbon under risk aversion and ambiguity aversion is slightly higher.

Thus assume now that emissions are the product of carbon intensity ψ_t and aggregate endowment C_t: $E_t = \psi_t C_t$. We calibrate the stochastic process for ψ_t such that in expectation emissions are similar to what they are in the non-stochastic emissions case. To bring this about we postulate that ψ_t declines at the rate $\psi_0 e^{-\alpha_\psi t} + \delta_\infty(1 - e^{-\alpha_\psi t})$ and set $\delta_0^\psi = -0.6\%$, $\delta_\infty^\psi = -6\%$ and $\alpha_\psi = 0.0045$. All other parameters are the same as in the exogenous case. The only difference is that future emissions are now also stochastic and correlated to output. The solution method is described in "Appendix H".

The results are given in Table 3. For zero risk aversion and in the absence of ambiguity aversion ($\gamma = 0$ and $\theta = 0$) the SCC is slightly smaller compared to the exogenous emissions case, although negligibly so: 352 $\$/tC$ instead of 363 $\$/tC$. But with endogenous and stochastic emissions, both risk and ambiguity aversion have larger effects on the social cost of carbon. For the high expected damages parametrization of the climate damages jump process $(\lambda_T, \eta) = (0.02, 30.25)$, column two in Table 3 shows that increasing γ from 0 to 5 leads to a 13 % increase in the SCC (go from the first to the second row in column two of Table 3). Adding ambiguity aversion (go from the second to the third row in column two of Table 3) leads to a further 69 % increase in the SCC. The combined impact of going from no risk/ambiguity aversion to our base case assumptions on the risk and ambiguity aversion parameters is a 91% increase in the SCC, up from 83% in the non-stochastic emissions case.

For the alternative low-expected-damages case $(\lambda_T, \eta) = (0.04, 61.5)$, column one in Table 3, the impact of increasing γ and θ is smaller although still sizable. Increasing γ from 0 to 5 leads to a 5 % increase in the SCC and subsequently increasing θ to 0.1 leads to a substantially larger increase in the SCC of 65%. The combined impact is in this case a somewhat smaller but still large: an increase in the SCC of 73%.

Overall Table 3 shows that our main results remain valid in the more realistic case where emissions are a function of aggregate endowment, while the impact of risk aversion and ambiguity aversion now is even a bit larger. These results should be intuitive since with endogenous emissions there is additional risk within the model: future emissions are now also stochastic. Especially risk aversion will therefore have a larger effect.

5.4 Ambiguity About the Climate Sensitivity Parameter α

So far we have assumed that the decision maker is only uncertain about the arrival rate and the size distribution of climate disasters and that all parameters in the entire climate system were known. In this section we go further and analyze the case where the decision maker is also uncertain about the climate sensitivity parameter α.

Like before we follow the multiple prior approach: the decision maker considers alternative probability distributions. However now every alternative probability measure is characterized by three instead of two parameters: a, b and c. Under the alternative measure $\mathbb{Q}^{a,b,c}$, the jump arrival rate and jump size are again given by $\lambda_t^{\mathbb{Q}^{a,b,c}} = a\lambda_T T_t^{\mathbb{Q}^{a,b,c}}$ and $\eta^{\mathbb{Q}^{a,b,c}} = b\eta$. However, future temperature levels are now different under alternative measures because the climate sensitivity is assumed to be unknown. Under the measure $\mathbb{Q}^{a,b,c}$, the temperature sensitivity becomes $\alpha^{\mathbb{Q}^{a,b,c}} = c\alpha$.

We next run into a new problem: with ambiguity about the climate sensitivity parameter and thus about the *future* temperature level T, instantaneous relative entropy is not a suitable distance measure anymore. We chose it before because of its tractability and because the ambiguity parameters directly affected the probability distribution. But ambiguity about the climate sensitivity parameter α does not alter the *current* probability distribution of climate damages, but only future pdf's after some time through different *future* temperature levels. Instantaneous entropy measures ignore future differences so we switch to lifetime entropy as distance measure instead of instantaneous entropy, like in Hansen and Sargent (2001).

Let $d_t(a, b, c)$ be the instantaneous distance measure that we used in previous sections. Lifetime entropy is then given by:

$$\tilde{d}(a, b, c) = \int_0^\infty e^{-\beta t} d_t(a, b, c) dt. \tag{20}$$

In the old situation with a distance measure that was not time-dependent, using instantaneous entropy was actually equivalent to using lifetime entropy. The only difference is that the ambiguity aversion parameter with lifetime entropy should be divided by β to make it comparable to the parameter with instantaneous entropy. This can be seen by calculating the integral of $d(a, b, c)$ when $d(a, b, c)$ does not depend on t. In this case $\tilde{d}(a, b, c) = \frac{d(a,b,c)}{\beta}$. We therefore divide the ambiguity parameter θ by the time preference parameter β in the lifetime constraint to make sure θ in the new situation is comparable to the way it was used in the previous sections.

Table 4 Social cost of carbon with ambiguity about α

Ambiguity aversion θ	a^*	b^*	c^*	SCC
$\theta = 0.1$	1.05	0.75	1.31	631
$\theta = 0.15$	1.07	0.71	1.35	697

We again use the minimax approach:

$$V_t = \min_{\mathbb{Q}^{a,b,c} \in \mathcal{P}^\theta} V_t^{\mathbb{Q}^{a,b,c}}$$
$$\text{where } \mathcal{P}^\theta = \{\mathbb{Q}^{a,b,c} : \tilde{d}(a,b,c) \leq \frac{\theta}{\beta}\}. \tag{21}$$

We assume that the decision makers chooses the optimal parameters (a, b, c) only once at the beginning of the problem. Previously a constant a^* and b^* were an optimal outcome of the model, but here we actually restrict the problem to constant parameters. This is necessary to be able to numerically solve the problem, because otherwise the parameters would have to be updated continuously. Not reoptimizing could create a time inconsistency problem which we choose to ignore since there is no clear a priori reason why the parameters should change over time given a constant "life time" budget θ.

We solve the problem as follows. We show in "Appendix C.2.2" that to find the optimal parameters, we need to minimize the following function:

$$\frac{\partial g^{\mathbb{Q}^{a,b,c}}(X_t)/\partial T_t}{g^{\mathbb{Q}^{a,b,c}}(X_t)(1-\gamma)} \frac{1}{\tau}\Big(F_t(c) - \upsilon T_t - \kappa(T_t - T_t^{oc})\Big) + a\lambda_t \frac{-1}{b\eta + 1 - \gamma}. \tag{22}$$

The parameters a and b show up in the last term which is directly related to climate disasters. The parameter c occurs in the optimization problem through radiative forcing F_t. The climate sensitivity parameter c only has an indirect effect on damages through temperature, which is captured in the first term of the equation. For a given combination of parameters (a, b, c), we can use numerical integration to find the unknown function g and after that we numerically differentiate g to calculate the derivative with respect to T. We also use numerical integration to evaluate the entropy constraint. We then use a constrained minimization solver to find the optimal combination of parameters.

5.4.1 Results

We first consider a run with the ambiguity aversion parameter $\theta = 0.1$. All other parameters are the same as in the default analysis.[12] From Table 4 we can see that this gives the optimal parameters (a^*, b^*, c^*) = (1.05, 0.75, 1.31). We can compare these numbers to the default case without an uncertain climate sensitivity. For the comparison with the default case we choose the same ambiguity budget: $\theta = 0.1$. The optimal parameters were (a^*, b^*) = (1.30, 0.75). So these new values indicate that the optimal b^* does not change, which implies that the worst-case disaster size stays the same. However, the introduction of ambiguity in the climate sensitivity does create a trade-off between direct ambiguity in

[12] We consider the case with higher frequency and lower damage size $(\lambda_T, \eta) = (0.04, 61.5)$.

the arrival rate through a^* and the indirect ambiguity through the temperature level. The optimal a^* is much lower than before, with as counterpart a higher value of c^*. The objective function turns out to be very flat in the a^*, c^* dimension because both parameters affect the arrival rate. This means that many (a^*, c^*) combinations will give a similar value of the objective function.

The social cost of carbon becomes 631 \$/tC, which is around 5% larger than in the corresponding default case without ambiguity on the climate sensitivity parameter α. Intuition behind this results is as follows. We calibrate $\beta = 2.25\%$, but the effective discount rate in the social cost of carbon becomes $CDR_0 = 1.5\%$ (cf Eq. 19 for the definition of the time independent part of the Consumption Discount Rate CDR_0). The effective discount rate is lower since we have assumed that the elasticity of intertemporal substitution $\epsilon > 1$. The higher discount rate β is used to discount future instantaneous entropy levels, which gives more room for deviations of the probability distributions in the future. Because the climate sensitivity alters the arrival rate only in the future, a large discount rate allows for a higher level of the climate sensitivity. The higher future damages are then discounted with a relatively lower rate in the social cost of carbon, which yields a higher social cost of carbon.

