Pricing Climate Risk*

Svenn Jensen[a] Christian Traeger[b]

Version of January 2026

Abstract: Climatic feedbacks make the warming response to greenhouse gas emissions highly uncertain. We derive a general analytic formula for the risk premium in the social cost of carbon (SCC). The formula decomposes how uncertainty shifts optimal mitigation through distinct channels of risk aversion, prudence, and damage convexity, and it extends standard asset-pricing logic by highlighting the dominant role of structural climate sensitivity. We show that policy-oriented Monte Carlo models overstate the SCC risk premium, while our responsive formula closely matches the dynamic programming solution in a DICE-based application. Quantitatively, the premium averages 20–30%, with substantially higher values only under strongly convex damages.

JEL Codes: Q54, Q00, D90, C63

Keywords: climate change, uncertainty, social cost of carbon, risk premium, precautionary savings, prudence, climate policy, dynamic programming, Monte Carlo, integrated assessment, DICE, recursive utility

*We are grateful for feedback from Buzz Brock, Raphael Calel, Ben Crost, Paul Fackler, Lars Hansen, Michael Hanemann, Bard Harstad, Larry Karp, David Kelly, Derek Lemoine, Matti Liski, Antony Millner, Mario Miranda, Eduardo Montoya, Rick van der Ploeg, Armon Rezai and the audiences at the annual meetings of the SAET, AERE, EAERE, SURED, workshops at FEEM, CREE, UCE^3, the University of East Carolina, UC Santa Barbara, and department seminars at the Universities of Berkeley, Gothenburg, and Oslo as well as at INSEAD, NMBU, Frisch Centre, OxCarre, Grantham Institute, and the SF Fed. Svenn Jensen thanks the Research Council of Norway for funding under grant 235573 of the Klimaforsk program. Related numeric results were circulated in Jensen's job market paper on Optimally Climate Sensitive Policy under Uncertainty and Learning, which has been replaced by the present paper (and an ongoing project on learning).

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

[b]Department of Economics, University of Oslo; ifo Institute, Munich; traeger@uio.no.

1 Introduction

Economic activity relies heavily on fossil fuel use. The resulting emissions alter our planet's climate, economic activity, and human well-being. A variety of feedback processes make the future temperature response to greenhouse gases highly uncertain. The World Economic Forum (2022, 2024) highlights extreme weather, a failure of climate action, and critical changes to Earth systems among the highest medium-run risks. Recent research identifies uncertainty about the temperature response to emissions as the climate uncertainty with the largest impact on economic policy (van den Bremer and van der Ploeg, 2021; Traeger, 2021). A central parameter in this debate is climate sensitivity, defined as the long-run temperature increase from a doubling of pre-industrial atmospheric CO_2 concentrations. The Intergovernmental Panel for Climate Change's estimate places expected climate sensitivity around 3°C, broadly in line with best-guess warming by 2100 under current emission projections (Masson-Delmotte et al., 2021; UN Environment Programme, 2024). Yet the panel emphasizes substantial uncertainty around this mean, with probability mass spread across both lower and much higher outcomes. This uncertainty directly governs not only the expected level of future warming but also how strongly each additional ton of CO_2 translates into higher temperatures.

Most existing models analyzing the optimal carbon price and its risk premium impose restrictive functional forms and rely on numerical approximations. Their estimates of the policy implications for climate uncertainty vary widely—from only a few cents in (Rudik and Lemoine, 2017) to potentially doubling the optimal tax (van den Bremer and van der Ploeg, 2021; Traeger, 2021)—underscoring limited knowlege about the true policy implications. We develop a general analytic model that identifies and quantifies the conceptual drivers of the climate risk premium in the social cost of carbon. Without committing to narrow parametric assumptions, we decompose how uncertainty shifts optimal mitigation through distinct channels driven by risk aversion, prudence, the convexity of production damages, and the responsiveness of future policy. This perspective clarifies the classic precautionary savings logic - where prudence matters - versus asset-pricing style covariance terms, and it pinpoints which characteristics of damages and preferences actually drive the premium. The result is a transparent map from climate uncertainty to policy, designed to sharpen economic intuition and guide both structural modeling and empirical research on climate damages.

Our decomposition naturally connects to an asset-pricing view of risk premia. In that framework, an asset's value reflects how its payoff covaries with the stochastic discount factor, and, in related formulations, with exposure to market risk. Applied to climate policy, this perspective has led researchers to study how damages correlate with consumption risk,

or with stochastic discount factors implied by richer preference structures. We show that this view is incomplete. Beyond the covariance of damages with marginal utility, a quantitatively more important channel arises from the correlation of marginal damages with the realization of climate sensitivity. When climate sensitivity is high, an additional unit of emissions both (i) causes a larger marginal temperature increase and (ii) occurs at a higher temperature level that amplifies damages and reduces consumption. The covariance between these two effects—the simultaneous dependence of marginal warming and marginal damages on climate sensitivity—accounts for the dominant share of the risk premium. We formalize this reasoning in an extended pricing formula with three distinct covariance terms, expressed in terms of the model's fundamentals and complemented by prudence effects reminiscent of precautionary savings.

A further implication concerns identification. The classic covariance between marginal damages and the stochastic discount factor can, in principle, be proxied from market data.[1] By contrast, our two novel covariance terms hinge on structural objects and are not directly recoverable from return covariances, so we discipline them structurally. Our decomposition provides a simple structural proxy: the new covariances reduce to a few primitives — relative risk aversion, damage convexity, and elasticities translating climate sensitivity into consumption and temperature.

Policy-oriented projects, such as the recent revision of the US federal Social Cost of Carbon (SCC), rely on "Monte Carlo" models that sample deterministic worlds to estimate the mean SCC (Rennert et al., 2022; EPA, 2023). While convenient, this approach neglects that future policymakers will strengthen or relax mitigation as new information about climate sensitivity becomes available, a response that requires complex dynamic programming to capture. We extend our analytic formula to approximate this anticipation of future policy adjustment. Using a DICE-based IAM, we show that our "responsive formula" closely matches the full dynamic programming solution with learning, while revealing that the Monte Carlo approach overstates the risk premium from temperature uncertainty by about 20%. Our approach thus offers a tractable way to embed future policy responsiveness into IAMs without the computational burden of high-dimensional stochastic control.

Our quantitative application finds a baseline risk premium of about 26%, with results generally ranging between 20% and 30% across variations in time preference, risk aversion,

[1]Conceptually, this would be a CCAPM-style exercise that relates asset payoffs to a consumption-based pricing kernel using variation from realized temperature shocks and macro covariation; it does not identify belief updates about structural climate sensitivity. By contrast, our two new covariance terms load on belief shocks about climate sensitivity and the warming-per-ton response. In the absence of traded claims that span such shocks (e.g., liquid long-dated temperature- or SCC-linked securities), they are not directly identifiable from return co-movements. We do not pursue empirical identification here; instead we quantify these channels in a stochastic IAM and show how a simple structural proxy can discipline them.

Epstein–Zin preferences, and other economic updates. Substantially larger premia arise only under major changes to the damage function, particularly when higher damage convexity raises the optimal carbon tax by more than 50%. The dominant driver of the premium is the structural interaction of climate sensitivity with the convexity of production damages. By contrast, standard risk aversion contributes only modestly, and prudence effects are negligible.

Literature. The literature studying how uncertainty shapes climate-economy interactions has expanded considerably over the past two decades. Jensen and Traeger (2024) provide a comprehensive review of the state of knowledge.

Kelly and Kolstad (1999) initiated this line of research by studying uncertainty in climate feedbacks in a stochastic dynamic optimization setting. Their pioneering analysis, building on the DICE model by Nordhaus (1993), set a template for subsequent work by providing the first quantitative assessment of how climate uncertainty affects optimal policy and learning about climate sensitivity. Leach (2007) extended this by adding a second uncertain parameter, showing that learning may slow to economically negligible rates. Kelly and Tan (2015) emphasized tail risk by introducing a right-skewed, fat-tailed climate sensitivity distribution. The fat tail initially induces a strong policy response, yet their learning dynamics quickly rule out (or confirm) extreme outcomes. In fixed-policy settings, simulating alternative uncertainty realizations sheds light on the distribution of potential outcomes. Monte Carlo studies, which often use rich structural models that do not permit stochastic dynamic optimization, consistently confirm the importance of climate uncertainty, typically finding sizable upward adjustments to the social cost of carbon (Nordhaus, 2008; Rennert et al., 2022; EPA, 2023).

The numerical component of the present paper is closely related to Kelly and Kolstad (1999). Using an updated DICE model and leveraging modern computational power, we can calculate the social cost of carbon under climate uncertainty with substantially improved numerical precision. Where our qualitative results differ, we conjecture that this enhanced precision is responsible. We also quantify the impact of responding optimally to the uncertainty, by comparing fixed and optimal policies in the same framework.

Given the complexity of integrating climate and economic dynamics into a unified stochastic system, and the many defensible modeling choices, quantitative estimates of uncertainty impacts vary widely. Reported effects range from negligible (Rudik and Lemoine, 2017) to increases that potentially double the social cost of carbon (van den Bremer and van der Ploeg, 2021). Qualitative contradictions also arise: for example, Daniel et al. (2019) find a declining social cost of carbon as uncertainty resolves, a result attributed to Epstein-Zin pref-

erences and their implied desire for early resolution of uncertainty (Giglio et al., 2021). Using the same preference structure but a less stylized model, we obtain the opposite qualitative pattern. These challenges explain the desire to clarify the underlying mechanisms, which has spurred a resurgence of analytic and semi-analytic work. Some recent analytic models are highly stylized, in the spirit of Heal (1984). Dietz et al. (2018) apply asset-pricing logic to such a model and emphasize the correlation between the marginal utility of consumption and marginal climate damages, which they call the 'climate beta.' Their complementing Monte Carlo analysis in a richer model version shows that uncertainty in temperature responses substantially raises the SCC under standard social preferences. Lemoine (2021a) likewise highlights the central role of this correlation and identifies several channels through which climate uncertainty affects the SCC, again finding quantitatively meaningful effects.

Compared to this literature, the present paper imposes fewer simplifications on the climate-economy dynamics. This matters: the correlation emphasized above is only one of five channels in our model through which climate uncertainty influences the social cost of carbon.

A parallel line of work seeks analytic tractability in full climate-economy models without relying on the linear-quadratic structure that enabled earlier analytic contributions (Hoel and Karp, 2001, 2002). Building on Golosov et al. (2014), who introduce a new class of more flexible analytic IAMs, Traeger (2015, 2021) shows how to analyze uncertainty under less restrictive assumptions using affined stochastic processes and incorporating fait-tailed uncertainty. Although such models still impose functional-form constraints, many assumptions are now more consistent with climate science.

van den Bremer and van der Ploeg (2021) advance the field by introducing perturbation methods to the analysis of climate uncertainty. Starting from a closed-form exogenous growth economy, they treat the climate as a first-order perturbation to this economy. When climate damages are sufficiently small, and the economy follows an AK-structure, this approach yields a good approximation to a full climate-economy model and enables analytic expressions for the SCC under various uncertainties. Quantitatively, they find substantial increases in the SCC. Qualitatively, their semi-analytic formula requires the climate sensitivity distribution to be skewed: without it, uncertainty has no effect.

In contrast, the present paper models the climate system explicitly and avoids imposing any restrictive functional forms for its qualitative analysis. This enables us, for example, to disentangle the roles of risk aversion and prudence, which are captured by a single parameter under standard expected utility. Our general expression for the SCC characterizes the qualitative drivers of uncertainty effects without relying on specific distributional assumptions about climate sensitivity. While our small risk approximation does not allow us to quantify

the role of skewness or higher moments, we nonetheless find significant effects of climate uncertainty even in their absence.

Uncertainty about future economic conditions and about the severity of climate damages is often of comparable quantitative importance to climate uncertainty, yet these sources operate through different channels. Growth uncertainty, for instance, directly affects damages and social welfare but does not alter the marginal temperature response to an additional ton of carbon. Jensen and Traeger (2014) examine the optimal SCC under growth uncertainty in a stylized stochastic IAM, highlighting the roles of precautionary savings and the curvature of the damage function. Extending the framework to Epstein-Zin preferences, they show that the impact of growth uncertainty becomes negative when the elasticity of intertemporal substitution (EIS) exceeds unity, as commonly used in the long-run risk literature. In contrast, the present paper finds that the risk premium associated with climate uncertainty remains positive even for EIS values greater than one, underscoring a fundamental difference in how growth and climate uncertainties propagate through the SCC. Cai and Lontzek (2019) confirm and extend the numeric findings of Jensen and Traeger (2014) using a richer implementation of DICE, further incorporating potential regime shifts in the damage function. van den Bremer and van der Ploeg (2021) derive an analytic representation of growth uncertainty effects within their perturbation-based SCC formula, obtaining the same qualitative dependence on the EIS as Jensen and Traeger (2014). See Jensen and Traeger (2024) for a more extensive discussion.

Rudik (2020) studies damage uncertainty in a recursive DICE framework, deriving a semi-analytic SCC expression that decomposes its effects into uncertainty, insurance, and precautionary savings motives. In his specification, damage uncertainty yields a modest positive impact on the current SCC. In contrast, Lemoine (2021a) finds that damage uncertainty can exert the largest effect among climate, damage, and growth uncertainties, though this result arises in a deliberately stylized setting and under comparatively extreme damage assumptions. Just as with their treatment of climate uncertainty, the effect of damage uncertainty in van den Bremer and van der Ploeg (2021) depends on the skewness of the underlying distribution. In Lemoine (2021a) all three uncertainties - climate, damages, and growth - operate through the same channels, due to the stylized nature of the model.

With analytic tractability as their main objective, van den Bremer et al. (2024) extend the perturbation approach and develop a parameterized SCC expression that incorporates not only damage uncertainty but also climate catastrophes - large, discrete shocks to output or welfare - and cascading tipping points. Catastrophic climate risks have also been examined through asset-pricing and macroeconomic lenses, where Poisson jump processes embedded in continuous-time Lucas-tree economies allow for parsimonious modeling of rare but severe

events. For a comprehensive overview of this growing research area at the intersection of climate economics and asset pricing, see Giglio et al. (2021). A related branch of the literature examines climate tipping points more directly, with a richer structural description of climate and economy. Following the seminal contribution of Keller et al. (2004), both single and multiple tipping points have been studied using numerical and semi-analytic methods (Lemoine and Traeger, 2014, 2016a; Cai and Lontzek, 2019; Dietz et al., 2021). Not all uncertainties are readily quantifiable. Some are better characterized as Knightian or ambiguous, and social preferences may exhibit aversion to ambiguity above and beyond standard risk aversion. Such ambiguity has been explored from various angles within climate-economy models (Millner et al., 2013; Lemoine and Traeger, 2016b; Barnett et al., 2020).

2 A Generic Integrated Assessment Model

This section introduces a simple yet general integrated assessment model of climate change, linking a growing world economy to a climate model through greenhouse gas emissions and the economic impacts of climate change (see Figure 1).

2.1 Analytic Model Structure

Figure 1 presents a schematic of our model. The right-hand side depicts a standard Ramsey growth framework, where world output is given by $F(K_t, T_t, E_t, t)$, a function of capital K_t, temperature T_t, emissions E_t, and exogenous processes such as labor and technology evolving with time t. The temperature increase T_t, measured in degrees Celsius above 1900 levels, reduces economic productivity. Output is divided between consumption C_t and investment in capital.

We model carbon dioxide emissions as an input into production, representing an energy sector in which emissions trade off against capital, labor, and a freely available renewable resource such as wind or solar power. Appendix D.1 shows the equivalence of this formulation to the abatement framework used in integrated assessment models like DICE. Emissions accumulate as atmospheric carbon M_t. Together with other greenhouse gases, atmospheric carbon traps outgoing infrared radiation, shifting the planet's energy balance and initiating warming that is amplified by feedback processes.

Climate sensitivity s measures the medium- to long-run increase in global temperature from a doubling of pre-industrial atmospheric CO_2, with the current best estimate around 3°C. The approximate "climate change law" states that each subsequent doubling (or fraction

thereof) yields an additional increase of about 3°C.[2] Reaching the new equilibrium temperature takes decades to centuries.[3] Rising temperatures reduce economic output through damages to production. Social welfare is defined as the discounted sum of per-period utility $u_t(C_t)$, and β denotes the pure time preference discount factor.

In summary, the social planner solves the Bellman equation

$$V(K_t, M_t, T_t, t) = \max_{C_t, E_t} u_t(C_t) + \beta \, \mathbb{E}_t \, V(K_{t+1}, M_{t+1}, T_{t+1}, t+1) \tag{1}$$

subject to the equations of motion

$$\begin{aligned} K_{t+1} &= (1-\delta)K_t + F(K_t, T_t, E_t, t) - C_t \\ M_{t+1} &= M_{t+1}(M_t, E_t) \quad \equiv M(M_t, E_t, t) \\ T_{t+1} &= T_{t+1}(T_t, M_{t+1}) \quad \equiv T(T_t, M_{t+1}, s). \end{aligned} \tag{2}$$

At this level of generality, each equation of motion can be subjected to stochastic shocks. We focus on the uncertainty generated by equation (2), which results from both the uncertainty in climate sensitivity s and stochastic temperature shocks. Two alternative characterizations of this uncertainty are detailed in Assumptions 1 and 1' below.

We assume that today's emissions add one-for-one to next period's carbon stock, so that $\frac{\partial M_1}{\partial E_0} = 1$.[4] For simplicity, we use M_t to denote both the atmospheric carbon stock and the function $M(\cdot, t)$ describing its evolution over time. This notation allows us to write the impact of a change in the carbon stock at period $\tau < t$ on the carbon stock at period t as

$$\frac{\partial M_t}{\partial M_\tau} = \prod_{i=\tau+1}^{t} \frac{\partial M_i}{\partial M_{i-1}} \equiv \prod_{i=\tau+1}^{t} \frac{\partial M(M_{i-1}, E_{i-1}, i-1)}{\partial M_{i-1}}.$$

Similarly, the impact of a change in the warming level at period $\tau < t$ on the warming level at period t can be written as

$$\frac{\partial T_t}{\partial T_\tau} = \prod_{i=\tau+1}^{t} \frac{\partial T_i}{\partial T_{i-1}} \equiv \prod_{i=\tau+1}^{t} \frac{\partial T(T_{i-1}, M_i, i-1)}{\partial T_{i-1}}.$$

We assume the optimization problem has a well-defined solution and that the discounted marginal value at period τ of a change in atmospheric carbon or temperature at time $t < \tau$

[2] Radiative forcing is logarithmic in atmospheric CO_2, and warming is approximately proportional to radiative forcing. Climate sensitivity is the proportionality constant. Our model also assumes that short-run temperature responses are proportional to climate sensitivity, though smaller in magnitude, consistent with models such as DICE and MAGICC (Traeger, 2014).

[3] We note that the TCRE model (Dietz and Venmans, 2019) suggests a rapid temperature response, relying on an approximate canceling of effects from carbon removal and warming delay to link temperature closely to *cumulative emissions* along certain paths. Our focus, however, is on temperature adjustment to atmospheric *concentrations*, which takes time and generates uncertainty about climate sensitivity. If TCRE were exact, climate sensitivity could be determined instantly, but significant uncertainty remains.

[4] Without this assumption, the SCC would contain an additional factor of $\frac{\partial M_1}{\partial E_0}$.

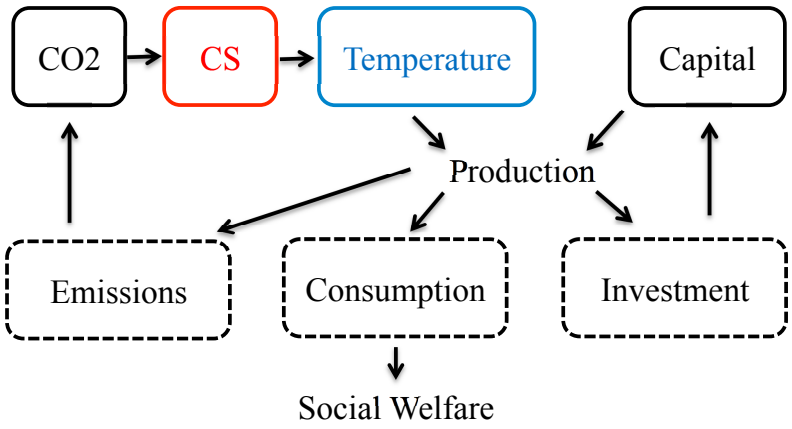


Figure 1: The main relations in the climate-economy model. Dashed rectangles represent control variables. Solid rectangles depict the main state variables of the system. Climate sensitivity ('CS') is uncertain. The model with learning represents climate sensitivity by a Bayesian prior (2 state variables). Temperature is stochastic.

approaches zero as τ approaches infinity: $\lim_{\tau\to\infty} \beta^{\tau-t}\, \mathbb{E}_t \frac{\partial V}{\partial M_\tau}\frac{\partial M_\tau}{\partial M_t}$ and $\lim_{\tau\to\infty} \beta^{\tau-t}\, \mathbb{E}_t \frac{\partial V}{\partial T_\tau}\frac{\partial T_\tau}{\partial T_t}$. Our analytic results rely on this general model structure. Assumption 1 in Section 3 refine the temperature's equation of motion (2). AppendixD.1 details our numeric DICE implementation as a special case of this model.

2.2 The Social Cost of Carbon

An optimal mitigation policy equates the current benefit of releasing a ton of CO_2 with the present value of the resulting future damages, known as the social cost of carbon (SCC). We denote by $\frac{\partial M_\tau}{\partial E_0}$ the effect of an additional ton of emissions today on the carbon stock in period τ. This increase in M_τ raises temperatures in later periods $t > \tau$ by $\frac{\partial T_t}{\partial M_\tau}$ through radiative forcing and feedbacks. Higher temperatures reduce output $Y_t = F_t(K_t, T_t, E_t, t)$, with marginal damages $-\frac{\partial F_t}{\partial T_t}$. The resulting loss lowers period-t welfare in proportion to $u_t'(c_t)$. Summing these discounted welfare losses over the infinite horizon and expressing them in present-day consumption units yields an analytic expression for the SCC.

Proposition 1. *The social cost of carbon of the generic integrated assessment model is*

$$SCC_0 = -\,\mathbb{E}_0 \sum_{t=1}^{\infty} \sum_{\tau=1}^{t} \underbrace{\beta^t \frac{u_t'(c_t)}{u_0'(c_0)}}_{\substack{\text{consumption} \\ \text{discount} \\ \text{factor}}} \underbrace{\frac{\partial F_t}{\partial T_t} \frac{\partial T_t}{\partial M_\tau} \frac{\partial M_\tau}{\partial E_0}}_{\substack{\text{marginal} \\ \text{emission} \\ \text{damage}}} . \tag{3}$$

Under certainty, the optimal carbon tax is the SCC evaluated along the optimal trajectory. In the stochastic setting, the optimal carbon tax evaluates equation (3) *on a decision tree where every period's policy on every branch is optimized conditional on past realizations.*

We obtain a double sum because of the delay between initial radiative forcing and the eventual temperature response.[5] Numerical integrated assessment models typically calculate the SCC without analytic expressions, either by perturbing emissions (e.g. FUND) or by relying on software to compute shadow values (e.g. DICE). Appendix A derives equation (3) from the Bellman equation. Our intuitive derivation above held policies fixed, while the Bellman-based derivation incorporates optimal future responses. The two approaches coincide along the optimal trajectory by the envelope theorem: the decision maker's future adjustment to a marginal ton of emissions does not change the program's value in first order, so the marginal value of an emission unit remains the same.

3 The Price of Uncertainty

This section presents an analytic formula for and interpretation of the SCC's response to uncertainty governing the climate's sensitivity to greenhouse gas emissions. This uncertainty results from a variety of feedback processes: more heat implies more water vapor, which is itself an important greenhouse gas;[6] higher temperatures melt glaciers and ice caps, reducing the reflectivity of our planet's surface, further increasing its uptake of solar radiation; soils currently frozen will release large amounts of methane, a powerful greenhouse gas; other feedbacks change the vertical temperature profile in the atmosphere and cloud formation, again feeding back onto our surface temperature. As a consequence of these feedbacks, the actual warming resulting from a given greenhouse gas concentration is highly uncertain.

As a result of this uncertainty, a deterministic SCC formula by itself is of limited practical use. The Interagency Working Group on the Social Cost of Carbon (2013); Interagency Working Group on Social Cost of Carbon (2016); Interagency Working Group on Social

[5] Gerlagh and Liski (2018) and Traeger (2023) provide examples of such double sums and show that in special cases they can be written in closed form by inverting transition matrices.

[6] Currently, our planet has a natural greenhouse effect of approximately 18°C mostly due to water vapor.

Cost Greenhouse Gases (2021) and the US federal SCC update (EPA, 2023) responded by simulating the underlying deterministic integrated assessment models for a wide range of climatic responses, deriving a probability distribution for the SCC and reporting its mean and the 95^{th} percentile. Yet, the optimal price on carbon has to be a single value that anticipates the stochastic evolution of the future, and it generally differs from the average optimal response across different deterministic worlds. We now derive a formula that quantifies and interprets the SCC's response to uncertainty. Section 3.1 reviews the precautionary savings concept of prudence and extends it to models with economic production. Section 3.2 states the formula and discusses the different channels that increase (or decrease) the optimal carbon tax under uncertainty about the climate's sensitivity to greenhouse gas emissions. Section 3.3 shows the related but distinct impact of atmospheric temperature fluctuations on the SCC, which is sometimes confused with climate uncertainty.

3.1 Precautionary Motives and the Climate-Economy Interaction: Setting the Stage

Uncertainty over the climate sensitivity s to greenhouse gas emissions affects the future climate (temperature response), which in turn affects future production (economic response) and human well-being (welfare response). Based on equation (3), a given year's contribution to the current SCC is proportional to

$$u_t'\big(C_t(T_t(s))\big)\ \left[-\frac{\partial F_t}{\partial T_t}(T_t(s))\right]\ \left[\sum_{\tau=1}^{t}\frac{\partial T_t}{\partial M_\tau}(s)\frac{\partial M_\tau}{\partial E_0}\right],$$

where the dependencies on the climate sensitivity s are explicitly marked. Starting from the end, a higher climate sensitivity s means an *additional* unit of emissions results in a larger *marginal temperature increase* $\frac{\partial T_t}{\partial M_\tau}(s)$. This effect is generally linear in climate sensitivity. Second, a higher climate sensitivity raises the temperature *level* due to the *prevailing concentration* of atmospheric greenhouse gases, i.e., $T_t(s)$ increases. This higher temperature level increases the *marginal production damage* from an additional unit of temperature increase $\frac{\partial F_t}{\partial T_t}(T_t(s))$. Third, higher temperatures and damages lead to lower consumption levels, which imply a higher *marginal welfare loss* from a given consumption unit $u_t'(c_t(s))$.

By Jensen's inequality, uncertainty increases the SCC and the optimal carbon tax if the right side of equation (3) is convex in the uncertain climate sensitivity. Intuitively, this convexity means that a high realization of climate sensitivity increases the social cost of a marginal emission unit more than a low realization decreases it. Unlike a simple precautionary savings model, climate sensitivity affects the right side of equation (3) through various

channels. We now introduce the terminology to discuss these uncertainty contributions.

Welfare response. We briefly recall the intuition behind the precautionary savings model of Leland (1968), as formalized by Kimball (1990). Consider an agent facing stochastic shocks to her savings returns. A risk-averse agent dislikes such shocks, but risk aversion alone does not change behavior. Empirical evidence shows that absolute risk aversion decreases with wealth. Then the agent can reduce expected utility losses—the pain from risk exposure—by saving more and reaching a higher wealth level. Kimball (1990) characterizes this mechanism in terms of prudence, analogous to the way Arrow (1965) and Pratt (1964) define risk aversion:

$$\text{RRA} = -\frac{u''(c_t)}{u'(c_t)}c_t \quad \text{and} \quad \text{Prud} = -\frac{u'''(c_t)}{u''(c_t)}c_t.$$

Relative risk aversion (RRA) reflects the curvature of utility, while relative prudence captures how risk aversion changes with consumption. In our notation, RRA and Prud may vary with c_t; in the quantitative analysis, however, we assume constant relative risk aversion, which makes these measures independent of the consumption level.

Economic response. In the precautionary savings example, uncertainty directly affects consumption and marginal utility. In the climate-economy setting, uncertainty also affects economic productivity through temperature. We therefore define second- and third-order measures of the production function in temperature, analogous to Arrow–Pratt risk aversion and Kimball prudence:

$$\text{Dam}_2 = \frac{-\frac{\partial^2 F_t}{\partial T_t^2}}{-\frac{\partial F_t}{\partial T_t}}\,T_t \quad \text{and} \quad \text{Dam}_3 = \frac{-\frac{\partial^3 F_t}{\partial T_t^3}}{-\frac{\partial^2 F_t}{\partial T_t^2}}\,T_t.$$

Here Dam_2 captures how marginal damages $-\frac{\partial F_t}{\partial T_t}$ change with temperature (damage convexity), while Dam_3 reflects how this convexity itself rises with higher temperatures. For simplicity, our notation suppresses dependencies on production levels and prevailing temperatures.

Temperature response. For clarity of exposition, we impose a condition satisfied by our DICE-based numerical integrated assessment model (see equation 27 in Appendix D.2):

Assumption 1. *The temperature evolution is linear in climate sensitivity s:*

$$T_{t+1} = \gamma T_t + s\,\Gamma_1(M_{t+1}) + \Gamma_2(t) + \epsilon_t, \tag{4}$$

where $\gamma \in (0,1)$, s is uncertain, Γ_1 and Γ_2 are real functions, and $(\epsilon_t)_{t\in\mathbb{N}}$ is an i.i.d. process of mean-zero temperature shocks.

Climate sensitivity is the proportionality factor between radiative forcing and temperature, capturing the net heat influx from greenhouse gases that gradually warms the planet.[7] The AR(1) process in equation (4) models the stochastic and time-dependent nature of heating. Corollary 1 in Appendix B presents a more general case without Assumption 1. *Because temperature evolves linearly in climate sensitivity, no additional convexity or prudence-type curvature measures are required for this direct temperature channel.*

Investment and emission response. We first consider the policy-oriented literature that applies Monte Carlo analyses of the SCC under *fixed policy scenarios*.[8] To provide a benchmark, we start from a hypothetical *deterministic* scenario in which climate sensitivity and temperature shocks are held constant at their expected values, for example 3°C for climate sensitivity and zero for shocks.

Assumption 2. *The policy-maker sets investment and emissions to the deterministically optimal levels.*

The essence of this assumption is that the current policy-maker does not anticipate future adjustments in response to new information. Whether the policy is fixed at the deterministic optimum or at another baseline path is secondary; our choice minimizes arbitrariness and simplifies the analysis. Section 4.2 will relax this assumption and incorporate policy responsiveness. A key tool for that step is the set of elasticities introduced next.

Response elasticities or "uncertainty conversion factors". Uncertainty is governed by climate sensitivity, but we must still translate the magnitude of this uncertainty into the implied variation in temperature and consumption. We capture these responses through two elasticities. The sensitivity of consumption to climate sensitivity is

$$\epsilon_{c,s} = -\frac{dc}{ds}\frac{s}{c} > 0,$$

and the sensitivity of global average surface temperature to climate sensitivity is

$$\epsilon_{T,s} = \frac{dT}{ds}\frac{s}{T} > 0.$$

Our notation suppresses time and other dependencies.[9] The temperature elasticity $\epsilon_{T,s}$ follows directly from recursive evaluation of equation (4) and climate dynamics. Section 4.2 explains how to derive $\epsilon_{T,s}$ numerically when policy responsiveness is reintroduced.

[7] It describes the change in net heat influx assuming all other conditions are at preindustrial levels. The system eventually converges to a new equilibrium where temperatures are higher, and energy in- and outflows are balanced.

[8] Such Monte Carlo runs must assume a fixed policy because (individually) re-optimizing in each model run and then averaging can produce incorrect signs for the climate policy risk premium (Crost and Traeger, 2013).