Additionally we also consider a case with a higher ambiguity budget. The ambiguity budget captures ambiguity about all three parameters at the same time. If we actually want to allow for additional ambiguity aversion because of the uncertain climate sensitivity, we should increase the ambiguity budget θ. Increasing θ with 50% to 0.15 leads to a social cost of carbon of 697\$/tC, which is around 10% higher than the social cost of carbon with $\theta = 0.1$.

Summarizing we can conclude that introducing fundamental uncertainty (ambiguity) on the climate sensitivity parameter α increases the social cost of carbon. The increase is small if we do not allow for a larger ambiguity budget. If we do allow for more uncertainty by increasing the ambiguity budget, then the social cost of carbon increases more.

Our result that b^* is unaffected is driven by the assumption that the arrival rate depends on temperature, but the size of disasters does not. If we would also allow for a temperature dependent disaster size, then there would be a trade-off between both (a^*, c^*) and (a^*, b^*). However, we do not expect the results to change qualitatively. Allowing for additional ambiguity would anyhow, as we have just seen, probably also lead to a relatively small increase in the SCC unless one additionally allows for a higher ambiguity budget θ.

6 Conclusions

Climate change will beyond reasonable doubt have a large impact on the economy in the future. However, because of the complex nature of the problem and the lack of data, it is not possible yet to accurately estimate the timing and extent of its impact. But we do know that potentially large and irreversible consequences are likely to take place unless mitigating policies are implemented in time. But these changes will happen possibly far into the future, while mitigating policies are (or should be) under consideration right now. That discrepancy should put the discussion on discounting at the center of the debate about the social cost of carbon and what we should do about climate change: to compare uncertain future damages with costs today, those future damages need to be discounted back towards today. The debate in the literature has largely zeroed in on the rate of time preference; the problem there is that to be consistent with capital market data, discount rates must be relatively high which in turn does not leave much once climate change consequences a century out are discounted back towards today (cf Weitzman (2007) for a very lucid overview of this debate). In this paper we also focus on the discounting question and its implication for the SCC, but we take a different approach. Rather than discussing numerical values of certain parameters, we explore alternative specifications of preferences with respect to risk and more fundamental uncertainty. We explicitly introduce not just risk (i.e. stochastic outcomes with known probability distribution) but also ambiguity (stochastic outcomes with unknown distribution), and show the implications for the social cost of carbon of risk and ambiguity aversion under different and independent assumptions about the intertemporal rate of substitution.

To do so we focus on the effect of Epstein–Zin recursive preferences on outcomes of the model, and on the impact of unmeasurable risk (ambiguity) and the interaction between those two. Both breaking the link between γ and the EIS (by introducing Epstein–Zin utility) and introducing ambiguity aversion are conceptually relevant in the climate change setting. The first extension is relevant because climate change problems have a very long horizon and therefore the elasticity of intertemporal substitution (EIS) unavoidably plays an important role. Arbitrarily restricting its value to $1/\gamma$ is then surely unsatisfactory. Second, conceptually ambiguity aversion is a logical extension, since we have no accurate estimation of climate damages nor in particular of their probability density function in the future. The assumption of unmeasurable risk ("Knightian uncertainty") then is a natural framework to use. Finally we highlight the sometimes complicated interactions between ambiguity aversion and intertemporal substitution elasticities for the value of the Social Cost of Carbon.

To do all this we set up an analytic IAM by extending a disaster risk model with a climate change model and a temperature dependent arrival rate. Furthermore, we model climate risk as disaster risk instead of assuming that temperature increases generate a certain amount of damage every year. The model is transparent because we manage to derive closed form solutions for the social cost of carbon. Where stochastic numerical IAMs can take hours to be solved, solving our model only requires numerical integration and is therefore solved within seconds.

Our base calibration generates a substantial social cost of carbon. Most importantly, we use our model to highlight how ambiguity aversion changes the social cost of carbon. The social cost of carbon including ambiguity aversion is between \$599 and \$664 per ton of carbon (\$163–\$181 per ton CO_2) with non-stochastic emissions, and slightly more for the stochastic case (between \$609 and \$673 per ton of carbon). Ambiguity aversion is responsible for an increase of the SCC with 65% up to 91% depending on the model specification.

Analysing the effect of ambiguity aversion on the SCC is not a trivial exercise since multiple potentially offsetting effects play a role: we show that ambiguity aversion has both a *direct* effect on the arrival rate and certainty equivalent of disasters for given discount rates (more ambiguity aversion leads to a higher certainty equivalent) and an *indirect* effect on the discounting component. The effect of ambiguity aversion on discounting depends on the intertemporal rate of substitution ϵ. When $\epsilon < 1$, increasing ambiguity aversion leads to a smaller effective discount rate on climate damages, making for a higher SCC since both the direct and the indirect effect work in the same direction. For the arguably interesting (because empirically supported) case $\epsilon > 1$, increasing ambiguity aversion has two offsetting effects on the SCC, the direct and indirect effects actually work in different direction. However, we show that even then the direct effect dominates when evaluated numerically and therefore that the presence of ambiguity aversion leads to a (substantially) higher social cost of carbon.

Lastly, we also show the importance of fully considering the impact of the consumption discount rate on the social cost of carbon, not just the impact of the rate of time preference. It is of course well known that the social cost of carbon is very sensitive to changes in the discount rate, but we stress that analyzing the discount rate impact of climate change involves more than a discussion of the pure rate of time preference on the discount rate; a low discount rate can also be caused by a high elasticity of intertemporal substitution, and additionally the appropriate discount rate depends in elaborate ways on the growth rate of the economy, volatility, risk aversion, climate disaster risk and ambiguity aversion. Disentangling these various effects, their interactions and their impact on the SCC is the key contribution of this paper. One major theme emerges: proper risk pricing and incorporating ambiguity aversion leads to much higher estimates of the Social Cost of Carbon.

Appendix A: Relative Entropy and the Distance Measure

For each a and b we define the measure $\mathbb{Q}^{a,b}$ which is equivalent to $\mathbb{P}$ and has Radon–Nikodym derivative $\xi_t^{a,b}$ where $\xi_t^{a,b}$ follows:

$$d\xi_t^{a,b} = (\lambda_t - \lambda_t^{\mathbb{Q}^{a,b}})\xi_t^{a,b}dt + \left(\frac{\lambda_t^{\mathbb{Q}^{a,b}} f^{\mathbb{Q}^{a,b}}(J_t)}{\lambda_t f(J_t)} - 1\right)\xi_{t-}^{a,b}dN_t. \tag{23}$$

Under the alternative measure $\mathbb{Q}^{a,b}$ the arrival rate equals $\lambda_t^{\mathbb{Q}^{a,b}} = a\lambda_t$ and the jump size parameter equals $\eta^{\mathbb{Q}^{a,b}} = b\eta$. We can calculate in this case the fraction of the two probability distributions: $\frac{f^{\mathbb{Q}^{a,b}}(x)}{f(x)} = b(1+x)^{(b-1)\eta}$. Substituting this into (23) gives:

$$d\xi_t^{a,b} = (1-a)\lambda_t\xi_t^{a,b}dt + \left(ab(1+J_t)^{(b-1)\eta} - 1\right)\xi_{t-}^{a,b}dN_t. \tag{24}$$

The Radon–Nikodym derivative $\xi_t^{a,b}$ is the ratio between the alternative measure $\mathbb{Q}^{a,b}$ and the reference measure $\mathbb{P}$. The relative entropy between $\mathbb{Q}^{a,b}$ and $\mathbb{P}$ over time unit Δ is defined as $E_t^{\mathbb{Q}^{a,b}}\left[\log(\frac{\xi_{t+\Delta}^{a,b}}{\xi_t^{a,b}})\right]$. Here $E_t^{\mathbb{Q}^{a,b}}$ denotes the expectation with respect to the alternative measure $\mathbb{Q}^{a,b}$. Then divide by Δ and let $\Delta \to 0$ to obtain the instantaneous relative entropy: $RE(a,b) = \lim_{\Delta\to 0}\frac{1}{\Delta}E_t^{\mathbb{Q}^{a,b}}\left[\log\left(\frac{\xi_{t+\Delta}^{a,b}}{\xi_t^{a,b}}\right)\right]$.