[9] Page 45 in Appendix B provides closed-form expressions for these elasticities under Assumption 2 and a generalization.

In deriving equation (3), the envelope theorem ensures that marginal shifts between consumption, investment, and abatement do not affect welfare. However, the impact of risk is non-marginal. , and Jensen's inequality does not provide a convenient analog for dynamic equations. To address this,

3.2 Uncertain Climate Sensitivity

We now present our formula for the SCC's climate sensitivity risk premium. To isolate structural uncertainty, we set the temperature shocks ϵ_t in equation (4) to zero. The next section derives the additional premium from stochastic temperature shocks. We fix expected climate sensitivity at its best-guess value and introduce a probability distribution over its true value, capturing epistemic uncertainty about the underlying parameter. Appendix B derives the following proposition from the Bellman equation and, in Corollary 1, extends it to cases without Assumption 1 and with a variation of Assumption 2.

Proposition 2. *Under Assumptions 1 and 2, uncertainty over climate sensitivity increases the SCC contribution from period t if and only if*

$$X_t(\cdot) \equiv \underbrace{\mathrm{RRA}\,\epsilon_{c,s}\big[2}_{\substack{\text{direct}\\ \text{risk}\\ \text{aversion}}} + \underbrace{\mathrm{Prud}\,\epsilon_{c,s}}_{\substack{\text{welfare}\\ \text{prudence}}} + \underbrace{3\mathrm{Dam}_2\epsilon_{T,s}\big]}_{\substack{\text{welfare}\\ \text{economy}\\ \text{interaction}}} + \underbrace{\mathrm{Dam}_2\,\epsilon_{T,s}\big[2}_{\substack{\text{direct}\\ \text{damage}\\ \text{convexity}}} + \underbrace{\mathrm{Dam}_3\,\epsilon_{T,s}\big]}_{\substack{\text{economy}\\ \text{prudence}}} \tag{5}$$

is greater than zero, where arguments and their period-dependence are suppressed. Moreover, under a small risk approximation, the climate uncertainty premium is

$$\Delta SCC_0 \approx \sum_{t=1}^{\infty}\sum_{\tau=1}^{t} \underbrace{\beta^t \frac{u_t'(c_t)}{u_0'(c_0)}}_{\substack{\text{consumption}\\ \text{discount}\\ \text{factor}}} \underbrace{\frac{\partial F_t}{\partial T_t}\frac{\partial T_t}{\partial M_\tau}\frac{\partial M_\tau}{\partial E_0}}_{\substack{\text{marginal}\\ \text{emission}\\ \text{damage}}} \underbrace{\frac{\mathbb{V}\mathrm{ar}(s)}{2(\mathbb{E}\,s)^2}}_{\substack{\text{level of}\\ \text{uncertainty}\\ \text{(normalized)}}} X_t(\cdot)\ . \tag{6}$$

The SCC today ($t = 0$) equals the deterministic contribution (equation 3 evaluated at expected climate sensitivity) plus the uncertainty premium ΔSCC_0. Equation (6) shows that the structure of ΔSCC_0 mirrors that of the deterministic SCC: both involve summing future impacts and translating them into present-day consumption equivalents, tracing the marginal effect of a ton of CO_2 released today through future carbon stocks, temperatures, output, and marginal utility.

The last two terms in equation (6) capture the role of uncertainty. First, the cost of uncertainty rises with the variance of climate sensitivity, normalized by its expectation.[10] Holding expected sensitivity fixed, the contribution is proportional to this variance, reflecting the small-risk approximation that neglects higher-order terms. Second, the factor $X_t(\cdot)$ specifies the structural channels through which uncertainty affects climate policy per unit of normalized variance. We now unpack these channels as spelled out in equation (5).

Risk aversion. The first message of equation (5) is that risk aversion matters directly (2 RRA). This is unusual in precautionary savings models and overlooked in standard asset-pricing approaches focusing only on correlations between marginal utility (or the stochastic discount factor) and marginal damages. The contribution arises from the covariance of marginal utility with the structural uncertainty in climate sensitivity. A high realization of climate sensitivity aggravates both channels simultaneously, raising the level of warming that lowers consumption as well as the marginal warming from an additional ton of emissions. The elasticity $\epsilon_{c,s}$ translates realizations of climate sensitivity into corresponding consumption effects. Our reasoning uses Assumption 1, under which climate sensitivity directly scales warming in degrees Celsius; relaxing this assumption slightly modifies the mechanism (see Corollary 1 in Appendix B).

(Welfare-) Prudence. The second contribution in equation (5) is the standard precautionary savings term, here applied to investment in a better climate. It operates through the uncertain effect of climate sensitivity on baseline warming. The elasticity $\epsilon_{c,s}$ converts this uncertainty into variation in consumption. By abating more CO_2, a prudent decision maker raises expected future consumption and thereby reduces effective risk aversion and the welfare costs of climate uncertainty.

Damage convexity. On the production side, damage convexity plays the same role that risk aversion does on the preference side. Whereas risk aversion reflects the curvature of utility, the moment Dam_2 captures the curvature of the production function with respect to temperature. The direct $2\,\mathrm{Dam}_2$ term in equation (5) parallels the 2 RRA contribution: both arise from covariances that are overlooked in standard asset-pricing perspectives focusing only on the correlation between marginal damages and marginal utility. Here, the relevant covariance is between marginal damages and the additional marginal warming generated by high climate sensitivity.

Intuitively, damage convexity reflects that early degrees of warming are considered less harmful, while additional warming becomes increasingly destructive. A high realization of climate sensitivity simultaneously raises the prevailing temperature level and the marginal

[10] The normalization arises because $X_t(\cdot)$ is expressed in terms of *relative* risk aversion, prudence, and damage measures as well as elasticities.

warming from a new ton of emissions, so that extra warming occurs precisely where damages are most severe. Conversely, a low realization reduces marginal warming but does so at a low baseline temperature, where damages are more modest.

Economy-prudence. The "economy-prudence" term proportional to Dam_3 is the production-side analogue of welfare prudence. Just as precautionary savings arise because higher wealth reduces the marginal pain of shocks, economy-prudence arises when higher temperatures amplify the convexity of damages. In this case, rising convexity with temperature amplifies the effect of uncertainty, creating an additional incentive for the policy maker to limit warming. Because damages are a bad, an increasing convexity of damages plays the same role as a decreasing convexity of welfare.

Whether damage convexity rises with temperature is ultimately an empirical question. Physical damages are bounded by total production (or at least the capital stock), so convexity should eventually decline and even turn concave. Indeed, even under optimal climate policy and relatively low temperature levels, we find that the sign of Dam_3 varies across versions of the DICE model.

Welfare-economy interaction. Finally, the interaction of risk aversion with damage convexity ($3\,\mathrm{RRA}\,\mathrm{Dam}_2$) captures the standard asset-pricing channel: the covariance between marginal utility and marginal damages (Jensen and Traeger, 2014; Dietz et al., 2018; Lemoine, 2021b). Under climate sensitivity uncertainty, both relative risk aversion and damage convexity are positive, so the covariance is positive. Intuitively, convex damages create expected production losses, and a risk-averse policy maker weighs these losses more heavily at low consumption (high temperature) levels, strengthening the case for mitigation. By contrast, under pure growth uncertainty the covariance can turn negative, since higher growth raises both consumption and damages, making losses occur when the economy is better off (Jensen and Traeger, 2014). Dietz et al. (2018), Lemoine (2021b), and van den Bremer and van der Ploeg (2021) analyzy how growth and climate uncertainties interact in this channel.

In the case of **risk neutrality**, the risk premium reduces to the last two terms in equation (5), those proportional to damage convexity and its change with temperature.

3.3 Temperature Stochasticity

Temperatures fluctuate naturally from year to year and over longer horizons. The empirical literature and some simulation models use such fluctuations as proxies for climate change when estimating economic damages. This section explains why the economic response to stochastic temperature variation differs fundamentally from the response to uncertain climate change.

Temperature stochasticity shifts the level of future warming but *does not affect the marginal warming from an additional ton of CO_2 released today.* As a result, the analysis below is closer to the classic precautionary savings and asset pricing case. The elasticity $\epsilon_{c,T}$ translates temperature changes into their consumption impact. The next proposition identifies the *additional* contribution to the SCC from stochastic temperature dynamics, holding climate sensitivity fixed.

Proposition 3. *Under Assumptions 1 and 2, temperature stochasticity increases the social cost of carbon contribution from a given period if and only if*

$$Z_t(\cdot) \equiv \text{RRA}\,\epsilon_{c,T}\big[\underbrace{\text{Prud}\,\epsilon_{c,T}}_{\substack{\text{welfare}\\ \text{prudence}}} + \underbrace{3\,\text{Dam}_2}_{\substack{\text{welfare}\\ \text{economy}\\ \text{interaction}}}\big] + \underbrace{\text{Dam}_2\,\text{Dam}_3}_{\substack{\text{economy}\\ \text{prudence}}} \tag{7}$$

is greater than zero. Arguments and their period-dependence are suppressed.
Moreover, under a small risk approximation, the temperature stochasticity contribution is

$$\hat{\Delta}SCC_0 \approx \sum_{t=1}^{\infty}\sum_{\tau=1}^{t} \underbrace{\beta^t \frac{u_t'(c_t)}{u_0'(c_0)}}_{\substack{\text{consumption}\\ \text{discount}\\ \text{factor}}} \underbrace{\frac{\partial F_t}{\partial T_t}\frac{\partial T_t}{\partial M_\tau}\frac{\partial M_\tau}{\partial E_0}}_{\substack{\text{marginal}\\ \text{emission}\\ \text{damage}}} \underbrace{\frac{1-\gamma^{2t}}{1-\gamma^2}}_{\substack{\text{intertemporal}\\ \text{temperature}\\ \text{correlation}}} \underbrace{\frac{\mathbb{V}\text{ar}(\epsilon)}{2(\mathbb{E}\,T_t)^2}}_{\substack{\text{level of}\\ \text{uncertainty}\\ \text{(normalized)}}} Z_t(\cdot). \tag{8}$$

Equation (7) shows that the welfare-prudence and economy-prudence channels both continue to generate a precautionary response. In addition, the interaction of risk aversion and damage convexity contributes through the classic covariance between marginal utility and marginal damages.

Most importantly, comparing equation (7) to equation (5) shows that the novel covariance terms we identified earlier—the direct contributions of risk aversion and of damage convexity—drop out under temperature stochasticity. These terms are the dominant drivers of the total uncertainty premium, yet they are entirely absent when risk is modeled solely through stochastic temperature paths. On the upside, the stochasticity premium can, at least in principle, be identified from market-based observations of risk premia, offering a clear empirical handle. But overall, this highlights a key limitation of the stochastic approach, which remains common in the literature: while it captures prudence-type channels and the classical covariance between marginal utility and marginal damages, it misses the structurally most important contributions to the climate risk premium.

Equation (8) shows that the quantification of the stochasticity premium contains a novel factor arising from the autoregressive nature of global temperatures. Although the shocks ϵ_t are independent, the temperature realizations across periods are correlated (equation (4)).

This intertemporal correlation introduces the multiplier $\frac{1-\gamma^{2t}}{1-\gamma^2}$, which amplifies the effective variance of shocks in each period. Absent this factor, the variance from temperature stochasticity would be very small relative to the variance associated with climate sensitivity.

Propositions 2 and 3 provide a useful perspective for interpreting much of the literature on climate policy under uncertainty. It clarifies how different modeling assumptions determine whether uncertainty affects policy levels or only welfare comparisons. For example, the first generation of analytic IAMs consisted of risk neutral linear-quadratic models inspired by Weitzman's (1974)'s seminal prices-versus-quantities paper. These models adapted his framework for stock pollutants to the climate context, providing tractable tools to compare alternative policy instruments such as carbon taxes and emissions trading schemes (Hoel and Karp, 2002; Karp and Zhang, 2006; Newell and Pizer, 2008; Zhang, 2012; Weitzman, 2018; Karp and Traeger, 2018; Pizer and Prest, 2019; Heijmans and Gerlagh, 2019). Their appeal lay in analytical clarity: by assuming quadratic damages and risk neutrality, they delivered closed-form results for efficiency comparisons between instruments. Propositions 2 and 3 show that their policy prescriptions hinge on a knife-edge treatment of uncertainty: risk neutrality removes the welfare channels of risk aversion and prudence; quadratic damages rule out the economy-prudence effect; and representing risk as stochastic climate evolution rather than structural uncertainty eliminates the direct damage-convexity term (cf. Proposition 3 versus Proposition 2). Under these combined assumptions, uncertainty does not change optimal policy levels. It only matters through welfare losses across policy instruments, which remain linked to damage convexity.

More recently, a second generation of analytic IAMs has built on Golosov et al.'s (2014)'s breakthrough in modeling richer economy-climate interactions. Golosov et al. (2014) introduce uncertainty, but because their framework is linear in all relevant state variables, uncertainty has no effect on either welfare or policy. Much of the subsequent literature in this tradition abstracts from temperature altogether. By contrast, Traeger (2018) and van den Bremer and van der Ploeg (2021) extend the framework by incorporating explicit temperature dynamics and general risk preferences. These extensions embed structural uncertainty into a stochastic setting and reveal hybrid contributions relative to our equations (2) and (3). For instance, van den Bremer and van der Ploeg (2021) solve their model using perturbation methods and capture a stochastic analogue of our damage-convexity contribution. Traeger (2018) further discusses how Bayesian learning models of epistemic uncertainty can be translated into simplified stochastic formulations.

More recently, a second generation of analytic IAMs has built on Golosov et al.'s (2014)'s breakthrough in modeling richer economy-climate interactions. Golosov et al. (2014) introduce uncertainty, but because their framework is linear in all relevant state variables,

uncertainty has no effect on either welfare or policy. Much of the subsequent literature in this tradition abstracts from temperature altogether. By contrast, Traeger (2021) and van den Bremer and van der Ploeg (2021) extend the framework by incorporating explicit temperature dynamics and general risk preferences. These extensions embed structural uncertainty into a stochastic setting and reveal hybrid contributions relative to our equations (2) and (3). For instance, Van den Bremer and Van der Ploeg (2018) solve their model using perturbation methods and capture a stochastic analogue of our damage-convexity contribution. Traeger (2018) further discusses how Bayesian learning models of epistemic uncertainty can be translated into simplified stochastic formulations.

Both of these papers also introduce skewness in the temperature shocks to mimic the right-skew of climate sensitivity. This generates an additional "skew amplification" of the risk premium. Corollary 2 in Appendix C generalizes Proposition 3 to cover this case by modifying the temperature evolution to the following form.

Assumption 1'. *The temperature evolution is*

$$T_{t+1} = \gamma T_t + s\,\Gamma_1(M_{t+1}) + \Gamma_2(t) + \Gamma_3(\epsilon_t), \tag{9}$$

where $\gamma \in (0,1)$*,* s *may be uncertain,* Γ_1*,* Γ_2*, and* Γ_3 *are real functions, and* $(\epsilon_t)_{t\in\mathbb{N}}$ *is an i.i.d. process of mean-zero temperature shocks.*

The new element is the transformation $\Gamma_3(\cdot)$, which allows the shocks ϵ_t to generate skewed temperature dynamics. As shown in Corollary 2, the stochasticity premium in equation (8) then acquires a normalized measure of skewness, $\frac{\Gamma_3''(0)}{\Gamma_3'(0)}T_t$, closely analogous to Arrow-Pratt risk aversion and our measure of damage convexity. This term contributes an *additional* component to the stochastic risk premium in equation (7), interacting with both risk aversion and damage convexity:

$$\frac{T_t\Gamma_3''(0)}{\Gamma_3'(0)}\left[\mathrm{RRA}\,\epsilon_{c,T} + \mathrm{Dam}_2\right].$$

Corollary 1 in Appendix B establishes a corresponding result when temperature dynamics are nonlinear in climate sensitivity.[11]

[11] It is relatively straightforward to introduce skewness into a stochastic process, but much harder to introduce skewness into structural uncertainty about climate sensitivity while maintaining consistency with Bayesian learning. Kelly and Tan (2015) achieve this in a fat-tail setting with conjugate priors and approximate Bayesian updating.