Applying Itô's lemma for jump processes to $\xi_t^{a,b}$, we obtain the following dynamics for $\log(\xi_t^{a,b})$:

$$d\log(\xi_t^{a,b}) = (1-a)\lambda_t dt + \left(\log(ab) + (b-1)\eta\log(1+J_t)\right)dN_t. \tag{25}$$

Using integration by parts we can calculate that $E_t^{\mathbb{Q}^{a,b}}[\log(1+J_t)] = -\frac{1}{\eta^{\mathbb{Q}}}$. Therefore the (instantaneous) relative entropy at time t equals:

$$\begin{aligned} RE(a,b) &= \lim_{\Delta\to 0}\frac{1}{\Delta}E_t^{\mathbb{Q}^{a,b}}\left[\log\left(\frac{\xi_{t+\Delta}^{a,b}}{\xi_t^{a,b}}\right)\right] \\ &= (1-a)\lambda_t + a\lambda_t\left(\log(ab) + \frac{1}{b} - 1\right). \end{aligned} \tag{26}$$

Scaling relative entropy by the arrival rate λ_t yields our distance measure:

$$d(a,b) = \frac{RE(a,b)}{\lambda_t} = (1-a) + a\Big(\log(ab) + \frac{1}{b} - 1\Big). \tag{27}$$

Appendix B: Calculating the Detection Error Probability

After observing the process of consumption over a period N years, what is the probability of choosing the wrong model? Let us start with the case that the reference model $\mathbb{P}$ is the true model and the agent considers the alternative model $\mathbb{Q}^{a,b}$. Note that the Radon–Nikodym derivative informs us about the likelihood ratio of both models. When this derivative is larger than one after N years, the worst-case model $\mathbb{Q}^{a,b}$ is the most likely and the agent will choose the wrong model. The probability of making this error is equal to (see for example Maenhout (2006)):

$$P\Big(\xi_N^{a,b} > 1|\mathbb{P}\Big) = P\Big(\log(\xi_N^{a,b}) > 0|\mathbb{P}\Big). \tag{28}$$

We calculate this probability by simulating the process of $\log(\xi_t^{a,b})$ forward. Simulation is done via a standard Euler method. Similarly, we can define the opposite mistake where the alternative model is actually true and the agent chooses the reference model. We now define the inverse Radon–Nikodym derivative: $\frac{d\mathbb{P}}{d\mathbb{Q}^{a,b}} = \tilde{\xi}_t^{\,a,b}$ where $\tilde{\xi}_t^{\,a,b}$ follows:

$$d\tilde{\xi}_t^{\,a,b} = (a-1)\lambda_t\tilde{\xi}_t^{\,a,b}dt + \Big(\frac{1}{ab}(1+J)^{(1-b)\eta} - 1\Big)\tilde{\xi}_{t-}^{\,a,b}dN_t. \tag{29}$$

Applying Itô's lemma gives:

$$d\log(\tilde{\xi}_t^{a,b}) = (a-1)\lambda_t dt + \Big(-\log(ab) + (1-b)\eta\log(1+J_t)\Big)dN_t. \tag{30}$$

The probability of choosing the wrong model when actually the alternative model $\mathbb{Q}^{a,b}$ is true equals:

$$P\Big(\tilde{\xi}_N^{\,a,b} > 1|\mathbb{Q}^{a,b}\Big) = P\Big(\log(\tilde{\xi}_N^{\,a,b}) > 0|\mathbb{Q}^{a,b}\Big). \tag{31}$$

Again this probability can be calculated by simulating the process $\log(\tilde{\xi}_t)$ forward. The detection error probability is then defined as:

$$\frac{1}{2}P\Big(\log(\xi_N^{a,b}) > 0|\mathbb{P}\Big) + \frac{1}{2}P\Big(\log(\tilde{\xi}_N^{\,a,b}) > 0|\mathbb{Q}\Big). \tag{32}$$

Appendix C: Solving the Model

Appendix C.1: Hamilton–Jacobi–Bellman Equation

We will first derive the Hamilton–Jacobi–Bellman equation for every measure $\mathbb{Q}^{a,b}$. In the next subsection of the appendix we introduce ambiguity.

The value function $V_t^{\mathbb{Q}^{a,b}} = V^{\mathbb{Q}^{a,b}}(C_t, X_t)$ is a function of aggregate consumption C_t and the vector of climate state variables X_t. Let $V_C^{\mathbb{Q}^{a,b}}$ denote the first derivative of the value function with respect to aggregate consumption, similar notation is used for the second derivative. For notational purposes, define the vector of climate state variables:

$$X_t = [g_{E,t}\ E_t\ M_{0,t}\ M_{1,t}\ M_{2,t}\ M_{3,t}\ F_{E,t}\ T_t\ T_t^{oc}]'. \tag{33}$$

The vector of state variables then follows: $dX_t = \mu_X(X_t)dt$. Denote by $V_X^{\mathbb{Q}^{a,b}}$ the row vector of partial derivatives of the value function $V_t^{\mathbb{Q}^{a,b}}$ with respect to the vector of state variables X_t: $V_X^{\mathbb{Q}^{a,b}} = \left[\frac{\partial V^{\mathbb{Q}^{a,b}}(C_t,X_t)}{\partial g_{E,t}} \dots \frac{\partial V^{\mathbb{Q}^{a,b}}(C_t,X_t)}{\partial T_t^{oc}}\right]$.

Duffie and Epstein (1992b) show that the HJB-equation for stochastic differential utility equals:

$$0 = f(C_t, V_t^{\mathbb{Q}^{a,b}}) + \mathcal{D}\mathcal{V}^{\mathbb{Q}^{a,b}}. \tag{34}$$

Here $\mathcal{D}\mathcal{V}^{\mathbb{Q}^{a,b}}$ is the drift of the value function. In order to calculate the drift of the value function, we will apply Itô's lemma. By Itô's lemma for jump processes we have:

$$\begin{aligned} dV_t^{\mathbb{Q}^{a,b}} &= V_C^{\mathbb{Q}^{a,b}}\left(\mu C_t dt + \sigma C_t dZ_t\right) + V_X^{\mathbb{Q}^{a,b}}\mu_X(X_t)dt + \frac{1}{2}V_{CC}^{\mathbb{Q}^{a,b}}\sigma^2 C_t^2 dt \\ &\quad + \left(V^{\mathbb{Q}^{a,b}}\left((1+J_t)C_{t-}, X_t\right) - V^{\mathbb{Q}^{a,b}}(C_{t-}, X_t)\right)dN_t. \end{aligned} \tag{35}$$

Then the drift under $\mathbb{Q}^{a,b}$ equals:

$$\begin{aligned} \mathcal{D}\mathcal{V}^{\mathbb{Q}^{a,b}} &= V_C^{\mathbb{Q}^{a,b}}\mu C_t + V_X^{\mathbb{Q}^{a,b}}\mu_X(X_t) + \frac{1}{2}V_{CC}^{\mathbb{Q}^{a,b}}\sigma^2 C_t^2 \\ &\quad + \lambda_t^{\mathbb{Q}^{a,b}} E^{\mathbb{Q}^{a,b}}\left[V^{\mathbb{Q}^{a,b}}\left((1+J_t)C_{t-}, X_t\right) - V^{\mathbb{Q}^{a,b}}(C_{t-}, X_t)\right]. \end{aligned} \tag{36}$$

This gives the following Hamilton–Jacobi–Bellman equation:

$$\begin{aligned} 0 &= f(C_t, V_t^{\mathbb{Q}^{a,b}}) + V_C^{\mathbb{Q}^{a,b}}\mu C_t + V_X^{\mathbb{Q}^{a,b}}\mu_X(X_t) + \frac{1}{2}V_{CC}^{\mathbb{Q}^{a,b}}\sigma^2 C_t^2 \\ &\quad + \lambda_t^{\mathbb{Q}^{a,b}} E^{\mathbb{Q}^{a,b}}\left[V^{\mathbb{Q}^{a,b}}\left((1+J_t)C_{t-}, X_t\right) - V^{\mathbb{Q}^{a,b}}(C_{t-}, X_t)\right]. \end{aligned} \tag{37}$$

We conjecture and verify that the value function is of the following form:

$$V^{\mathbb{Q}^{a,b}}(C_t, X_t) = g^{\mathbb{Q}^{a,b}}(X_t)\frac{C_t^{1-\gamma}}{1-\gamma}, \tag{38}$$

where $g^{\mathbb{Q}^{a,b}}(X_t)$ is some function of X_t. Substituting our conjecture $V^{\mathbb{Q}^{a,b}}(C_t, X_t) = \frac{g^{\mathbb{Q}^{a,b}}(X_t)C_t^{1-\gamma}}{1-\gamma}$ into $f(C_t, V_t)$ gives:

$$\begin{aligned} f(C_t, V^{\mathbb{Q}^{a,b}}(C_t, X_t)) &= \frac{\beta}{1-1/\epsilon}\frac{C_t^{1-1/\epsilon} - \left(g^{\mathbb{Q}^{a,b}}(X_t)C_t^{1-\gamma}\right)^{\frac{1}{\zeta}}}{\left(g^{\mathbb{Q}^{a,b}}(X_t)C_t^{1-\gamma}\right)^{\frac{1}{\zeta}-1}} \\ &= \frac{\beta}{1-1/\epsilon}\left(g^{\mathbb{Q}^{a,b}}(X_t)^{1-\frac{1}{\zeta}}C_t^{1-\gamma} - g^{\mathbb{Q}^{a,b}}(X_t)C_t^{1-\gamma}\right) \\ &= \beta\zeta\left(g^{\mathbb{Q}^{a,b}}(X_t)^{-\frac{1}{\zeta}} - 1\right)V^{\mathbb{Q}^{a,b}}(C_t, X_t). \end{aligned} \tag{39}$$

The partial derivatives of V are given by:

$$V_C^{\mathbb{Q}^{a,b}} = g^{\mathbb{Q}^{a,b}}(X_t)C_t^{-\gamma}, \quad V_{CC}^{\mathbb{Q}^{a,b}} = -\gamma g^{\mathbb{Q}^{a,b}}(X_t)C_t^{-\gamma-1},$$
$$V_X^{\mathbb{Q}^{a,b}} = \frac{g_X^{\mathbb{Q}^{a,b}}(X_t)C_t^{1-\gamma}}{1-\gamma}. \tag{40}$$