4 Quantification

This section quantifies the climate risk premium using a DICE-based integrated assessment model, solved via either a Monte Carlo approach (no policy response) or a recursive dynamic programming approach (responsive policy). We use these numeric results to test the accuracy of our risk premium formula. The formula allows us to quantify each uncertainty channel separately, whereas the numeric model only provides the overall premium.

Section 4.1 outlines the numerical implementation and model calibration. Section 4.2 presents the risk premium for non-responsive policies. Section 4.3 incorporates policy responsiveness. Section 4.4 quantifies the risk premium across various scenarios, including a sensitivity analysis. Finally, Section 4.5 quantifies the different risk channels for selected scenarios.

4.1 Numerical Implementation and Calibration

The core components of our model - production, growth, emissions benefits, and climate damages - are based on the DICE model (Nordhaus and Sztorc, 2013). [12] DICE employs a quadratic damage function $D(T) = aT^2$, which specifies the percentage loss of potential world output due to a temperature increase T above preindustrial levels. Temperature follows equation (4).

Our CO2 dynamics are slightly simplified compared to DICE, easing computational demands and addressing the exaggerated sluggishness in DICE's carbon cycle (Traeger, 2014; Van der Ploeg et al., 2020). Given our focus on uncertainty, we use a constant relative risk aversion (RRA) of RRA = 2, differing from Nordhaus' RRA = 1.45 for a deterministic world. The estimate for temperature stochasticity is based on GISTEMP data (GISTEMP Team, 2024), confirming the quantification used by Kelly and Kolstad (1999).

We approximate the equilibrium climate sensitivity uncertainty by a normal distribution with a mean of $3°C$ and a variance of $3°C^2$. We use a normal distribution because our stochastic dynamic programming model with Bayesian updating requires a conjugate prior. The mean corresponds to the best guess of the latest, sixth IPCC assessment report (Arias et al., 2021). The variance is higher than AR6's to compensate for the normal approximation lacking higher-order moments.[13] The model's climate dynamics generate transient climate

[12] We prefer the 2013 vintage of DICE over the 2016 version for two reasons: (i) Van der Ploeg et al. (2020) show that its carbon cycle performs substantially better compared to scientific models used by the Stocker et al. (2013), and (ii) we find the sharp initial decline in the rate of technological progress in DICE 2016 descriptively unconvincing.

[13] It matches the estimate in AR6 derived from climate models alone, which is higher than that derived using all methods.

sensitivity (TCS) intervals of 1 to 2.6 degrees Celsius (66%). The interval in AR6 is narrower, consistent with climate model estimates, although AR6 reports a narrower range.

Our robustness scenarios vary the time preference, risk aversion, and damage specifications. Our robustness scenarios vary time preference, risk aversion, and damage specifications. We also include a scenario adjusting DICE's economic parameters to match the 2019 Penn World Tables (Feenstra et al., 2015), updating global output, capital-output ratio, rate of return, and investment rate.

The premium under a non-responsive policy is generated using the Monte Carlo method. The decision maker implements the optimal policy for a climate sensitivity of 3. We draw 1000 realizations of climate sensitivity from a normal distribution with a mean and variance of 3.[14] The model is then simulated for each realization, and the 1000 time paths are averaged to calculate the optimal non-responsive policy under uncertainty.

For the responsive policy scenario, we employ recursive stochastic dynamic programming techniques. Appendix D details our numerical implementation, which solves the stochastic infinite horizon fixed-point problem using a Chebyshev approximation of the value function, following the approach pioneered by Kelly and Kolstad (1999). The social planner observes temperature realizations, which inform Bayesian updating of the climate sensitivity distribution. The optimal policy anticipates these learning dynamics and the resulting future policy adjustments.

4.2 Monte Carlo

In the Monte Carlo context, we derive and interpret our risk premium formula under non-responsive policies (Assumption 2). The formula sums the contributions from climate sensitivity uncertainty (equation (6)) and temperature stochasticity (equation 8). Aside from the calibrated variances, our formula relies on the optimal deterministic time path from the numeric integrated assessment model, which assumes a climate sensitivity of 3°C. Our formula yields an SCC risk premium of 19.1 USD/tC, with temperature stochasticity contributing only 25 cents. The increase over the deterministic SCC is 31%.

We depict the decomposition of the risk premium from climate sensitivity uncertainty, following equation (5), in Figure 2. The left graph shows the contributions to the 2020 SCC from each future period. The primary contributions to the risk premium arise from the latter half of this century, increasing as uncertainty unfolds but eventually decreasing due to discounting. The right graph illustrates the cumulative contributions. In the reference

[14]To avoid scientifically implausible negative values, any such realizations are set to zero. Negative sensitivity is an artifact of the normal distribution, necessary for our stochastic dynamic programming model. DICE's quadratic damage function would imply 'cooling' damages for negative values.

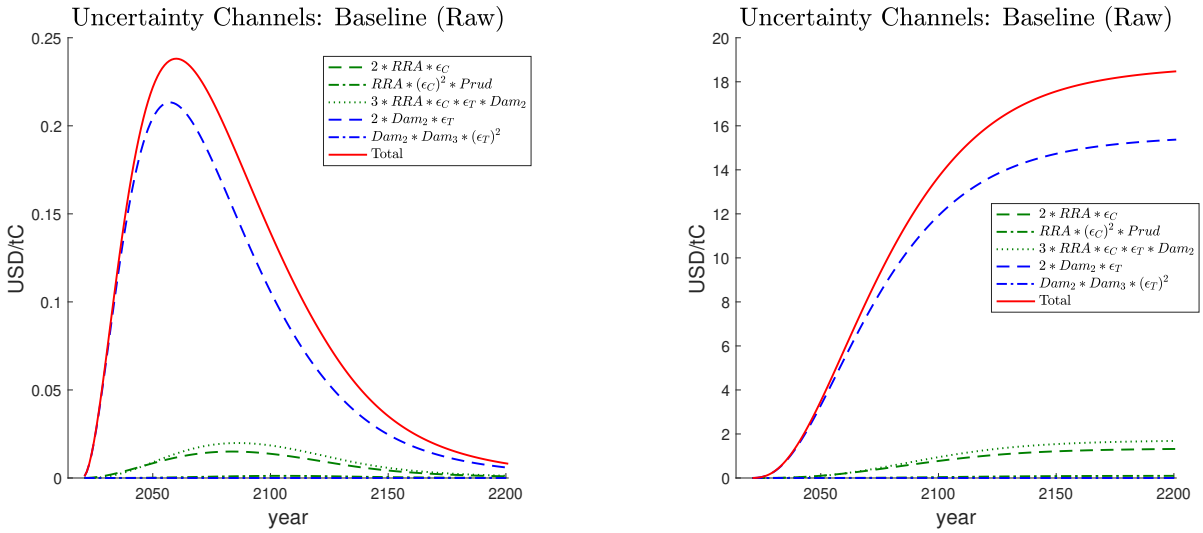


Figure 2: Contributions of climate sensitivity uncertainty to the SCC for our base formula and our reference scenario of DICE 2013. The left panel shows the contributions to the SCC in 2020 from a given year in the future, whereas the right panel shows the cumulative contributions corresponding to the depicted time horizon.

scenario, discounting leads the policymaker to disregard risks beyond a few centuries into the future.

The damage convexity channel contributes the most to the risk premium (82%), reflecting the structural nature of the uncertainty, which is absent in models that only represent uncertainty via a stochastic shock process. Structural uncertainty affects both the baseline warming from existing atmospheric CO_2 and the marginal warming from an additional ton of carbon released today. Risk aversion contributes 9% through its interaction with damage convexity and 7% through the direct risk aversion effect. Contributions from prudence and economy-prudence effects are negligible.

4.3 Policy Response and Stochastic Dynamic Programming

The numeric stochastic dynamic programming model delivers a policy risk premium of 15.8 USD/tC, a 26% increase over the deterministic SCC. Compared to the base formula, anticipating future policy responses to climate sensitivity updates reduces the risk premium on the SCC by 17%. If future climate sensitivity turns out to be higher than expected, future policy makers will strengthen climate policy, reducing the costs of a high realization of climate sensitivity. The current policy maker anticipates that response. If climate sensitivity turns out low, the current policy maker anticipates that future policy makers will relax the climate policy, partly mitigating the costs of too strict policy. Being able to respond to deviations from the expected climate sensitivity will reduce future costs and, thus, make climate change

less costly today.

How can our formula account for future policy responses? The lack of anticipation manifests in our base formula in the temperature elasticity to climate sensitivity. In the base formula we calculate $\epsilon_{T,s}$ based on a recursive evaluation of the temperature equation of motion (4). This evaluation assumes that higher climate sensitivity translates straight into higher temperatures, as it solely relies on the climate system. By contrast, if future policy makers react to a high climate sensitivity with stricter policy, this moderates the effect on temperature. With responsive policies, the temperature elasticity to climate sensitivity can no longer be deduced from climate system characteristics alone.

An analytic approach to calculating the temperature elasticity to climate sensitivity fully incorporating all future policy responses is infeasible; it would entail solving the complete stochastic dynamic programming problem analytically. But we can use the deterministic numeric model to create a proxy for policy responsiveness. For this purpose, we calculate $\epsilon_{T,s}$ by evaluating the deterministic model along the optimal path for expected climate sensitivity as well as along a small perturbation where climate sensitivity is a tenth of a degree higher. Figure 3 plots the resulting temperature elasticity (dashed blue line) and contrasts it with the analytically calculated elasticity neglecting the policy response (solid black line). We also depict the elasticity resulting from our numeric solution of the stochastic model (red dash-dotted line).[15] The graph confirms that future policy responses substantially decrease the medium and long-term temperature response to a change in climate sensitivity. The graph also shows that the elasticity hardly differs between the stochastic and the deterministic model. Employing the deterministic (blue-dashed) elasticity $\epsilon_{T,s}$ in our analytic formula delivers a policy risk premium of 16 USD/tC, which deviates by less than 2% from the premium obtained from the stochastic dynamic programming model with Bayesian learning (the deviation is merely 0.8% with respect to the overall SCC). [16] We can focus our attention

[15]Here, we calculate $\epsilon_{T,s} = \frac{dT}{ds} \frac{s}{T}$ by changing both expected and true climate sensitivity by a tenth of a degree and assuming that realized temperature shocks are zero. True climate sensitivity and actual shocks are unkown to the decision maker, see Appendix E.2 for details. In the case of uncertainty, we evaluate the elasticity along the path of expected realizations.

[16]For our reference scenario, evaluating the elasticity merely along the expected path works well. As a higher order effect, the response elasticity itself will vary depending on the realized climate sensitivity. For scenarios delivering high risk premia, we can slightly improve the approximation of $\epsilon_{T,s}$ by using the average response across two different (deterministic) scenarios, one with a high and one with a low realization of climate sensitivity, see section 4.4. Table 1 shows that these scenarios are the ones varying the damage function and changing the preference structure more substantially (Epstein-Zin preferences). In these scenarios, marked by an asterisk, we use the two Gaussian quadrature nodes and weights of the climate sensitivity distribution. Gaussian quadrature picks discrete nodes and weights to approximate the moments of a continuous distribution. The expectations calculated from two Gaussian nodes already match the first 4 moments of a distribution exactly. We only use them to calculate the elasticity $\epsilon_{T,s}$ from two different realizations of climate sensitivity using the deterministic model.

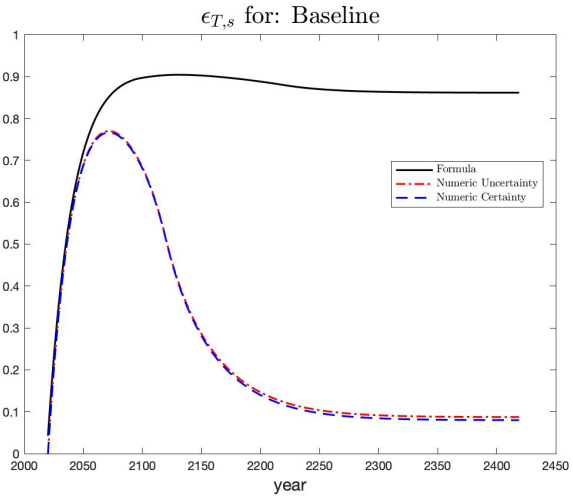


Figure 3: Elasticity of the temperature response to changes in climate sensitivity.
(i) solid line: without optimal policy response (Assumption 2);
(ii) blue-dashed line: including optimal policy response in the deterministic model;
(iii) red-dash-dotted line: including optimal response in the stochastic model.

on $\epsilon_{T,s}$ because the dominant contributions to the risk premium rely on $\epsilon_{T,s}$; contributions relying on the consumption elasticity $\epsilon_{c,s}$ are minimal even without anticipating future policy response.[17]

We believe approximating policy responsiveness in this way can have a positive impact on applied modeling. Currently, the only IAM for which the literature offers fully numeric stochastic dynamic programming solutions governing climate sensitivity is the low-complexity DICE model. Most models are too complex for dynamic programming solutions due to computational requirements. Our approach suggests a way forward to incorporate the anticipation of future policy response into these complex IAMs.

4.4 Quantifying the Risk Premium

Baseline. Table 1 presents the overall climate risk premium for our reference calibration and for variations of key parameters and functional forms. We find that our responsive formula performs well for most scenarios. The climate risk premium of our reference scenario is 16 USD/tC, which is about one quarter of its deterministic value. An annual rate of pure time preference of $\rho = 1.5\%$ might be a somewhat high pick. Reducing it to the median estimate of the recent expert survey by Drupp et al. (2018), $\rho = 0.5\%$, increases the risk premium to 22 USD/tC, but slightly decreases the risk premium as a fraction of the deterministic SCC. Our major update of DICE to match the 2019 Penn World Table reduces the absolute

[17]The consumption elasticity to climate sensitivity is less straight-forward to approximate numerically because consumption itself is a control variable. Given the corresponding channels are minute already without the anticipation of future policy updates, we opt for simplicity and adjust only the quantitatively relevant elasticity.

Risk Premium (USD/tC) / Scenario	Analytic Formula Raw	Adjusted	Error	Full Model	Stochastic Fraction	Fraction of Cert
RRA=2, $\rho = 1.5$, DICE13	19.1	16.0	0.8%	15.8	1.6%	26%
RRA = 1.45	29.8	21.3	0.7%	21.4	1.4%	21%
PRTP $\rho = 0.5$	34.3	22.6	2.1%	23.0	2.1%	21%
Update PWT 2019 (RRA=3)	13.7	12.7	5.2%	13.4	2.2%	28%
DICE 2007 Damages	16.0	13.4*	2.8%	13.0	0.4%	21%
Howard & Sterner Damages	100	74.4*	8.1%	80.9	6.7%	35%
Cubic Damages	70.8	46.8*	3.8%	48.6	8.7%	76%
Cubic Damages, $\rho = 0.5$	122	71.2*	1.5%	70.1	10.6%	63%
Epstein-Zin: $\eta = 2, RRA = 6$	26.4	21.8*	10.1%	19.8	3.7%	32%
Epstein-Zin: $\eta = \frac{2}{3}, RRA = 6$	87.7	57.5*	12.1%	51.3	5.7%	20%

Table 1: Total risk premium for formula without policy response, with policy response adjustment, and from the recursive dynamic programming model. The error column specifies the deviation of the adjusted formula from the numeric solution. The column on the right presents the relative risk premium, i.e., the risk premium relative to the deterministic SCC based on the full stochastic model. The asterix * denotes results using two Gaussian-quadrature nodes to calculate the temperature elasticity $\epsilon_{T,s}$ instead of merely using the expected climate sensitivity.

risk premium and SCC as a result of matching a higher rate of return, implying higher discounting. In relative terms, the risk premium increases, which might not be surprising as we increase risk aversion to $RRA = 3$ to match returns. Without changing the damage convexity or the fundamental preference structure, we find a risk premium of approximately one quarter of the deterministic SCC.