Here $g_X^{\mathbb{Q}^{a,b}}$ denotes the row vector with partial derivatives to each of the state variables, similar to $V_X^{\mathbb{Q}^{a,b}}$. Additionally we can calculate the expected impact of a jump on the value function:

$$E^{\mathbb{Q}^{a,b}}\left[V^{\mathbb{Q}^{a,b}}\left((1+J_t)C_{t-}, X_t\right) - V^{\mathbb{Q}^{a,b}}(C_{t-}, X_t)\right] = \frac{E^{\mathbb{Q}^{a,b}}\left[(1+J_t)^{1-\gamma}\right] - 1}{1-\gamma} g^{\mathbb{Q}^{a,b}}(X_t)C_t^{1-\gamma}$$
$$= \frac{\frac{b_t\eta}{b_t\eta+1-\gamma} - 1}{1-\gamma} g^{\mathbb{Q}^{a,b}}(X_t)C_t^{1-\gamma} = \frac{-1}{b_t\eta + 1 - \gamma} g^{\mathbb{Q}^{a,b}}(X_t)C_t^{1-\gamma}. \tag{41}$$

Substituting $f(C_t, V^{\mathbb{Q}^{a,b}}(C_t, X_t))$ together with the partial derivatives of $V_t^{\mathbb{Q}^{a,b}}$ and the expectation into (37) yields the following equation:

$$0 = \frac{\beta}{1-1/\epsilon}\left(g^{\mathbb{Q}^{a,b}}(X_t)^{-\frac{1}{\zeta}} - 1\right)g^{\mathbb{Q}^{a,b}}(X_t)C_t^{1-\gamma} + \mu g^{\mathbb{Q}^{a,b}}(X_t)C_t^{1-\gamma}$$
$$- \frac{\gamma}{2}\sigma^2 g^{\mathbb{Q}^{a,b}}(X_t)C_t^{1-\gamma} + \frac{g_X^{\mathbb{Q}^{a,b}}(X_t)C_t^{1-\gamma}}{1-\gamma}\mu_X(X_t) + a_t\lambda_t\frac{-1}{b_t\eta + 1 - \gamma} g^{\mathbb{Q}^{a,b}}(X_t)C_t^{1-\gamma}. \tag{42}$$

Dividing by $g^{\mathbb{Q}^{a,b}}(X_t)C_t^{1-\gamma}$ gives:

$$0 = \frac{\beta}{1-1/\epsilon}\left(g^{\mathbb{Q}^{a,b}}(X_t)^{-\frac{1}{\zeta}} - 1\right) + \mu - \frac{\gamma}{2}\sigma^2 + \frac{g_X^{\mathbb{Q}^{a,b}}(X_t)}{g^{\mathbb{Q}^{a,b}}(X_t)(1-\gamma)}\mu_X(X_t)$$
$$+ a_t\lambda_t\frac{-1}{b_t\eta + 1 - \gamma}. \tag{43}$$

Appendix C.2: Finding the Ambiguity Aversion Parameters

In this section we show how to derive the ambiguity aversion parameters that correspond to the worst case outcome while just staying within the ambiguity budget. We first consider the case where there is only ambiguity about the climate shock parameters: there is ambiguity about the the arrival rate and magnitude of the climate Poisson shock.

Appendix C.2.1: Optimal *a* and *b*

Given a probability measure $\mathbb{Q}^{a,b}$, we can solve Eq. (43) to find $g^{\mathbb{Q}^{a,b}}(X_t)$. Now let us return to the problem with ambiguity. We are not interested in the solution for every single measure $\mathbb{Q}^{a,b}$. The maxmin procedure advocated by Gilboa and Schmeidler (1989) that we apply in this paper requires us to focus on the worst case distribution, which leads to the following minimization problem:

$$V_t = \min_{\mathbb{Q}^{a,b} \in \mathcal{P}^\theta} V_t^{\mathbb{Q}^{a,b}}. \tag{44}$$

And since every probability measure $\mathbb{Q}^{a,b}$ that we consider in our set $\mathcal{P}^{\theta}$ is uniquely defined by the parameters a and b, minimizing over $\mathbb{Q}^{a,b}$ is equivalent to minimizing over the parameters a and b. So we can replace the global minimization problem of Eq. (44) by an instantaneous optimization problem over a and b. The HJB-equation of the problem with ambiguity then becomes:

$$\begin{aligned} 0 = \min_{(a,b)\ s.t.\ d(a,b)\leq\theta} \Bigg\{ & \frac{\beta}{1-1/\epsilon}\left(g^{\mathbb{Q}^{a,b}}(X_t)^{-\frac{1}{\zeta}} - 1\right) + \mu - \frac{\gamma}{2}\sigma^2 \\ & + \frac{g_X^{\mathbb{Q}^{a,b}}(X_t)}{g^{\mathbb{Q}^{a,b}}(X_t)(1-\gamma)}\mu_X(X_t) + a\lambda_t \frac{-1}{b\eta + 1 - \gamma} \Bigg\}. \end{aligned} \tag{45}$$

The optimal a and b can thus be found by solving:

$$\min_{(a,b)} a\lambda_t \frac{-1}{b\eta + 1 - \gamma} \quad s.t.\ \ d(a,b) \leq \theta. \tag{46}$$

We can drop λ_t and write this problem as a constrained optimization problem with Lagrangian:

$$L(a,b,l) = a\frac{-1}{b\eta + 1 - \gamma} - l\Big(d(a,b) - \theta\Big). \tag{47}$$

Here l is the Lagrange multiplier. a^* and b^* and the Lagrange-multiplier l are the solutions to the following first order conditions:

$$\begin{aligned} \frac{\partial}{\partial a}L(a,b,l) &= \frac{-1}{b\eta + 1 - \gamma} - l\Big(\log(ab) + \frac{1}{b} - 1\Big) = 0, \\ \frac{\partial}{\partial b}L(a,b,l) &= a\frac{\eta}{(b\eta + 1 - \gamma)^2} - la\frac{b-1}{b^2} = 0, \\ \frac{\partial}{\partial l_t}L(a,b,l) &= \theta - (1-a) - \Big(\log(ab) + \frac{1}{b} - 1\Big) = 0. \end{aligned} \tag{48}$$

From now on, we use the notation V_t for the optimal value function ($V_t = V_t^{\mathbb{Q}^{a^*,b^*}}$). Similar notation is used for g_t.

Appendix C.2.2: Optimization with Additional Ambiguity About the Climate Sensitivity Parameter α: Finding a^*, b^* and c^*

In Sect. 5.4 we analyze a setting with ambiguity about the climate sensitivity parameter. The HJB-equation for this problem is equal to:

$$\begin{aligned} 0 = \min_{(a,b,c)\ s.t.\ \tilde{d}(a,b,c)\leq\frac{\theta}{\beta}} \Bigg\{ & \frac{\beta}{1-1/\epsilon}\left(g^{\mathbb{Q}^{a,b,c}}(X_t)^{-\frac{1}{\zeta}} - 1\right) + \mu - \frac{\gamma}{2}\sigma^2 \\ & + \frac{g_X^{\mathbb{Q}^{a,b,c}}(X_t)}{g^{\mathbb{Q}^{a,b,c}}(X_t)(1-\gamma)}\mu_X(X_t) + a\lambda_t \frac{-1}{b\eta + 1 - \gamma} \Bigg\}. \end{aligned} \tag{49}$$

Similar to the two parameter case, we can drop all terms unrelated to a, b and c when we solve for the optimal parameters. $\mu_X(X_t)$ is a vector with drifts of all the state variables. Note that c is part of $\mu_X(X_t)$ because the climate sensitivity determines the change in temperature in relation to an increasing carbon concentration. The optimal parameters can therefore be found by solving the following problem:

$$\min_{(a,b,c)\ s.t.\ \tilde{d}(a,b,c)\leq\frac{\theta}{\beta}} \left\{ \frac{\partial g^{\mathbb{Q}^{a,b,c}}(X_t)/\partial T_t}{g^{\mathbb{Q}^{a,b,c}}(X_t)(1-\gamma)}\frac{1}{\tau}\left(F_t - \upsilon T_t - \kappa(T_t - T_t^{oc})\right) + a\lambda_t\frac{-1}{b\eta+1-\gamma} \right\}. \tag{50}$$

Radiative forcing F_t is a function of the climate sensitivity parameter α. It is not possible to obtain closed form first order conditions like in the two parameter case. We therefore numerically solve for the optimal parameters by minimizing this equation.

Appendix C.3: Solving the Model

It is typically not possible to solve the partial differential equation of the problem with climate state variables unless one would make the highly restrictive assumption assumption of a unit *EIS*, which we choose not to do. However we are able to obtain exact solutions for the value function and the consumption-to-wealth ratio without making restrictive assumptions like $EIS = 1$, and the consumption-to-wealth ratio is what we need for assessing the SCC. We will now sketch our approach.