Damage Specification. The most substantial changes of the risk premium derive from changes in the damage function, in line with our initial quantification of the different risk channels in Section 4.2. DICE 2013's damage specification $D(T) = aT^2$ assumes that the world loses $a = 0.266\%$ of potential output at a 1°C warming (today) and nine times as much at a 3° warming. Howard and Sterner's (2017)'s meta analysis of damage estimates suggests five times those damages, somewhat in agreement with a recent survey among climate scientists and economists by Pindyck (2020). Howard and Sterner's (2017)'s preferred estimate is a loss of $a = 1.15\%$ of potential output at a 1°C warming increasing quadratically in temperature just as in the DICE model. Under such an increase of the damage magnitude, the risk premium increases to approximately one third of the deterministic SCC.

Several authors have suggested that damages might be much more convex than assumed in these quadratic specifications (e.g. Weitzman 2012, Millner 2013). Changing damages to a cubic formulation, $D^{cub}(T) = a^{cub}T^3$, we adjust the linear coefficient a^{cub} such that both functions imply the same damages at a 2.5°C warming, the main calibration point of

Nordhaus' DICE model. Damages are initially smaller, but eventually larger, at least with some probability, compared to the baseline. In this scenario, the risk premium increases to 76% for our base specification and to 63% in the case of the reduced time preference.

Risk Preferences. Finally, we consider the decision maker's risk aversion. Merely reducing it $RRA = 2$ to DICE's original value of 1.45 reduces the relative premium to about 20%, but increases the absolute risk premium. In the standard economic model, the parameter RRA serves two purposes. First, it characterizes the decision maker's risk aversion. Second, it characterizes the decision maker's desire to smooth consumption over time. In a growing world economy, a preference for consumption smoothing translates into growth discounting. The weight on the future increases when we lower the desire to redistribute from the rich future to the poor present. As a result, the concern for climate change and the risk premium increase despite the reduction in risk aversion. We observe the same finding, but in the opposite direction, in the scenario matching Penn World Table data, where the risk aversion increases; the relative premium increases but the absolute premium and SCC fall.

Epstein and Zin's (1989) preferences disentangle the concerns for consumption smoothing and risk aversion. A branch of the asset pricing literature shows that such preferences improve the explanation of observed discount rates and risk premia in asset markets.[18] Table 1's Epstein-Zin specifications increase risk aversion without simultaneously reducing the desire to smooth consumption over time. Based on the literature cited above, we increase risk aversion to $RRA = 6$. In this specification, the increase in risk aversion also increases the relative risk premium, which is now over 30%. However, the cited literature on long-run risk also suggests that Epstein-Zin preferences should use a higher elasticity of intertemporal substitution ($\eta = \frac{2}{3}$), which reduces the risk premium back to 20%. The premium is large in absolute terms, but the higher intertemporal elasticity reduces the growth-based discounting and substantially increases the deterministic SCC. The large positive risk premium contrasts sharply with the impact of growth uncertainty on optimal climate policy studied in Jensen and Traeger (2014) and Cai and Lontzek (2019), who find a negative risk premium under the same increase in the elasticity of intertemporal substitution under Epstein-Zin preferences. The difference arises because the structural climate sensitivity parameter impacts both marginal warming and warming levels (see section 4.2). Our analytic formula performs worse for Epstein-Zin preferences because it does not account for the recursive evaluation structure of such preferences.[19]

[18] Vissing-Jorgensen and Attanasio (2003); Bansal and Yaron (2004); Bansal et al. (2010); Chen et al. (2013); Bansal et al. (2012, 2014); Collin-Dufresne et al. (2016); Nakamura et al. (2017).

[19] Our formula can disentangle risk aversion entering the risk premia $X_t(\cdot)$ in equation (5) and $Z_t(\cdot)$ in equation (7) from consumption smoothing, which enters the other terms of equations (6) and (8). We

Policy Response. If policies do not respond to updated knowledge about climate sensitivity risk premia are higher for all scenarios. The difference between Table 1's "Base" and "Responsive" columns characterizes the cost reductions achieved byn in the future responding optimally to new information about climate uncertainty. Per ton of carbon emitted, the savings are by far the largest in the case of cubic damages under a low rate of pure time preference (50 USD/tC). In the baseline, the difference is merely a few dollars, and in the case of our Penn World Table match the difference is negligible. Also in relative terms the impact of an optimal policy response varies widely. The impact is highest for low discounting and high or more convex damages, both of which trigger more responsiveness to changes in long-run warming.

Expected Climate Sensitivity and Variance of the Prior. Figure 4 shows how the reference scenario's risk premium (black lines) varies in response to the expected climate sensitivity and the Bayesian prior's variance for the numeric stochastic dynamic programming implementation. It also graphs these dependencies for the scenario with cubic damages (red dash-dotted lines). Our formula's risk premium increases linearly in the variance as a result of the small risk approximation. Figure 4, on the left, shows that for the numeric model the SCC's risk premium increases linearly also in the Bayesian prior's variance. This linearity still holds for large risk, despite the model's non-linearities. Our normal distribution favors this result, but it enables us to implement a consistent Bayesian learning model as part of our stochastic dynamic programming implementation. Under DICE's quadratic damages, the risk premium increases approximately 5 USD/tC per unit variance. Under cubic damages, it increases three times as steeply.

The right graph of Figure 4 shows that a higher expected climate sensitivity does not necessarily increase the SCC's risk premium. For quadratic damages, the premium decreases slightly as the expected climate sensitivity increases. To reconcile this numeric result with the discussion in Section 3.2, note that we in section 4.2 derived that in the reference scenario the dominant contribution to the risk premium is proportional to damage convexity. With quadratic damages, the relative convexity measure (Dam_2) is independent of the warming level.[20] Under a cubic damage specification, the convexity increases in the expected warming

obtain the analytic formula only by omitting the recursive structure of Epstein-Zin preferences. A recursive evaluation tends to imply substantial premia for an early resolution of uncertainty. Epstein et al. (2014) show that these might be unreasonably high. A nice feature of the formula is that it does not involve such premia.

[20]That the line is almost exactly constant also relies on the normalization factors and the response elasticities. They vary, but in ways that mostly offset each other. We note that the *absolute level of abatement expenditure* as a result of uncertainty increases with a higher expected climate sensitivity. The deterministic SCC is already higher under a higher expected climate sensitivity, and adding (less) additional units of abatement at the higher abatement cost comes at a higher additional expenditure.

level. The next section discusses the different risk channels for selected scenarios in more detail.

Related Quantifications. Kelly and Tan (2015) find a risk premium of 24% using RRA = 1.5 and a much higher rate of pure time preference of 5% in combination with fat tails. Rudik and Lemoine (2017) find a risk premium of 1% using a time preference of 1.5% (and of 10% in the absence of learning). Both papers rely on DICE-based recursive stochastic dynamic programming models.[21] Keller et al. (2004) and Lemoine & Traeger (2014, 2016) analyze the consequences of an unknown threshold at which a discrete regime shift increases climate sensitivity (tipping point). Lemoine and Traeger (2014) show that such tipping risk increases the SCC by approximately 20%. This effect, however, combines risk with a higher *expected* climate sensitivity and such "tipping premia" operate mostly through the increase of the expected climate sensitivity rather than through the uncertainty modeled here (Lemoine and Traeger, 2016a; Taconet et al., 2021). Like the present paper, the preceding studies focus on epistemological uncertainty governing the climate's sensitivity to greenhouse gas emissions.

Hambel et al.'s (2021) continuous time AK model identifies a risk premium of 4% under DICE damages that increases to 12% under a damage function with higher and more increasing convexity, using Epstein-Zin preferences and a pure time preference of 1.5%. Van den Bremer and Van der Ploeg (2018) find an approximate risk premium of 22% in their base case, which is based on damage convexity and right skew, and Traeger (2018) finds a risk premium of 35% in the base case, falling to 20% for a reduction in risk aversion making the result more comparable to our base case. These settings rely on a stochastic climate rather than epistemological uncertainty and skewness of the distributions plays a more crucial role (see Section 3.3).

4.5 Risk Channels: Sensitivity and Insights

Figure 5 quantifies the different uncertainty channels using the policy responsive formula evaluation for selected scenarios. For comparison purposes, the top left panel depicts the contributions for the reference scenario, which are similar to the non-responsive evaluation in Figure 2.[22]

Time Preference. The bottom left panel of Figure 5 displays the contributing channels under the reduction of the annual rate of pure time preference from $\rho = 1.5\%$ to $\rho =$

[21] Whereas Kelly and Tan (2015), like us, use a state-reduced version of DICE that some of DICE's exagaggerated sluggishness in the climate dynamics, Rudik and Lemoine (2017) use the full DICE model and a Smolyak grid approximation to deal with state space complexity.

[22] The only notable difference is that the (minor) direct risk aversion contribution and welfare-economy-interaction contribution are now indistinguishable in their contributions to the risk premium.

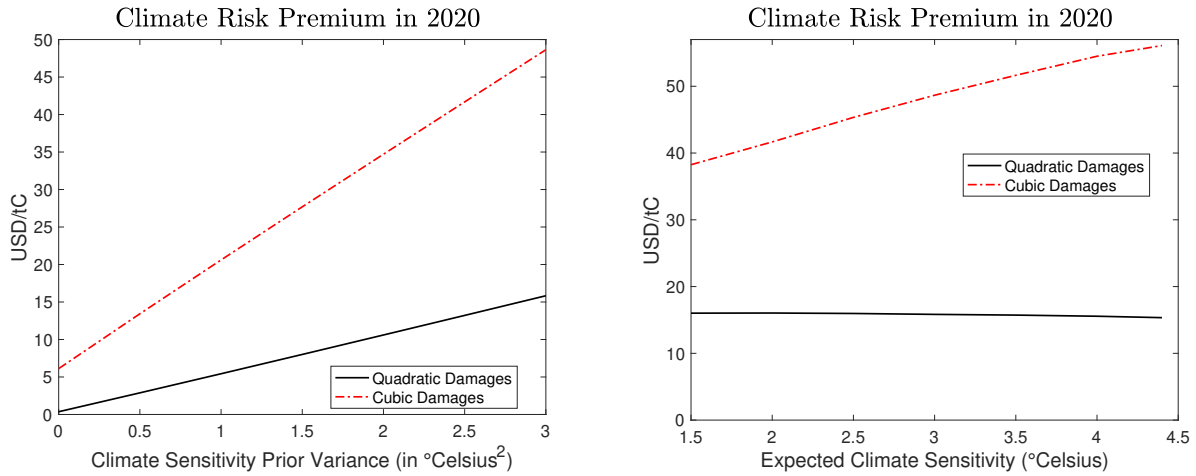


Figure 4: Dependence of the SCC's risk premium in 2020 on the prior's variance (left) and the expected climate sensitivity (right).

0.5%. We find that this reduction of the discount rate hardly changes the composition of the contributing factors. The only difference is a tiny increase of the direct risk aversion contribution relative to the interaction term. The cumulative contribution flattens out a little slower than in the baseline, but what happens after the year 2200 still hardly matters for the (SCC and its) risk premium.

Damage normalization. The right panels of Figure 5 change the shape of the damage function and alter the contributing channels notably. Nordhaus (2008) normalizes quadratic damages as $D(T) = \frac{aT^2}{1+aT^2}$, which avoids that damages can exceed total production. Nordhaus and Sztorc (2013) and Nordhaus (2017) still use this formulation in the model description, but the actual implementations of DICE 2013 and DICE 2016 now use the simpler formulation $D(T) = aT^2$, which has the advantage that a directly characterizes the fraction of world GDP lost at a 1°C warming. The top right panel changes the damage function "back" to its normalized version. This minor change under certainty implies a notable distinction for the risk premium. This formulation implies a falling damage convexity. As a result, the economy-prudence term contributes negatively; uncertainty's production loss caused by the damage convexity falls in temperature. The welfare prudence contribution is zero and the welfare-economy interaction contributes positively with a similar magnitude as the economy-prudence effect. As a result, we only find a positive risk premium because of the structural uncertainty.

Cubic damages. The lower right panel of Figure 5 analyzes the channels for our cubic damage specification ($D^{cub}(T) = a^{cub}T^3$). We note that the scale of this lower right panel is substantially higher. All non-zero contributions increase. In absolute terms, the damage convexity contribution increases the most. In relative terms, the *economy-prudence* channel

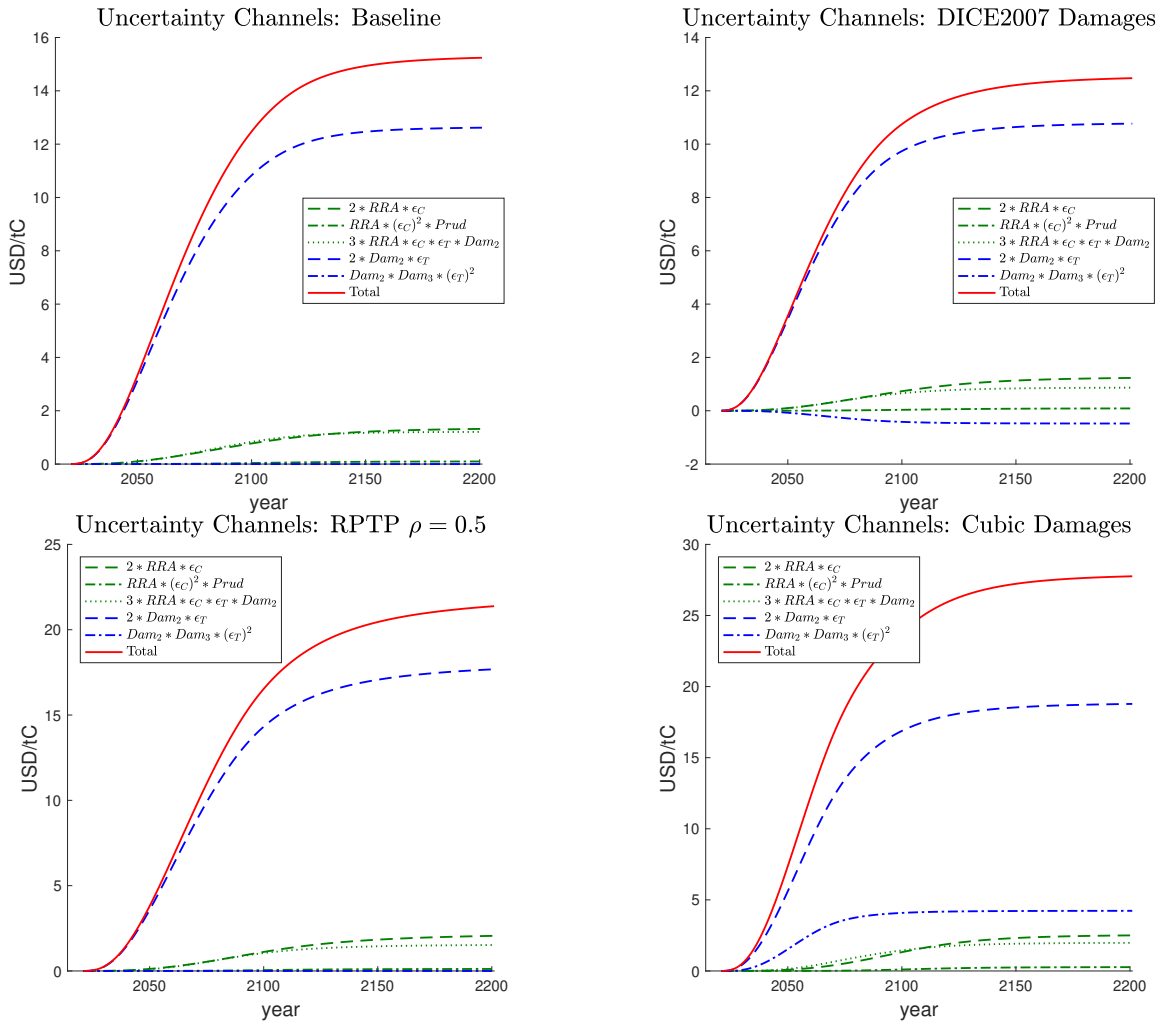


Figure 5: Uncertainty about climate sensitivity. Cumulative contribution to the SCC's risk premium in 2020 by channel for different scenarios.

catches up the most; it is zero or even negative in the other specifications and now contributes a substantial part of the risk premium. Also the welfare-economy interaction contribution grows together with the damage convexity.

Stochasticity premium. Figure 6 presents the corresponding graphs for the stochasticity premium. The two left panels represent the base scenario and the scenario with a reduced time preference. In both scenarios the *welfare-economy interaction* is the close-to-sole driver of the stochasticity premium. By Proposition 3, the channel contributing most under climate sensitivity uncertainty, the damage convexity channel, drops out for temperature stochasticity, and so does the direct risk aversion channel. The reduction in time preference pushes the time horizon of relevant contributions out further in the case of temperature stochasticity than in the case of climate sensitivity uncertainty; the curve is steeper and the relative contributions are more relevant around 2200. The top right panel normalizes damages to

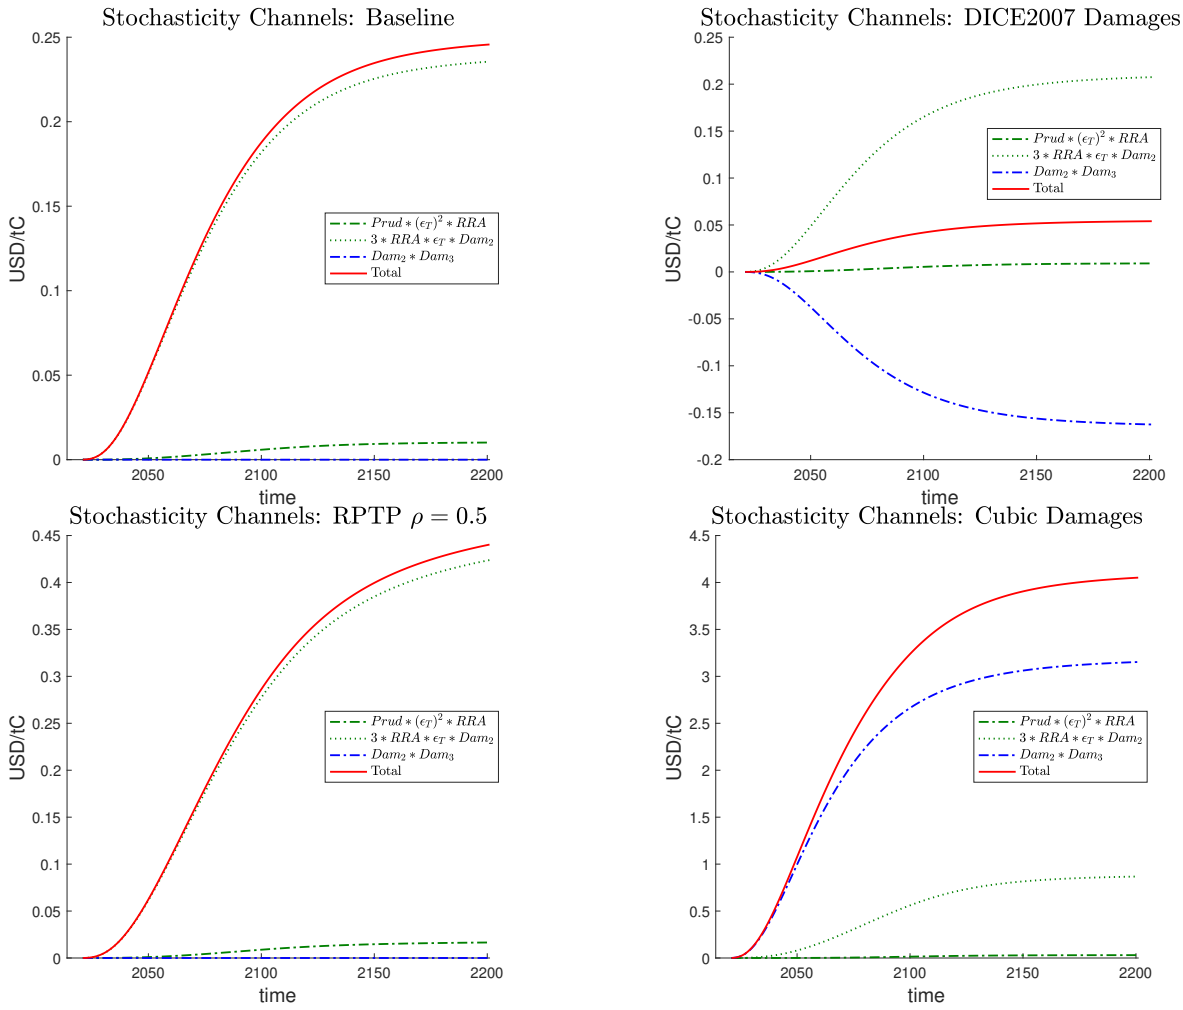


Figure 6: Stochastic temperature evolution. Cumulative contribution to the SCC's risk premium in 2020 by channel for different scenarios.

their DICE2007 form. As a result, the damage convexity falls in temperature and the now *negative economy-prudence term* almost cancels the positive contribution, leaving a negligible total stochasticity premium. It is interesting that a minor normalization under certainty, also performed in some later DICE-code-vintages without mentioning, can mostly eliminate the stochasticity premium. In contrast, the cubic damage specification results in a *positive economy-prudence contribution that dominates the stochasticity premium*, which now delivers a more notable contribution to the overall risk premium. These results illustrate that models merely relying on a stochastic temperature evolution, without modeling structural uncertainty, obtain notable risk premia only as a result of more convex damages or, similarly, skewed temperature distributions (see discussion of equation 9).

5 Conclusions

We analyze the climate policy risk premium in response to climate uncertainty. We focus on the epistemoligical uncertainty governing the temperature's response to greenhouse gas emissions (climate sensitivity). Using a simple but general integrated assessment model of climate change, we present an analytic formula approximating this risk premium. This formula enables an understanding of the different risk channels that drive the policy risk premium and allows us to quantify each of them. At the same time, we calculate the exact risk premium in a DICE-based numeric stochastic dynamic programming implementation.

We first derive the formula under the Monte Carlo approach most common in policy-oriented research. In the Monte Carlo setting, policies are fixed. We next show how a simple adjustment allows us to incorporate the policy maker's optimal response to uncertainty and its resolution into the formula. This adjustment calculates the approximate temperature response elasticity to climate sensitivity numerically using only the same deterministic model run we require to quantify the formula in the first place. With this modification, the formula matches the risk premium of the full stochastic dynamic programming model closely for a wide range of parameterizations.

Integrated assessment of climate change augments the familiar precautionary savings reasoning in two fundamental ways, giving rise to two additional sets of risk channels. Precautionary savings are driven by prudence: as we become richer, absolute risk aversion falls and we have an additional incentive to save under risk. In our setting, future temperatures additionally have a non-linear impact on the economy that interacts with the marginal appreciation for consumption. As a result, we identify the additional economy-prudence and welfare-economy-interaction channels. First, climate damages are convex and this convexity makes uncertainty costly. If this costly convexity increases in temperature, then policy has an additional incentive to keep temperatures down under uncertainty, establishing the economy-prudence effect. We show that the economy-prudence effect is either negative or zero in DICE, depending on the model version. It delivers a positive premium if damages are more convex (cubic rather than quadratic), in which case is grows substantially larger than the welfare prudence effect ("standard precautionary savings motive"). Second, the welfare-economy-interaction channel is always positive for a risk averse decision maker and convex climate damages. This interaction channel dominates the risk premium for a merely stochastic evolution of temperature (no epistemological uncertainty). It is small compared to the risk premium resulting from climate sensitivity uncertainty.

The second set of novel channels arises from the fact that climate sensitivity is a structural parameter that governs both the *marginal* impact of CO_2 released today and the long-run

temperature *level*. We show that, (only!) as a result of the structural parameter's "dual impact", risk aversion and damage convexity have a direct impact on the risk premium. Quantitatively, the damage convexity channel dominates in all of our scenarios. In contrast to common belief, the welfare based prudence and risk aversion effects contribute only a small fraction to the climate policy's risk premium. The empirical literature and some of the modeling literature uses stochastic temperature evolution as a proxy for uncertain future climate change. Our results raise awareness that the optimal response to an uncertain climate differs fundamentally from the response to a merely stochastic evolution of the climate. While our paper quantifies the impact for a global "mitigation investment", the basic insight carries over to adaptation investments.

We find risk premia around 25% for our base calibration and when varying time preference, risk aversion, and updating DICE to more current economic data. Changing DICE's quadratic damage function to a cubic form substantially increases the risk premium to around 65-75% of its deterministic value (depending on time preference). Epstein-Zin preferences disentangle risk aversion from the intertemporal elasticity of substitution. They allow us to increase risk aversion more drastically (RRA=6) without excessively discounting the future. This increase in risk aversion raises the baseline premium to over 30%. However, the long-run risk literature also suggests that Epstein-Zin preferences should use a higher elasticity of intertemporal substitution, which reduces the risk premium back to 20%. The premium is large in absolute terms, but the higher intertemporal elasticity reduces the growth-based discounting and increases the deterministic SCC even more substantially than the risk premium. This large positive premium contrasts sharply with the impact of growth uncertainty on optimal climate policy studied in Jensen and Traeger (2014) and Cai and Lontzek (2019), who find a negative risk premium for the same increase in the elasticity of intertemporal substitution. The difference arises yet again because climate sensitivity is a structural parameter with "dual impact".

Our study emphasizes the analytic drivers of the risk premium and quantifies the premium for a variety of scenarios. We use a simple climate change model with reasonable performance. We leave more detailed climate change models that accurately generate higher order moments of the climate sensitivity distribution to future and paralleling research. We focus on climate sensitivity uncertainty. Several exciting papers are quantifying the climate policy impact of multiple interacting uncertainties and we hope that the concepts developed here will also help to better understand the interactions.

References

Arias, P., N. Bellouin, E. Coppola, R. Jones, G. Krinner, J. Marotzke, V. Naik, M. Palmer, G.-K. Plattner, J. Rogelj, M. Rojas, J. Sillmann, T. Storelvmo, P. Thorne, B. Trewin, K. Achuta Rao, B. Adhikary, R. Allan, K. Armour, G. Bala, R. Barimalala, S. Berger, J. Canadell, C. Cassou, A. Cherchi, W. Collins, W. Collins, S. Connors, S. Corti, F. Cruz, F. Dentener, C. Dereczynski, A. Di Luca, A. Diongue Niang, F. Doblas-Reyes, A. Dosio, H. Douville, F. Engelbrecht, V. Eyring, E. Fischer, P. Forster, B. Fox-Kemper, J. Fuglestvedt, J. Fyfe, N. Gillett, L. Goldfarb, I. Gorodetskaya, J. Gutierrez, R. Hamdi, E. Hawkins, H. Hewitt, P. Hope, A. Islam, C. Jones, D. Kaufman, R. Kopp, Y. Kosaka, J. Kossin, S. Krakovska, J.-Y. Lee, J. Li, T. Mauritsen, T. Maycock, M. Meinshausen, S.-K. Min, P. Monteiro, T. Ngo-Duc, F. Otto, I. Pinto, A. Pirani, K. Raghavan, R. Ranasinghe, A. Ruane, L. Ruiz, J.-B. SallÃ©e, B. Samset, S. Sathyendranath, S. Seneviratne, A. SÃ¶rensson, S. Szopa, I. Takayabu, A.-M. TrÃ©guier, B. van den Hurk, R. Vautard, K. von Schuckmann, S. Zaehle, X. Zhang, and K. Zickfeld (2021). *Technical Summary*, pp. 33â^'144. Cambridge, United Kingdom and New York, NY, USA: Cambridge University Press.

Arrow, K. J. (1965). *Aspects of the Theory of Risk Bearing*. Helsinki: Helsinki: Yrjö Jahnsonian Sätiö.

Bansal, R., D. Kiku, I. Shaliastovich, and A. Yaron (2014). Volatility, the macroeconomy, and asset prices. *The Journal of Finance 69*(6), 2471–2511.

Bansal, R., D. Kiku, and A. Yaron (2010). Long-run risks, the macro-economy and asset prices. *American Economic Review: Papers & Proceedings 100*, 542–546.

Bansal, R., D. Kiku, and A. Yaron (2012). An empirical evaluation of the long-run risks model for asset prices. *Critical Finance Review 1*, 183–221.

Bansal, R. and A. Yaron (2004). Risks for the long run: A potential resolution of asset pricing puzzles. *The Journal of Finance 59*(4), 1481–509.

Barnett, M., W. Brock, and L. P. Hansen (2020, 02). Pricing uncertainty induced by climate change. *The Review of Financial Studies 33*(3), 1024–1066.

Cai, Y. and T. S. Lontzek (2019). The social cost of carbon with economic and climate risks. *Journal of Political Economy 127*(6), 2684–2734.

Chen, X., J. Favilukis, and S. C. Ludvigson (2013). An estimation of economic models with recursive preferences. *Quantitative Economics 4*, 39–83.

Collin-Dufresne, P., M. Johannes, and L. A. Lochstoer (2016). Parameter Learning in General Equilibrium: The Asset Pricing Implications. *American Economic Review 106*(3), 664–698.

Crost, B. and C. P. Traeger (2013). Optimal climate policy: Uncertainty versus Monte-Carlo. *Economic Letters 120*, 552–558.

Daniel, K. D., R. B. Litterman, and G. Wagner (2019). Declining CO2 price paths. *Proceedings of the National Academy of Sciences 116*(42), 20886–20891.

Dietz, S., C. Gollier, and L. Kessler (2018). The climate beta. *Journal of Environmental Economics and Management 87*, 258–278.

Dietz, S., J. Rising, T. Stoerk, and G. Wagner (2021). Economic impacts of tipping points in the climate system. *Proceedings of the National Academy of Sciences 118*(34), e2103081118.

Dietz, S. and F. Venmans (2019). Cumulative carbon emissions and economic policy: in search of general principles. *Journal of Environmental Economics and Management 96*, 108–29.

Drupp, M. A., M. C. Freeman, B. Groom, and F. Nesje (2018). Discounting disentangled. *American Economic Journal: Economic Policy 10*(4), 109–134.

EPA (2023). Report on the social cost of greenhouse gases: Estimates incorporating recent scientific advances.

Epstein, L. G., E. Farhi, and T. Strzalecki (2014, 09). How much would you pay to resolve long-run risk? *The American Economic Review 104*, 2680–2697.

Epstein, L. G. and S. E. Zin (1989, July). Substitution, risk aversion, and the temporal behavior of consumption and asset returns: A theoretical framework. *Econometrica 57*(4), 937–69.

Feenstra, R. C., R. Inklaar, and M. P. Timmer (2015). The next generation of the Penn World Table. *American Economic Review 105*(10), 3150–3182. PWT 10.0, downloaded 3/05/2021 at http://www.ggdc.net/pwt/.

Gerlagh, R. and M. Liski (2018). Consistent climate policies. *Journal of the European Economic Association 16*(1), 1–44.

Giglio, S., B. Kelly, and J. Stroebel (2021). Climate finance. *Annual Review of Financial Economics 13*(Volume 13, 2021), 15–36.

GISTEMP, T. (2018). GISS surface temperature analysis (GISTEMP). *NASA Goddard Institute for Space Studies version 3*. https://data.giss.nasa.gov/gistemp/.

GISTEMP Team (2024). Giss surface temperature analysis (gistemp), version 4. `https://data.giss.nasa.gov/gistemp/`. Dataset accessed 20YY-MM-DD.

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

Hambel, C., H. Kraft, and E. Schwartz (2021). Optimal carbon abatement in a stochastic equilibrium model with climate change. *European Economic Review 132*, 103642.

Heal, G. (1984). Interactions between economy and climate: A framework for policy design under uncertainty. In V. K. Smith and A. D. White (Eds.), *Advances in Applied Microeconomics*, Volume 3, pp. 151–168. Greenwich, CT: JAI Press.

Heijmans, R. and R. Gerlagh (2019, 1). Regulating global externalities. Workingpaper, CentER, Center for Economic Research. CentER Discussion Paper Nr. 2019-001.

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

Hoel, M. and L. Karp (2002). Taxes versus quotas for a stock pollutant. *Resource and Energy Economics 24*, 367–384.

Howard, P. H. and T. Sterner (2017). Few and not so far between: A meta-analysis of climate damage estimates. *Environmental and Resource Economics 68*, 197–225.