Duffie and Epstein (1992a) derive that the pricing kernel (or stochastic discount factor) with stochastic differential utility equals $\pi_t = \exp\left\{\int_0^t f_V(C_s, V_s)ds\right\} f_C(C_t, V_t)$. However, the pricing kernel has to be adjusted for the ambiguity aversion preferences. Chen and Epstein (2002) show that the pricing kernel in the ambiguity setting should be multiplied by the Radon–Nikodym derivative $\xi_t^{a^*,b^*}$ of the measure corresponding to the optimal a^* and b^*. $\xi_t^{a,b}$ is defined in (23). When calculating the pricing kernel, we obtain an expression that depends on the unknown function $g(X_t)$. But by substituting the HJB-equation into the pricing kernel we obtain an expression that only depends on known parameters.

As an intermediate step it is helpful to introduce the concept of consumption strips. A consumption strip is an asset that pays a proportion of aggregate consumption C_s at the unique time $s > t$. Call its price at time t: $H(C_t, X_t, u)$, where u denotes the time to maturity; $u = s - t$. The price of a consumption strip paying out at time $s > t$ equals:

$$\begin{aligned} H_t &= H(C_t, X_t, u) \\ &= E_t\left[\frac{\pi_s}{\pi_t}C_s\right] = \exp\left\{-\int_t^{t+u} CDR_s ds\right\}C_t. \end{aligned} \tag{51}$$

We will refer to CDR_t as the consumption discount rate. Now define a stock S_t that gives a claim to consumption and therefore it pays a continuous stream of dividends C_t. The value of such a stock then obviously becomes:

$$S_t = \int_0^\infty H(C_t, X_t, u)du. \tag{52}$$

In equilibrium aggregate wealth must be equal to the value of the stock. The state-dependent consumption-wealth ratio therefore equals:

$$k(X_t) = \frac{C_t}{S_t} = \frac{C_t}{\int_0^\infty H(C_t, X_t, u)du} = \left(\int_0^\infty \exp\left\{-\int_t^{t+u} CDR_s ds\right\}du\right)^{-1}. \tag{53}$$

Using the expression for the consumption-wealth ratio, we can calculate the value function. At the optimum (see for example Munk (2015), Ch. 17), we have the envelope condition that $f_C = V_S$. Furthermore, we derived that $V(C_t, X_t) = \frac{g(X_t)C_t^{1-\gamma}}{1-\gamma}$. Using the chain rule we get:

$$V_S = V_C \frac{\partial C}{\partial S} = V_C k(X_t) = g(X_t) C_t^{-\gamma} k(X_t). \tag{54}$$

Also we have for the intertemporal aggregator:

$$f_C = \beta g(X_t)^{\frac{1/\epsilon-\gamma}{1-\gamma}} C_t^{-\gamma}. \tag{55}$$

Together this gives us:

$$g(X_t) = \left(\frac{k(X_t)}{\beta}\right)^{-\frac{1-\gamma}{1-1/\epsilon}}. \tag{56}$$

In "Appendix D" we derive an expression for the consumption discount rate CDR_t. Given the consumption discount rate, we can solve for the consumption-wealth ratio and therefore we know the value function.

Appendix D: Discount Rates

Appendix D.1: The Pricing Kernel

Duffie and Epstein (1992a) derive that the pricing kernel with stochastic differential utility equals $\pi_t = \exp\left\{ \int_0^t f_V(C_s, V_s) ds \right\} f_C(C_t, V_t)$. However, the pricing kernel has to be adjusted for the ambiguity aversion preferences. Chen and Epstein (2002) show that the pricing kernel in the ambiguity setting should be multiplied by the Radon–Nikodym derivative $\xi_t^{a^*,b^*}$ of the measure corresponding to the optimal a^* and b^*. $\xi_t^{a,b}$ is defined in (23).

We will start with deriving the explicit stochastic differential equation of the pricing kernel. First we calculate the derivatives of $f(C_t, V_t)$ with respect to C_t and V_t:

$$\begin{aligned} f_C(C, V) &= \frac{\beta C^{-1/\epsilon}}{\left((1-\gamma)V\right)^{\frac{1}{\zeta}-1}}, \\ f_V(C, V) &= \beta\zeta\left\{\left(1-\frac{1}{\zeta}\right)\left((1-\gamma)V\right)^{-\frac{1}{\zeta}} C^{1-1/\epsilon} - 1\right\}. \end{aligned} \tag{57}$$

Substituting $V_t = g(X_t)\frac{C_t^{1-\gamma}}{1-\gamma}$ into $f_C(C_t, V_t)$ and $f_V(C_t, V_t)$ we obtain:

$$\begin{aligned} f_C(C_t, V_t) &= \beta g(X_t)^{1-\frac{1}{\zeta}} C_t^{-\gamma}, \\ f_V(C_t, V_t) &= \beta\zeta\left\{ g(X_t)^{-\frac{1}{\zeta}}\left(1-\frac{1}{\zeta}\right) - 1\right\}. \end{aligned} \tag{58}$$

This gives:

$$\pi_t = \xi_t^{a^*,b^*} \exp\left(\int_0^t \beta\zeta\left(g(X_s)^{-\frac{1}{\zeta}}\left(1-\frac{1}{\zeta}\right)-1\right)ds\right)\beta g(X_t)^{1-\frac{1}{\zeta}} C_t^{-\gamma}. \tag{59}$$

Take the logarithm and write as a differential equation:

$$\begin{aligned} d\log(\pi_t) = {} & \beta\zeta\left(g(X_t)^{-\frac{1}{\zeta}}\left(1-\frac{1}{\zeta}\right)-1\right)dt - \gamma d\log(C_t) + d\log(\xi_t^{a^*,b^*}) \\ & + (1-\frac{1}{\zeta})d\log\left(g(X_t)\right). \end{aligned} \tag{60}$$

Apply Ito's lemma to $\log(C_t)$, $\log(\xi_t^{a^*,b^*})$ and $\log\left(g(X_t)\right)$ and substitute the results; this leads to the following differential equation:

$$\begin{aligned} d\log(\pi_t) = & \Big\{ \beta\zeta\Big(g(X_t)^{-\frac{1}{\zeta}}\big(1-\frac{1}{\zeta}\big)-1\Big) - \gamma\big(\mu-\frac{\sigma^2}{2}\big) + \lambda_t(1-a^*) \\ & + (1/\epsilon-\gamma)\frac{g_X(X_t)}{g(X_t)(1-\gamma)}\mu_X(X_t)\Big\}dt \\ & - \gamma\sigma dZ_t + \Big(\log(a^*b^*) + \big((b^*-1)\eta-\gamma\big)\log(1+J_t)\Big)dN_t. \end{aligned} \tag{61}$$

After applying Ito's lemma to $\log(\pi_t)$ we find:

$$\begin{aligned} d\pi_t = & \Big\{ \beta\zeta\Big(g(X_t)^{-\frac{1}{\zeta}}\big(1-\frac{1}{\zeta}\big)-1\Big) - \gamma\big(\mu-(\gamma+1)\frac{\sigma^2}{2}\big) + \lambda_t(1-a^*) \\ & + (1/\epsilon-\gamma)\frac{g_X(X_t)}{g(X_t)(1-\gamma)}\mu_X(X_t)\Big\}\pi_t dt + -\gamma\sigma\pi_t dZ_t \\ & + \Big(a^*b^*(1+J_t)^{(b^*-1)\eta-\gamma}-1\Big)\pi_{t-}dN_t. \end{aligned} \tag{62}$$

We can now substitute the HJB Eq. (43) into the pricing kernel. Several terms cancel out and we are left with:

$$\begin{aligned} d\pi_t = & \Big\{ -\beta - \frac{\mu}{\epsilon} + \Big(1+\frac{1}{\epsilon}\Big)\frac{\gamma}{2}\sigma^2 + \Big(\gamma-\frac{1}{\epsilon}\Big)\lambda_t^* \frac{-1}{b^*\eta+1-\gamma} \\ & + \lambda_t(1-a^*)\Big\}\pi_t dt - \gamma\sigma\pi_t dZ_t \\ & + \Big(a^*b^*(1+J_t)^{(b^*-1)\eta-\gamma}-1\Big)\pi_{t-}dN_t. \end{aligned} \tag{63}$$

Appendix D.2: The Interest Rate

Let B_t be the price of a risk-free asset with a return equal to the interest rate. By the no-arbitrage argument, the interest rate r_t should be such that $\pi_t B_t$ is a martingale,. Now write $d\pi_t = \mu_{\pi,t}\pi_t dt + \sigma_\pi \pi_t dZ_t + J_{\pi,t}\pi_{t-}dN_t$. The product with B_t then follows:

$$d\pi_t B_t = (r_t + \mu_{\pi,t})\pi_t B_t dt + \sigma_\pi \pi_t B_t dZ_t + J_{\pi,t}\pi_{t-}B_t dN_t. \tag{64}$$

This is a martingale if $r_t + \mu_\pi + \lambda_t E_t[J_{\pi,t}] = r_t + \mu_\pi + \lambda_t\Big(a^*\frac{b^*\eta}{b^*\eta-\gamma}-1\Big) = 0$. Therefore the equilibrium interest rate equals:

$$
\begin{aligned}
r_t =& -\mu_\pi - \lambda_t\left(a^* \frac{b^*\eta}{b^*\eta-\gamma} - 1\right) \\
=& \beta + \frac{\mu}{\epsilon} - \left(1+\frac{1}{\epsilon}\right)\frac{\gamma}{2}\sigma^2 - \left(\gamma - \frac{1}{\epsilon}\right)a^*\lambda_t \frac{-1}{b^*\eta+1-\gamma} \\
& - a^*\lambda_t\left(\frac{b^*\eta}{b^*\eta-\gamma} - 1\right).
\end{aligned} \tag{65}
$$

Substituting r_t into the pricing kernel gives:

$$
\begin{aligned}
d\pi_t =& \left\{ -r_t - \lambda_t\left(a^*\frac{b^*\eta}{b^*\eta-\gamma} - 1\right)\right\}\pi_t dt - \gamma\sigma\pi_t dZ_t \\
& + \left(a^* b^* (1+J_t)^{(b^*-1)\eta-\gamma} - 1\right)\pi_{t-} dN_t.
\end{aligned} \tag{66}
$$

Appendix D.3: The Equity Premium

Consider a stock that pays continuous dividends at a rate C_t and has ex-dividend price S_t. We denote the cum-dividend stock price by S_t^d. We use the expression for the consumption-wealth ratio in combination with the HJB-equation to derive the risk premium. An alternative derivation is to apply the no arbitrage condition. Using Eq. (53) we can write $S_t = \frac{C_t}{k(X_t)}$. The stock price then follows:

$$
\begin{aligned}
dS_t^d =& dS_t + C_t dt = \frac{1}{k(X_t)} dC_t - \frac{C_t}{k(X_t)^2} dk(X_t) + k(X_t) S_t dt \\
=& \left(\mu - \frac{k_X(X_t)}{k(X_t)}\mu_X(X_t) + k(X_t)\right) S_t dt + \sigma S_t dZ_t + J_t S_{t-} dN_t.
\end{aligned} \tag{67}
$$

From Eq. (67), we know that the drift of the stock equals $\mu_{S,t} = \mu - \frac{k_X(X_t)}{k(X_t)}\mu_X(X_t) + k(X_t)$. From (56) we have: $k(X_t) = \beta g(X_t)^{-\frac{1-1/\epsilon}{1-\gamma}}$. This gives: $\frac{k_X(X_t)}{k(X_t)} = -\frac{1-1/\epsilon}{1-\gamma}\frac{g_X(X_t)}{g(X_t)}$. Rewriting the HJB Eq. (43) gives:

$$
\begin{aligned}
\frac{1-1/\epsilon}{1-\gamma}\frac{g_X(X_t)}{g(X_t)}\mu_X(X_t) + k(X_t) =& \beta + (1/\epsilon - 1)\left(\mu - \frac{\gamma}{2}\sigma^2\right. \\
& \left. + a^*\lambda_t \frac{-1}{b^*\eta+1-\gamma}\right).
\end{aligned} \tag{68}
$$

Substituting this into $\mu_{S,t}$ gives:

$$
\mu_{S,t} = \mu - \frac{k_X(X_t)}{k(X_t)}\mu_X(X_t) + k(X_t) \tag{69}
$$

The risk premium is then equal to the excess return of the stock over the interest rate:

$$
\begin{aligned}
rp_t =& \mu_{S,t} + a^*\lambda_t \frac{-1}{b^*\eta+1} - r_t \\
=& \gamma\sigma^2 + a^*\lambda_t\left(\frac{-1}{b^*\eta+1} - \frac{b^*\eta}{b^*\eta+1-\gamma} + \frac{b^*\eta}{b^*\eta-\gamma}\right).
\end{aligned} \tag{70}
$$

Appendix D.4: Consumption Strips

Let $H_t = H(C_t, X_t, s-t) = E_t\left[\frac{\pi_s}{\pi_t} C_s\right]$ be the price of an asset that pays out a proportion of the aggregate consumption at time s. H_t is also called a consumption strip. Conjecture that $H(C_t, X_t, u) = \exp\left\{-\int_t^{t+u} CDR_s ds\right\} C_t$. u denotes the time to maturity of the consumption strip. Clearly, $H(C_t, X_t, 0) = C_t$. Applying Ito's lemma to H_t gives:

$$\begin{aligned} dH_t &= H_C dC_t + H_X dX_t - \frac{\partial H_t}{\partial u} dt = \frac{1}{C_t} H_t dC_t \\ &\quad - \frac{\partial}{\partial X_t}\left(\int_t^{t+u} CDR_s ds\right)\mu_X(X_t) H_t dt \\ &\quad + \frac{\partial}{\partial u}\left(\int_t^{t+u} CDR_s ds\right) H_t dt. \end{aligned} \tag{71}$$

We can calculate both derivatives:

$$\begin{aligned} \frac{\partial}{\partial X_t}\left(\int_t^{t+u} CDR_s ds\right)\mu_X(X_t) &= \frac{\partial}{\partial t}\left(\int_t^{t+u} CDR_s ds\right)\frac{\partial t}{\partial X_t}\mu_X(X_t) \\ &= \frac{\partial}{\partial t}\left(\int_t^{t+u} CDR_s ds\right) = CDR_{t+u} - CDR_t, \end{aligned} \tag{72}$$

$$\frac{\partial}{\partial u}\left(\int_t^{t+u} CDR_s ds\right) = CDR_{t+u}. \tag{73}$$

Therefore dH_t becomes:

$$dH_t = \left(\mu + CDR_t\right) H_t dt + \sigma H_t dZ_t + J_t H_{t-} dN_t. \tag{74}$$

Now define $dH_t = \mu_{H,t} H_t dt + \sigma H_t dZ_t + J_t H_{t-} dN_t$. By the no arbitrage condition, $\pi_t H_t$ must be a martingale:

$$\begin{aligned} d\pi_t H_t &= (\mu_{\pi,t} + \mu_H + \sigma\sigma_\pi)\pi_t H_t dt + (\sigma + \sigma_\pi)\pi_t H_t dZ_t \\ &\quad + \Big((1+J_t)(1+J_{\pi,t}) - 1\Big)\pi_{t-} H_{t-} dN_t. \end{aligned} \tag{75}$$

We can calculate the expectation of the jump term:

$$\begin{aligned} E_t[(1+J_t)(1+J_{\pi,t}) - 1] &= E_t[a^* b^* (1+J_t)^{(b^*-1)\eta+1-\gamma} - 1] \\ &= a^* \frac{b^*\eta}{b^*\eta + 1 - \gamma} - 1. \end{aligned} \tag{76}$$

Therefore $\pi_t H_t$ is a martingale if:

$$0 = \mu_\pi + \mu_H + \sigma\sigma_\pi + \lambda_t\left(a^* \frac{b^*\eta}{b^*\eta + 1 - \gamma} - 1\right). \tag{77}$$

Substituting μ_π, μ_H and $\sigma\sigma_\pi = -\gamma\sigma^2$ gives:

$$0 = \mu + CDR_t - r_t - \lambda_t\Big(a^* \frac{b^*\eta}{b^*\eta - \gamma} - 1\Big) - \gamma\sigma^2 + \lambda_t\Big(a^* \frac{b^*\eta}{b^*\eta + 1 - \gamma} - 1\Big). \tag{78}$$

Note that this implies that: $CDR_t = r_t + rp_t - (\mu + a^*\lambda_t \frac{-1}{b^*\eta+1})$. Lastly, we can substitute r_t and rp_t, which yields:

$$CDR_t = \beta + (1/\epsilon - 1)\Big(\mu - \frac{\gamma}{2}\sigma^2 + a^*\lambda_t \frac{-1}{b^*\eta + 1 - \gamma}\Big). \tag{79}$$