Interagency Working Group on Social Cost Greenhouse Gases, U. S. G. (2021). Technical support document: Social cost of carbon, methane, and nitrous oxide interim estimates under executive order 13990. *Department of Energy*.

Interagency Working Group on Social Cost of Carbon, U. S. G. (2016). Technical support document: Technical update of the social cost of carbon for regulatory impact analysis under executive order 12866. *Department of Energy*.

Interagency Working Group on the Social Cost of Carbon, U. S. G. (2013). Technical support document: Technical update of the social cost of carbon for regulatory impact analysis under executive order 12866. *Department of Energy*.

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

Jensen, S. and C. Traeger (2024). Uncertainty in climate-economic modeling. Volume 1 of *Handbook of the Economics of Climate Change*, pp. 351–423. North-Holland.

Karp, L. and C. Traeger (2018). Prices versus quantities reassessed. *CESifo Working Paper* (7331).

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

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

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

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

Kelly, D. L. and Z. Tan (2015). Learning and climate feedbacks: Optimal climate insurance and fat tails. *Journal of Environmental Economics and Management 72*, 98–122.

Kimball, M. S. (1990, January). Precautionary Saving in the Small and in the Large. *Econometrica 58*(1), 53–73.

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

Leland, H. E. (1968). Saving and uncertainty: The precautionary demand for saving. *The Quarterly Journal of Economics 82*(3), 465–473.

Lemoine, D. (2021a). The climate risk premium: how uncertainty affects the social cost of carbon. *Journal of the Association of Environmental and Resource Economists 8*(1), 27–57.

Lemoine, D. (2021b). The climate risk premium: how uncertainty affects the social cost of carbon. *Journal of the Association of Environmental and Resource Economists 8*(1), 27–57.

Lemoine, D. and C. Traeger (2014). Watch your step: Optimal policy in a tipping climate. *American Economic Journal: Economic Policy 6*(1), 137–166.

Lemoine, D. and C. Traeger (2016a). Economics of tipping the climate dominoes. *Nature Climate Change 6*, 514–519.

Lemoine, D. M. and C. P. Traeger (2016b). Ambiguous tipping points. *Jounral of Economic Behaviour & Organization 132B*(2667866), 5–18.

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

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

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

Nakamura, E., D. Sergeyev, and J. Steinsson (2017, January). Growth-rate and uncertainty shocks in consumption: Cross-country evidence. *American Economic Journal: Macroeconomics 9*(1), 1–39.

Newell, R. G. and W. A. Pizer (2008). Indexed regulation. *Journal of Environmental Economics and Management 56*, 221–233.

Nordhaus, W. and P. Sztorc (2013). *DICE2013R:Introduction and User's Manual*. dicemodel.net.

Nordhaus, W. D. (1993). Optimal greenhouse-gas reductions and tax policy in the "dice" model. *The American Economic Review 83*(2), 313–317.

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

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

Pindyck, R. S. (2020). The social cost of carbon revisited. *Journal of Environmental Economics and Management 94(C)*, 140–160.

Pizer, W. A. and B. C. Prest (2019). Prices versus quantities with policy updating. *Journal of the Association of Environmental and Resource Economists 7*(3), 483–518.

Pratt, J. W. (1964). Risk aversion in the small and in the large. *Econometrica 32*(1/2), 122Â–136.

Rennert, K., F. Errickson, B. C. Prest, L. Rennels, R. G. Newell, W. Pizer, C. Kingdon, J. Wingenroth, R. Cooke, B. Parthum, et al. (2022). Comprehensive evidence implies a higher social cost of co2. *Nature 610*(7933), 687–692.

Rudik, I. (2020). Optimal climate policy when damages are unknown. *American Economic Journal: Economic Policy 12*(2), 340–373.

Rudik, I. and D. Lemoine (2017). Managing climate change under uncertainty: Recursive integrated assessment at an inflection point. *Annual Review of Resource Economics 9*, 117–42.

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

Taconet, N., C. Guivarch, and A. Pottier (2021). Social cost of carbon under stochastic tipping points. *Environmental and Resource Economics 78*(4), 709–737.

Traeger, C. P. (2014). A 4-stated dice: Quantitatively addressing uncertainty effects in climate change. *Environmental and Resource Economics 59*(1), 1–37.

Traeger, C. P. (2015). Closed-form integrated assessment and uncertainty. *CESifo Working Paper 5464*. Split into Traeger (2021) and Traeger (2023).

Traeger, C. P. (2018). Uncertainty in the analytic climate economy. SSRN Working Paper 3832735, SSRN.

Traeger, C. P. (2021). Uncertainty in the analytic climate economy. SSRN Working Paper 3832735, SSRN.

Traeger, C. P. (2023). Aceâ€"analytic climate economy. *American Economic Journal: Economic Policy 15*(3), 372–406.

UN Environment Programme, U. (2024). Emissions Gap Report 2024: No more hot air â€¦ please! United Nations Environment Programme. Published October 24, 2024.

Van den Bremer and Van der Ploeg (2018). The risk-adjusted carbon price. *OxCarre Research Paper* (203).

van den Bremer, T., C. Hambel, and F. van der Ploeg (2024). Three reasons to price carbon under uncertainty: Accuracy of simple rules. CEPR Discussion Paper DP19645, CEPR, Paris and London.

van den Bremer, T. S. and F. van der Ploeg (2021). The risk-adjusted carbon price. *American Economic Review 111*(9), 2782–2810.

Van der Ploeg, R., S. Dietz, A. Rezai, and F. Venmans (2020). Are economists getting climate dynamics right and does it matter? Economics Series Working Papers 900, University of Oxford, Department of Economics.

Vissing-Jorgensen, A. and O. P. Attanasio (2003). Stock-market participation, intertemporal substitution, and risk-aversion. *The American Economic Review 93*(2), 383–391.

Weitzman, M. L. (1974). Prices versus quantities. *Review of Economic Studies 41*, 477–491.

Weitzman, M. L. (2012). Ghg targets as insurance against catastrophic climate damages. *Journal of Public Economic Theory 14*(2), 221–244.

Weitzman, M. L. (2018, January). Prices or quantities dominate banking and borrowing. Working Paper 24218, National Bureau of Economic Research.

World Economic Forum, T. (2022). *The Global Risks Report 2022* (17 ed.). Geneva: The World Economic Forum.

World Economic Forum, T. (2024). *The Global Risks Report 2024* (19 ed.). Geneva: The World Economic Forum.

Zhang, L. K. J. (2012). Taxes versus quantities for a stock pollutant with endogenous abatement costs and asymmetric information. *Economic Theory 49*, 371–409.

Appendix

A Proof of Proposition 1

The social planner solves the Bellman equation (1)

$$V(K_t, M_t, T_t, t) = \max_{C_t, E_t} u_t(C_t) + \beta\, \mathbb{E}_t\, V(K_{t+1}, M_{t+1}, T_{t+1}, t+1)$$

subject to the equations of motion

$$\begin{aligned} K_{t+1} &= (1-\delta)K_t + F(K_t, T_t, E_t, t) - C_t \\ M_{t+1} &= M_{t+1}(M_t, E_t) \quad \equiv M(M_t, E_t, t) \\ T_{t+1} &= T_{t+1}(T_t, M_{t+1}) \quad \equiv T(T_t, M_{t+1}, t). \end{aligned}$$

As stated in the text, we assume that today's emissions add directly to the next period's carbon stock: $\frac{\partial M_1}{\partial E_0} = 1$. Without this assumption, the derived formula for the social cost of carbon has to be multiplied by the factor $\frac{\partial M_1}{\partial E_0}$. We slightly violate notation in using M_t for the stock of carbon in the atmosphere and for the function $M(\cdot, t)$ describing its evolution over time. This notation enables us to abbreviate more intuitively the change of the carbon stock in period t as a result of a change of the carbon stock in period $\tau < t$ as $\frac{\partial M_t}{\partial M_\tau} = \prod_{i=\tau+1}^{t} \frac{\partial M_i}{\partial M_{i-1}} \equiv \prod_{i=\tau+1}^{t} \frac{\partial M(M_{i-1}, E_{i-1}, i-1)}{\partial M_{i-1}}$. We proceed analogously with temperature, abbreviating the change of the warming level in period t as a result of a change of the warming level in period $\tau < t$ as $\frac{\partial T_t}{\partial T_\tau} = \prod_{i=\tau+1}^{t} \frac{\partial T_i}{\partial T_{i-1}}$.

The first order conditions for consumption and emissions imply

$$u_t'(C_t) = \beta\, \mathbb{E}_t\, \frac{\partial V(K_{t+1}, M_{t+1}, T_{t+1}, t+1)}{\partial K_{t+1}} \tag{10}$$

and

$$\begin{aligned} &\beta\, \mathbb{E}_t\, \tfrac{\partial V(K_{t+1}, M_{t+1}, T_{t+1}, t+1)}{\partial K_{t+1}} \left(- \tfrac{\partial F(K_t, T_t, E_t, t)}{\partial E_t} \right) \\ &= \beta\, \mathbb{E}_t \left(\tfrac{\partial V(K_{t+1}, M_{t+1}, T_{t+1}, t+1)}{\partial M_{t+1}} + \tfrac{\partial V(K_{t+1}, M_{t+1}, T_{t+1}, t+1)}{\partial T_{t+1}} \tfrac{\partial T_{t+1}}{\partial M_{t+1}} \right) \tfrac{\partial M(M_t, E_t, t)}{\partial E_t}. \end{aligned} \tag{11}$$

The term $-\frac{\partial F(K_t, T_t, E_t, t)}{\partial E_t}$ denotes the benefits from another emission unit. In period t, this emission benefit is deterministic, and the remaining stochastic term on the left, $\beta \frac{\partial V(K_{t+1}, M_{t+1}, T_{t+1}, t+1)}{\partial K_{t+1}}$, equals by equation (10) the marginal utility in period t. Thus, making use of equation (10) when solving equation (11) for the optimality condition of the marginal benefit from another

emission unit in the present ($t = 0$) delivers the expression

$$-\frac{\partial F}{\partial E_0} = \underbrace{\frac{1}{u_0'(C_0)}\beta\,\mathbb{E}_0\left(\frac{\partial V}{\partial M_1}+\frac{\partial V}{\partial T_1}\frac{\partial T_1}{\partial M_1}\right)\frac{\partial M(M_0,E_0,0)}{\partial E_0}}_{\equiv SCC_0}\,. \tag{12}$$

In this social optimum, the marginal benefit from emitting another ton of carbon $-\frac{F(K_t,T_t,E_t,t)}{\partial E_t}$ has to equal the social cost of carbon given by the expression on the right hand side of equaton (12).

By the envelope theorem, the Bellman equation (1) delivers

$$\begin{aligned}\frac{\partial V}{\partial T_t} &= \beta\,\mathbb{E}_t\left(\frac{\partial V}{\partial K_{t+1}}\frac{\partial F}{\partial T_t}+\frac{\partial V}{\partial T_{t+1}}\frac{\partial T(T_t,M_{t+1},t)}{\partial T_t}\right) \\ &= u_t'(C_t)\frac{\partial F}{\partial T_t}+\beta\,\mathbb{E}_t\,\frac{\partial V}{\partial T_{t+1}}\frac{\partial T(T_t,M_{t+1},t)}{\partial T_t}\end{aligned} \tag{13}$$

where we obtain the second equality using that $\frac{\partial F}{\partial T_t}$ is known in period t so that we can employ equation (10).[23] Advancing equation (13) by one period and reinserting it into equation (13) delivers

$$\begin{aligned}\frac{\partial V}{\partial T_t} &= u_t'(C_t)\frac{\partial F}{\partial T_t}+\beta\,\mathbb{E}_t\,u_{t+1}'(C_{t+1})\frac{\partial F}{\partial T_{t+1}}\frac{\partial T(T_t,M_{t+1},t)}{\partial T_t} \\ &\quad+\beta\,\mathbb{E}_t\,\beta\,\mathbb{E}_{t+1}\,\frac{\partial V}{\partial T_{t+2}}\frac{\partial T(T_{t+1},M_{t+2},t+1)}{\partial T_{t+1}}\frac{\partial T(T_t,M_{t+1},t)}{\partial T_t} \\ &= u_t'(C_t)\frac{\partial F}{\partial T_t}+\beta\,\mathbb{E}_t\,u_{t+1}'(C_{t+1})\frac{\partial F}{\partial T_{t+1}}\frac{\partial T_{t+1}}{\partial T_t}+\beta^2\,\mathbb{E}_t\,\frac{\partial V}{\partial T_{t+2}}\frac{\partial T_{t+2}}{\partial T_t},\end{aligned} \tag{14}$$

where we made use of the law of iterated expectations (tower rule) and the notation $\frac{\partial T_t}{\partial T_\tau} \equiv \prod_{i=\tau+1}^{t}\frac{\partial T_i}{\partial T_{i-1}}$ for $\tau < t$. For simplicity in subsequent notation we "define" $\frac{\partial T_t}{\partial T_t} = 1$. Repeating the step of advancing equation (13) and reinserting it into equation (14) eventually delivers

$$\frac{\partial V}{\partial T_t} = \mathbb{E}_t\sum_{\tau=t}^{\infty}u_\tau'(C_\tau)\frac{\partial F}{\partial T_\tau}\beta^{\tau-t}\frac{\partial T_\tau}{\partial T_t}+\lim_{\tau\to\infty}\beta^{\tau-t}\,\mathbb{E}_t\,\frac{\partial V}{\partial T_\tau}\frac{\partial T_\tau}{\partial T_t}.$$

We assume (see main text) that the dynamic system is well-defined so that the shadow value of the temperature increase $\frac{\partial V}{\partial T_\tau}$ grows sufficiently slow along the optimal path to make the limit approach zero (which should hold along any reasonable policy scenario).

[23] Let $(\Omega,\mathcal{F},\mathcal{P})$ denote the underlying probability space. The equations of motions are conditional on the events $\omega\in\Omega$. Let $\{\mathcal{F}_t\}_{t\geq 0}$ denote the filtration generated by the underlying stochastic processes. Then, $\frac{\partial F}{\partial T_t}$ is measurable with respect to the sigma algebra $\mathcal{F}_t$.

By the envelope theorem, the Bellman equation (1) also delivers

$$\begin{aligned}\frac{\partial V}{\partial M_t} &= \beta\,\mathbb{E}_t\left(\frac{\partial V}{\partial M_{t+1}} + \frac{\partial V}{\partial T_{t+1}}\frac{\partial T(T_t, M_{t+1}, t)}{\partial M_{t+1}}\right)\frac{\partial M(M_t, E_t, t)}{\partial M_t} \\ &= \mathbb{E}_t\left(\beta\frac{\partial V}{\partial M_{t+1}} + \beta\sum_{\tau=t+1}^{\infty} u'_\tau(C_\tau)\frac{\partial F}{\partial T_\tau}\beta^{\tau-(t+1)}\frac{\partial T_\tau}{\partial T_{t+1}}\frac{\partial T_{t+1}}{\partial M_{t+1}}\right)\frac{\partial M_{t+1}}{\partial M_t}\end{aligned} \tag{15}$$

making use of equation (14). Using the definitions $\frac{\partial M_t}{\partial M_\tau} \equiv \prod_{i=\tau+1}^{t}\frac{\partial M_i}{\partial M_{i-1}}$ and $\frac{\partial M_t}{\partial M_t} = 1$ and repeatedly advancing equation (15) and reinserting it into itself yields

$$\frac{\partial V}{\partial M_t} = \mathbb{E}_t\sum_{j=t+1}^{\infty}\beta^{j-t}\sum_{\tau=j}^{\infty}u'_\tau(C_\tau)\frac{\partial F}{\partial T_\tau}\beta^{\tau-j}\frac{\partial T_\tau}{\partial T_j}\frac{\partial T_j}{\partial M_j}\frac{\partial M_j}{\partial M_t} + \lim_{\tau\to\infty}\beta^{\tau-t}\,\mathbb{E}_t\frac{\partial V}{\partial M_\tau}\frac{\partial M_\tau}{\partial M_t}.$$

In the text, we assumed that also the shadow value of carbon (weighted with its persistence) does not outgrow the exponential damping of the discount factor and that the limit approaches zero.