Appendix E: The Social Cost of Carbon

The Social Cost of Carbon is calculated as the derivative of the value function with respect to carbon emissions, scaled by instantaneous marginal utility. With a single carbon box, the marginal cost of increasing carbon emissions by one unit is the derivative of the value function with respect to the carbon concentration M_t: $\frac{\partial V_t}{\partial M_t}$. However, with multiple carbon boxes, emitting one unit of carbon leads to an increase of v_i units in box i, $i = 0, 1, 2, 3$. We slightly abuse notation and define $\frac{\partial}{\partial M_t} \equiv v_0 \frac{\partial}{\partial M_{0,t}} + v_1 \frac{\partial}{\partial M_{1,t}} + v_2 \frac{\partial}{\partial M_{2,t}} + v_3 \frac{\partial}{\partial M_{3,t}}$. In the derivation, we use the following relations: $g(X_t) = \Big(\frac{k(X_t)}{\beta}\Big)^{-\zeta}$, $\zeta = \frac{1-\gamma}{1-1/\epsilon}$, $k(X_t) = \Big(\int_0^\infty \exp\Big\{ - \int_t^{t+u} CDR_s ds\Big\} du\Big)^{-1}$ and $CDR_t = \beta + (1/\epsilon - 1)\Big(\mu - \frac{\gamma}{2}\sigma^2 + a^*\lambda_t \frac{-1}{b^*\eta+1-\gamma}\Big)$. Differentiation of the value function gives:

$$
\begin{aligned}
SCC_t &= -\frac{\partial V_t/\partial M_t}{f_C(C_t, V_t)} = -\frac{\partial\left(\frac{g(X_t)C_t^{1-\gamma}}{1-\gamma}\right)/\partial M_t}{\beta g(X_t)^{1-\frac{1}{\zeta}} C_t^{-\gamma}} = -\frac{\frac{\partial}{\partial M_t} g(X_t)}{(1-\gamma)\beta g(X_t)^{1-\frac{1}{\zeta}}} C_t \\
&= -\frac{\frac{\partial}{\partial M_t}\left(\frac{k(X_t)}{\beta}\right)^{-\zeta}}{(1-\gamma)\beta\left(\frac{k(X_t)}{\beta}\right)^{1-\zeta}} C_t = -\frac{-\zeta\left(\frac{k(X_t)}{\beta}\right)^{-\zeta-1}\frac{\partial}{\partial M_t}(\frac{k(X_t)}{\beta})}{(1-\gamma)\beta\left(\frac{k(X_t)}{\beta}\right)^{1-\zeta}} C_t \\
&= \frac{\zeta\frac{\partial}{\partial M_t} k(X_t)}{(1-\gamma)\beta^2\left(\frac{k(X_t)}{\beta}\right)^2} C_t = \frac{\frac{\partial}{\partial M_t} k(X_t)}{(1-1/\epsilon)k(X_t)^2} C_t \\
&= \frac{\frac{\partial}{\partial M_t}\left(\int_0^\infty \exp\left\{-\int_t^{t+u} CDR_s ds\right\} du\right)^{-1}}{(1-1/\epsilon)\left(\int_0^\infty \exp\left\{-\int_t^{t+u} CDR_s ds\right\} du\right)^{-2}} C_t \\
&= \frac{1}{1/\epsilon - 1}\frac{\partial}{\partial M_t}\int_0^\infty \exp\left\{-\int_t^{t+u} CDR_s ds\right\} du\ C_t \\
&= \frac{1}{1/\epsilon - 1}\int_0^\infty \frac{\partial}{\partial M_t}\exp\left\{-\int_t^{t+u} CDR_s ds\right\} du\ C_t \\
&= -\frac{1}{1/\epsilon - 1}\int_0^\infty \exp\left\{-\int_t^{t+u} CDR_s ds\right\}\frac{\partial}{\partial M_t}\left(\int_t^{t+u} CDR_s ds\right) du\ C_t \\
&= -\frac{1}{1/\epsilon - 1}\int_0^\infty \exp\left\{-\int_t^{t+u} CDR_s ds\right\}\int_t^{t+u}\frac{\partial}{\partial M_t} CDR_s ds\ du\ C_t \\
&= -\frac{1}{1/\epsilon - 1}\int_0^\infty \exp\left\{-\int_t^{t+u} CDR_s ds\right\} \\
&\quad \int_t^{t+u}\frac{\partial}{\partial M_t}(1/\epsilon - 1)a^*\lambda_T T_s \frac{-1}{b^*\eta + 1 - \gamma} ds\ du\ C_t \\
&= \int_0^\infty \exp\left\{-\int_t^{t+u} CDR_s ds\right\}\int_t^{t+u} a^*\lambda_T \frac{\partial T_s}{\partial M_t} ds \frac{1}{b^*\eta + 1 - \gamma} du\ C_t.
\end{aligned}
\tag{80}
$$

Appendix F: Calibration of Climate Model

See Table 5.

Table 5 Parameters for the Climate model

Par	Description	Value
E_0	Initial level of total emissions (in *GtC*, 2015)	10.45
g_0^E	Initial growth rate of emissions (2015)	0.017
g_∞^E	Long-run growth rate of emissions	− 0.02
δ_{g^E}	Speed of convergence of growth rate of emissions	0.0075
M_0	Initial carbon concentration compared to pre-industrial (in *GtC*, 2015)	263
M_{pre}	Pre-industrial atmospheric carbon concentration (in *GtC*)	588
$M_{0,0}$	Initial carbon concentration box 0 (in *GtC*, 2015)	139
$M_{1,0}$	Initial carbon concentration box 1 (in *GtC*, 2015)	90
$M_{2,0}$	Initial carbon concentration box 2 (in *GtC*, 2015)	29
$M_{3,0}$	Initial carbon concentration box 3 (in *GtC*, 2015)	4
$\delta_{M,0}$	Decay rate of carbon box 0	0
$\delta_{M,1}$	Decay rate of carbon box 1	0.0025
$\delta_{M,2}$	Decay rate of carbon box 2	0.027
$\delta_{M,3}$	Decay rate of carbon box 3	0.23
v_0	Fraction of emissions carbon box 0	0.217
v_1	Fraction of emissions carbon box 1	0.224
v_2	Fraction of emissions carbon box 2	0.282
v_3	Fraction of emissions carbon box 3	0.276
F_0^E	Initial level of exogenous forcing (in W/m^2, 2015)	0.5
F_∞^E	Long-run level of exogenous forcing (in W/m^2)	1
δ_F	Speed of convergence exogenous forcing	0.02
T_0	Initial surface temperature compared to pre-industrial (in °C, 2015)	0.85
T_0^{oc}	Initial ocean temperature compared to pre-industrial (in °C, 2015)	0.0068
κ	Speed of temperature transfer between upper and deep ocean	0.73
υ	Equilibrium temperature response to radiativeforcing	1.13
α	Equilibrium temperature impact of CO_2 doubling (in °C)	3.05
τ	Heat capacity of the surface	7.34
τ_{oc}	Heat capacity of the oceans	105.5

Appendix G: Simple Climate Model

Table 6 gives the results for the simple climate model with a Temperature Response to cumulative emissions of 1.75 $TtC/$°C. The SCC is a bit higher with the simple linear climate model compared to the non-linear climate model, especially in the case with $\epsilon = 1.5$.

Table 6 Social cost of carbon as function of risk aversion and ambiguity aversion

Social cost of carbon	$(\lambda_T, \eta) = (0.04, 61.5)$	$(\lambda_T, \eta) = (0.02, 30.25)$
$\epsilon = 1, \gamma = 0, \theta = 0$	184	184
$\epsilon = 1, \gamma = 5, \theta = 0$	200	219
$\epsilon = 1, \gamma = 5, \theta = 0.1$	354	400
$\epsilon = 1.5, \gamma = 0, \theta = 0$	419	419
$\epsilon = 1.5, \gamma = 5, \theta = 0$	411	446
$\epsilon = 1.5, \gamma = 5, \theta = 0.1$	680	753

In the simple climate model, a pulse of emissions today will lead to the same amount of warming in any future period. While in the non-linear climate model, the change in temperature because of an emissions pulse will somewhat decline over time due to decay of carbon from the atmosphere. Especially with a low discount rate (which is the case with $\epsilon = 1.5$), the simple model therefore leads to a SCC that is a bit higher.

The relative differences of adding risk aversion and ambiguity aversion are however very similar. So the different climate model mostly affects the level of the SCC in the no-risk and no-ambiguity case. But the increase of the SCC when risk aversion and ambiguity aversion are added does not change much.

Appendix H: Stochastic Emissions

The HJB-equation for this problem becomes:

$$\begin{aligned} 0 = \min_{(a,b)\ s.t.\ d(a,b)\leq\theta} & \Big\{ f(C_t, V_t^{\mathbb{Q}}) + V_C^{\mathbb{Q}}\mu C_t dt + \frac{1}{2}V_{CC}^{\mathbb{Q}}\sigma^2 C_t^2 + V_X^{\mathbb{Q}}\mu_X(X_t, C_t) \\ & + \lambda_t^{\mathbb{Q}} E_t^{\mathbb{Q}}\left[V^{\mathbb{Q}}\big((1+J_t)C_{t-}, X_t\big) - V^{\mathbb{Q}}(C_{t-}, X_t)\right] \Big\}. \end{aligned} \tag{81}$$

The main difference with the HJB-equation without stochastic emissions is that now, the drift of the state variables μ_X also depends on aggregate endowment C_t. It is therefore not possible anymore to substitute out the variable C_t. We thus have to solve a seven dimensional model numerically. We use the stochastic grid method to numerically solve the model, as described in Olijslagers (2021). Similar to value function iteration, the time step is discretized and the problem is solved backwards. The stochastic grid method simulates random grid points every time period and uses regressions with basis functions to approximate the value function. The main advantage is that this method can handle high-dimensional problems while avoiding the curse of dimensionality (computing time growing exponentially along with dimensionality) and that derivatives of the value function can be calculated easily. This is useful to calculate the social cost of carbon, and also to solve the first order conditions.

The first order conditions for optimal consumption are:

$$\begin{aligned} & E_t^{\mathbb{Q}}\left[V^{\mathbb{Q}}\big((1+J_t)C_{t-}, X_t\big) - V^{\mathbb{Q}}(C_{t-}, X_t)\right] - l_t\Big(\log(ab) + \frac{1}{b} - 1\Big) = 0, \\ & a\frac{\partial E_t^{\mathbb{Q}}\left[V^{\mathbb{Q}}\big((1+J_t)C_{t-}, X_t\big)\right]}{\partial b_t} - l_t a\frac{b-1}{b^2} = 0, \\ & \theta - (1-a) - a\Big(\log(ab) + \frac{1}{b} - 1\Big) = 0. \end{aligned} \tag{82}$$

Acknowledgements We thank Bob Pindyck, Nick Stern, Rick van der Ploeg, the late Martin Weitzman and seminar participants at the EEA meetings in 2020, DNB and the University of Amsterdam for helpful comments and pension administrator and wealth manager MN for financial support

References

Anderson EW, Hansen LP, Sargent TJ (2003) A quartet of semigroups for model specification, robustness, prices of risk, and model detection. J Eur Econ Assoc 1(1):68–123

Bansal R, Kiku D, Ochoa M (2019) Climate change risk. preprint. https://www.frbsf.org/economic-research/events/2019/november/economics-of-climate-change/files/P

Barnett M, Brock W, Hansen LP (2020) Pricing uncertainty induced by climate change. Rev Financ Stud 33(3):1024–1066

Barnett M, Brock W, Hansen LP (2022) Climate change uncertainty spillover in the macroeconomy. NBER Macroecon Annu 36(1):253–320

Barro RJ (2015) Environmental protection, rare disasters and discount rates. Economica 82(325):1–23

Berger L, Bosetti V (2020) Are policymakers ambiguity averse? Econ J 130(626):331–355

Bretscher L, Vinogradova A (2018) Escaping damocles' sword: endogenous climate shocks in a growing economy. Economics working paper series, 18

Brock WA, Hansen LP (2017) Wrestling with uncertainty in climate economic models

Cai Y (2020) The role of uncertainty in controlling climate change. In: Oxford research Encyclopedia of economics and finance

Cai Y, Lontzek TS (2019) The social cost of carbon with economic and climate risks. J Political Econ 127(6):2684–2734

Cai Y, Lenton TM, Lontzek TS (2016) Risk of multiple interacting tipping points should encourage rapid CO_2 emission reduction. Nat Clim Change 6(5):520–525

Chen Z, Epstein L (2002) Ambiguity, risk, and asset returns in continuous time. Econometrica 70(4):1403–1443

Cochrane JH (2009) Asset pricing, Revised. Princeton University Press, Princeton

Dell M, Jones BF, Olken BA (2012) Temperature shocks and economic growth: evidence from the last half century. Am Econ J Macroecon 4(3):66–95

Dietz S, Venmans F (2019) Cumulative carbon emissions and economic policy: in search of general principles. J Environ Econ Manag 96:108–129

Dietz S, van der Ploeg F, Rezai A, Venmans F (2021) Are economists getting climate dynamics right and does it matter? J Assoc Environ Resour Econ 8(5):895–921

Duffie D, Epstein LG (1992a) Asset pricing with stochastic differential utility. Rev Financ Stud 5(3):411–436

Duffie D, Epstein LG (1992b) Stochastic differential utility. Econom J Econom Soc 60:353–394

Ellsberg D (1961) Risk, ambiguity, and the savage axioms. Q J Econ 75:643–669

Epstein LG, Zin SE (1989) Substitution, risk aversion, and the temporal behavior of consumption and asset returns: a theoretical framework. Econom J Econom Soc 99:937–969

Gilboa I, Schmeidler D (1989) Maxmin expected utility with non-unique prior. J Math Econ 18(2):141–153. https://doi.org/10.1016/0304-4068(89)90018-9

Golosov M, Hassler J, Krusell P, Tsyvinski A (2014) Optimal taxes on fossil fuel in general equilibrium. Econometrica 82(1):41–88

Goosse H (2015) Climate system dynamics and modelling. Cambridge University Press. Retrieved from https://books.google.at/books?id=2-NBCgAAQBAJ

Hambel C, Kraft H, Schwartz E (2021) Optimal carbon abatement in a stochastic equilibrium model with climate change. Eur Econ Rev 132:103642

Hansen LP, Miao J (2018) Aversion to ambiguity and model misspecification in dynamic stochastic environments. Proc Natl Acad Sci 115(37):9163–9168

Hansen LP, Sargent TJ (2001) Robust control and model uncertainty. Am Econ Rev 91(2):60–66

Hansen LP, Sargent TJ (2008) Robustness. Princeton University Press, Princeton

Hope C (2006) The marginal impact of CO2 from PAGE2002: an integrated assessment model incorporating the IPCC's five reasons for concern. Integr Assess 6(1):1–16

IPCC (2013) Climate change 2013: the physical science basis. Contribution of working group i to the fifth assessment report of the intergovernmental panel on climate change [Book]. Cambridge, United Kingdom and New York, NY, USA: Cambridge University Press. Retrieved from www.climatechange2013.orghttps://doi.org/10.1017/CBO9781107415324

IPCC (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 [Book]. Cambridge, United Kingdom and New York, NY, USA: Cambridge University Press
Jensen S, Traeger CP (2014) Optimal climate change mitigation under longterm growth uncertainty: stochastic integrated assessment and analytic findings. Eur Econ Rev 69:104–125
Karydas C, Xepapadeas A (2019) Climate change risks: pricing and portfolio allocation (Technical report). CER-ETH-Center of Economic Research (CER-ETH) at ETH Zurich
Klibanoff P, Marinacci M, Mukerji S (2005) A smooth model of decision making under ambiguity. Econometrica 73(6):1849–1892
Knight FH (1921) Risk, uncertainty and profit. Hart, Schaffner and Marx, New York
Lemoine D, Traeger C (2014) Watch your step: optimal policy in a tipping climate. Am Econ J Econ Policy 6(1):137–66
Lemoine D, Traeger CP (2016) Ambiguous tipping points. J Econ Behav Organ 132:5–18
Liu J, Pan J, Wang T (2004) An equilibrium model of rare-event premia and its implication for option smirks. Rev Financ Stud 18(1):131–164
Maenhout PJ (2004) Robust portfolio rules and asset pricing. Rev Financ Stud 17(4):951–983
Maenhout PJ (2006) Robust portfolio rules and detection-error probabilities for a mean-reverting risk premium. J Econ Theory 128(1):136–163
Mattauch L, Millar R, van der Ploeg R, Rezai A, Schultes A, Venmans F et al (2018) Steering the climate system: an extended comment
Mehra R, Prescott EC (1985) The equity premium: a puzzle. J Monet Econ 15(2):145–161
Millner A, Dietz S, Heal G (2013) Scientific ambiguity and climate policy. Environ Resour Econ 55(1):21–46
Munk C (2015) Dynamic asset allocation
Nordhaus W (2014) Estimates of the social cost of carbon: concepts and results from the dice-2013r model and alternative approaches. J Assoc Environ Resour Econ 1(1/2):273–312
Nordhaus WD (2017) Revisiting the social cost of carbon. Proc Natl Acad Sci 114:1518–1523
Olijslagers S (2021) Solution methods for DSGE models in continuous time: application to a climate-economy model. University of Amsterdam, Amsterdam
Olijslagers S, van der Ploeg F, van Wijnbergen S (2023) On current and future carbon prices in a risky world. J Econ Dyn Control 146:104569
Pindyck RS (2017) The use and misuse of models for climate policy. Rev Environ Econ Policy 11(1):100–114
Ramsey FP (1928) A mathematical theory of saving. Econ J 38(152):543–559
Rezai A, Van der Ploeg F (2016) Intergenerational inequality aversion, growth, and the role of damages: Occam's rule for the global carbon tax. J Assoc Environ Resour Econ 3(2):493–522
Rudik I (2020) Optimal climate policy when damages are unknown. Am Econ J Econ Policy 12(2):340–73
Tol RS (2002) Estimates of the damage costs of climate change. Part 1: benchmark estimates. Environ Resour Econ 21(1):47–73
Traeger CP (2021) Uncertainty in the analytic climate economy. CEPR discussion paper
Traeger CP (2023) Ace—analytic climate economy. Am Econ J Econ Policy 15(3):372–406
Van Binsbergen JH, Fernández-Villaverde J, Koijen RS, Rubio-Ramírez J (2012) The term structure of interest rates in a DSGE model with recursive preferences. J Monet Econ 59(7):634–648
Van den Bremer TS, Van der Ploeg F (2021) The risk-adjusted carbon price. Am Econ Rev 111(9):2782–2810
Venmans F, Groom B (2021) Social discounting, inequality aversion, and the environment. J Environ Econ Manag 109:102479
Vissing-Jørgensen A, Attanasio OP (2003) Stock-market participation, intertemporal substitution, and risk-aversion. Am Econ Rev 93(2):383–391
Weitzman ML (2007) A review of the stern review on the economics of climate change. J Econ Lit 45(3):703–724
Weitzman ML (2012) GHG targets as insurance against catastrophic climate damages. J Public Econ Theory 14(2):221–244