Reinserting our findings into equation (12) implies

$$-\frac{\partial F}{\partial E_0} = SCC_0 = \frac{1}{u'_0(C_0)}\,\mathbb{E}_0\sum_{j=1}^{\infty}\sum_{\tau=j}^{\infty}\beta^\tau u'_\tau(C_\tau)\frac{\partial F}{\partial T_\tau}\frac{\partial T_\tau}{\partial T_j}\frac{\partial T_j}{\partial M_j}\frac{\partial M_j}{\partial M_1}\frac{\partial M_1}{\partial E_0}\ . \tag{16}$$

We assumed that $\frac{\partial M_1}{\partial E_0} = 1$. The indices τ of the second sum are always larger than the j-index of the first sum. For every τ the double-sum runs through the j indices lower or equal to τ, capturing the marginal emission unit's impact on warming in all previous periods. Thus, we can reorder the sum to the form

$$-\frac{\partial F}{\partial E_0} = SCC_0 = \frac{1}{u'_0(C_0)}\,\mathbb{E}_0\sum_{t=1}^{\infty}\sum_{\tau=1}^{t}\beta^t u'_t(C_\tau)\frac{\partial F}{\partial T_t}\frac{\partial T_t}{\partial T_\tau}\frac{\partial T_\tau}{\partial M_\tau}\frac{\partial M_\tau}{\partial M_1} \tag{17}$$

where we relabled $\tau \to t$ and $j \to \tau$ with respect to equation (16). In the main text, we abbreviated $\frac{\partial T_\tau}{\partial M_j} \equiv \frac{\partial T_\tau}{\partial T_j}\frac{\partial T_j}{\partial M_j}$.

B Proof of Proposition 2

Under Assumption A2, the emission and investment levels do not respond to the realizations of climate sensitivity. As a consequence, future capital levels remain independent of the climate sensitivity realizations, and all damage to production translates straight into consumption and welfare. By the envelope theorem we know that this assumption does not affect the first order calculations and the derivation of the social cost of carbon in section A. However, the assumption will matter further below when we take higher order derivatives of

the social cost of carbon. As a result of the assumption, the impact of climate sensitivity on the social cost of carbon is entirely transmitted through temperature, production increase, and consumption loss. Marking these, and only these dependencies restates equation (17) as

$$SCC_0 = \frac{1}{u_0'(C_0)}\ \mathbb{E}\left[\sum_{t=1}^{\infty}\beta^t \underbrace{u_t'\big(C_t(T_t(\mathrm{s}))\big)\ \left[-\tfrac{\partial F_t}{\partial T_t}(T_t(\mathrm{s}))\right]\ \left[\sum_{\tau=1}^{t}\tfrac{\partial T_t}{\partial M_\tau}(\mathrm{s})\tfrac{\partial M_\tau}{\partial E_0}\right]}_{\equiv z_t(\mathrm{s})}\right].$$

For now, we analyze the period t contribution individually. By Jensen's inequality, the uncertainty over climate sensitivity will increase the SCC's period t contribution if and only if the term $z_t(\mathrm{s})$ is convex. Moreover, the variance-based small-risk approximation adds a contribution to the deterministic $z_t(\mathbb{E}\,s)$ that is proportional to the climate sensitivity's variance and to $z_t''(s)$:

$$\begin{aligned}\mathbb{E}\,z_t(s) &\approx z_t(\mathbb{E}\,s) + \mathbb{E}\,z_t'([\mathbb{E}\,s])(s-\mathbb{E}\,s) + \frac{1}{2}\,\mathbb{E}\,z_t''([\mathbb{E}\,s])(s-\mathbb{E}\,s)^2\\ &= z_t(\mathbb{E}\,s) + \frac{1}{2}z_t''([\mathbb{E}\,s])\,\mathbb{V}\mathrm{ar}(s).\end{aligned}$$

Because s is a common uncertain parameter for all periods, the climate sensitivity's variance simultaneously multiplies the sum over per period contributions: $\mathbb{V}\mathrm{ar}(s)\ \sum_{t=1}^{\infty}\beta^t z_t''([\mathbb{E}\,s])$. Thus, for small risks, the present SCC's risk premium per unit of variance is

$$\frac{\Delta SCC_0}{\mathbb{V}\mathrm{ar}(s)} = \frac{1}{u_0'(C_0)}\sum_{t=1}^{\infty}\beta^t\frac{\mathbb{E}\,z_t(s) - z_t(\mathbb{E}\,s)}{\mathbb{V}\mathrm{ar}(s)} \approx \frac{1}{2u_0'(C_0)}\sum_{t=1}^{\infty}\beta^t z_t''([\mathbb{E}\,s]). \tag{18}$$

In the following, we calculate the second order derivative of the term $z_t(s)$, abbreviating under slight violation of notation $\mathbb{E}\,s$ by s. Our assumption that investment remains constant while consumption absorbs the climate sensitivity uncertainty translates into $C_t(T_t(\mathrm{s})) = F(T_t(\mathrm{s})) - I_t$, where investment $I_t = K_{t+1} - (1-\delta)K_t$ does not respond to climate sensitivity. As a consequence we have $\frac{\partial C_t(T_t(\mathrm{s}))}{\partial s} = \frac{\partial F(T_t(\mathrm{s}))}{\partial s}$ and consumption absorbs all the uncertainty.

Without Assumption A2, the optimal policy would spread the uncertainty between consumption and investment, slightly reducing the cost of uncertainty. By contrast, Assumption A2 slightly exaggerates the consumption response to uncertainty. Alternatively, we can assume that only a fraction χ of the production impact translates into consumption change and assume

$$\frac{\partial C_t(T_t(\mathrm{s}))}{\partial s} = \chi\,\frac{\partial F(T_t(\mathrm{s}))}{\partial s},$$

$0 < \chi \leq 1$. Picking χ as the consumption rate, we would more appropriately approximate the immediate cost of uncertainty for consumption but neglect the consequences for investment

that would respond proportionally to $1-\chi$. This approach would undervalue the costs of uncertainty. We derive the subsequent analysis using a general χ, and Proposition 2 corresponds to the case where $\chi = 1$. Then

$$
\begin{aligned}
z_t(s) &= -u_t'\big(F(T_t(\mathrm{s})) - I_t\big)\; F_t'(T_t(\mathrm{s}))\;\left[\sum_{\tau=1}^{t} \tfrac{\partial T_t}{\partial M_\tau}(\mathrm{s})\tfrac{\partial M_\tau}{\partial E_0}\right] \\
z_t'(s) &= -u_t''\big(F(T_t(\mathrm{s})) - I_t\big)\,\chi\, F_t'(T_t(\mathrm{s}))^2\, T_t'(\mathrm{s})\;\left[\sum_{\tau=1}^{t} \tfrac{\partial T_t}{\partial M_\tau}(\mathrm{s})\tfrac{\partial M_\tau}{\partial E_0}\right] \\
&\quad -u_t'\big(F(T_t(\mathrm{s})) - I_t\big)\, F_t''(T_t(\mathrm{s}))\, T_t'(\mathrm{s})\left[\sum_{\tau=1}^{t} \tfrac{\partial T_t}{\partial M_\tau}(\mathrm{s})\tfrac{\partial M_\tau}{\partial E_0}\right] \\
&\quad -u_t'\big(F(T_t(\mathrm{s})) - I_t\big)\, F_t'(T_t(\mathrm{s}))\;\left[\sum_{\tau=1}^{t} \tfrac{\partial^2 T_t}{\partial M_\tau \partial s}(\mathrm{s})\tfrac{\partial M_\tau}{\partial E_0}\right].
\end{aligned}
$$

Defining $A_t(s) = \sum_{\tau=1}^{t} \frac{\partial T_t}{\partial M_\tau}(\mathrm{s})\frac{\partial M_\tau}{\partial E_0}$ we obtain the second order derivative

$$
\begin{aligned}
z_t''(s) = &-u_t'''(\cdot)\;\chi^2\; F_t'(\cdot)^3\, T_t'(\cdot)^2 A_t(s) - 2u_t''(\cdot)\;\chi\; F_t'(\cdot)F_t''(\cdot)\, T_t'(\cdot)^2 A_t(s) \\
&-u_t''(\cdot)\;\chi\; F_t'(\cdot)^2\, T_t''(\cdot)A_t(s) - u_t''(\cdot)\;\chi\; F_t'(\cdot)^2\, T_t'(\cdot)A_t'(s) \\
&-u_t''(\cdot)\;\chi\; F_t'(\cdot)F_t''(\cdot)T_t'(\cdot)^2 A_t(s) - u_t'(\cdot)F_t'''(\cdot)T_t'(\cdot)^2 A_t(s) \\
&-u_t'(\cdot)F_t''(\cdot)T_t''(\cdot)A_t(s) - u_t'(\cdot)F_t''(\cdot)T_t'(\cdot)A_t'(s) \\
&-u_t''(\cdot)\;\chi\; F_t'(\cdot)^2\, T_t'(\cdot)A_t'(s) - u_t'(\cdot)\, F_t''(\cdot)T_t'(\cdot)A_t'(s) - u_t'(\cdot)\, F_t'(\cdot)A_t''(s)
\end{aligned}
$$

Collecting terms delivers

$$
\begin{aligned}
z_t''(s) = &-u_t'''(\cdot)\;\chi^2\; F_t'(\cdot)^3\, T_t'(\cdot)^2 A_t(s) - 3u_t''(\cdot)\;\chi\; F_t'(\cdot)F_t''(\cdot)\, T_t'(\cdot)^2 A_t(s) \\
&-u_t''(\cdot)\;\chi\; F_t'(\cdot)^2\, T_t''(\cdot)A_t(s) - 2u_t''(\cdot)\;\chi\; F_t'(\cdot)^2\, T_t'(\cdot)A_t'(s) - u_t'(\cdot)F_t'''(\cdot)T_t'(\cdot)^2 A_t(s) \\
&-u_t'(\cdot)F_t''(\cdot)T_t''(\cdot)A_t(s) - 2u_t'(\cdot)F_t''(\cdot)T_t'(\cdot)A_t'(s) - u_t'(\cdot)F_t'(\cdot)A_t''(s).
\end{aligned}
$$

Dividing by $u_t'(C_t)$, $-F'(T_t(s)) > 0$, and A_t results in

$$
\begin{aligned}
&\frac{z_t''(s)}{u_t'(\cdot)(-F'(\cdot))A_t} \\
&= -\tfrac{u_t'''(\cdot)}{u_t''(\cdot)} C_t \tfrac{-F_t'(\cdot)\,T_t'(\cdot)}{C_t} \left(-\tfrac{u_t''(\cdot)}{u_t'(\cdot)} C_t\right) \tfrac{-F_t'(\cdot)\,T_t'(\cdot)}{C_t}\,\chi^2 - 3\tfrac{u_t''(\cdot)}{u_t'(\cdot)} C_t \tfrac{-F_t'(\cdot)\,T_t'(\cdot)}{C_t} \tfrac{-F_t''(\cdot)}{-F_t'(\cdot)} T_t \tfrac{T_t'(\cdot)}{T_t}\,\chi \\
&\quad - \tfrac{u_t''(\cdot)}{u_t'(\cdot)} C_t \tfrac{-F_t'(\cdot)T'(\cdot)}{C_t} \tfrac{T_t''(\cdot)}{T_t'(\cdot)}\,\chi - 2\tfrac{u_t''(\cdot)}{u_t'(\cdot)} C_t \tfrac{-F_t'(\cdot)T_t'(\cdot)}{C_t} \tfrac{A_t'(\cdot)}{A_t}\,\chi + \tfrac{-F_t'''(\cdot)}{-F_t''(\cdot)} T_t \tfrac{T_t'(\cdot)}{T_t} \tfrac{-F_t''(\cdot)}{-F_t'(\cdot)} T_t \tfrac{T_t'(\cdot)}{T_t} \\
&\quad + \tfrac{-F_t''(\cdot)}{-F_t'(\cdot)} T_t \tfrac{T_t'(\cdot)}{T_t} \tfrac{T_t''(\cdot)}{T_t'(\cdot)} + 2\tfrac{-F_t''(\cdot)}{-F_t'(\cdot)} T_t \tfrac{T_t'(\cdot)}{T_t} \tfrac{A_t'(\cdot)}{A_t} + \tfrac{A_t''(\cdot)}{A_t'(\cdot)} \tfrac{A_t'(\cdot)}{A_t} \\
&= \mathrm{Prud}\,\hat{\epsilon}_{c,s}\,\mathrm{RRA}\,\hat{\epsilon}_{c,s} + 3\,\mathrm{RRA}\,\hat{\epsilon}_{c,s}\,\mathrm{Dam}_2\,\hat{\epsilon}_{T,s} + \mathrm{RRA}\,\hat{\epsilon}_{c,s} \tfrac{T_t''(\cdot)}{T_t'(\cdot)} + 2\,\mathrm{RRA}\,\hat{\epsilon}_{c,s} \tfrac{A_t'(\cdot)}{A_t} \\
&\quad + \mathrm{Dam}_3\,\hat{\epsilon}_{T,s}\,\mathrm{Dam}_2\,\hat{\epsilon}_{T,s} + \mathrm{Dam}_2\,\hat{\epsilon}_{T,s} \tfrac{T_t''(\cdot)}{T_t'(\cdot)} + 2\,\mathrm{Dam}_2\,\hat{\epsilon}_{T,s} \tfrac{A_t'(\cdot)}{A_t} + \tfrac{A_t''(\cdot)}{A_t'(\cdot)} \tfrac{A_t'(\cdot)}{A_t} \\
&= \mathrm{RRA}\,\hat{\epsilon}_{c,s} \left[2\tfrac{A_t'(\cdot)}{A_t} + \mathrm{Prud}\,\hat{\epsilon}_{c,s} + 3\,\mathrm{Dam}_2\,\hat{\epsilon}_{T,s}\right] + \mathrm{Dam}_2\,\hat{\epsilon}_{T,s} \left[2\tfrac{A_t'(\cdot)}{A_t} + \mathrm{Dam}_3\,\hat{\epsilon}_{T,s}\right] \\
&\quad + \tfrac{T_t''(\cdot)}{T_t'(\cdot)} \left[\mathrm{RRA}\,\hat{\epsilon}_{c,s} + \mathrm{Dam}_2\,\hat{\epsilon}_{T,s}\right] + \tfrac{A_t''(\cdot)}{A_t'(\cdot)} \tfrac{A_t'(\cdot)}{A_t} \\
&= \tfrac{1}{s^2}\Big\{ \mathrm{RRA}\,\epsilon_{c,s} \left[2\tfrac{sA_t'(\cdot)}{A_t} + \mathrm{Prud}\,\epsilon_{c,s} + 3\,\mathrm{Dam}_2\,\epsilon_{T,s}\right] \\
&\qquad + \mathrm{Dam}_2\,\epsilon_{T,s} \left[2\tfrac{sA_t'(\cdot)}{A_t} + \mathrm{Dam}_3\,\epsilon_{T,s}\right] \Big\} \\
&\quad + \tfrac{1}{s^2} \left\{ \tfrac{sT_t''(\cdot)}{T_t'(\cdot)} \left[\mathrm{RRA}\,\epsilon_{c,s} + \mathrm{Dam}_2\,\epsilon_{T,s}\right] + \tfrac{sA_t''(\cdot)}{A_t'(\cdot)} \tfrac{sA_t'(\cdot)}{A_t} \right\}
\end{aligned}
\tag{19}
$$

where we used the definitions $\mathrm{RRA} = -\frac{u_t'''(\cdot)}{u_t''(\cdot)} c$, $\mathrm{Prud} = -\frac{u_t'''(\cdot)}{u_t''(\cdot)} c$, $\mathrm{Dam}_2 = \frac{-F_t''(\cdot)}{-F_t'(\cdot)} T$, $\mathrm{Dam}_3 = \frac{-F_t'''(\cdot)}{-F_t''(\cdot)} T$, the semi-elasticities $\hat{\epsilon}_{c,s} = -\frac{dc}{ds}/c = -\chi\, F_t'(\cdot) T_t'(\cdot) \frac{1}{c}$ and $\hat{\epsilon}_{T,s} = \frac{dT}{ds}/T$, and the regular elasticities $\epsilon_{c,s} = -\frac{dc}{ds}/\frac{c}{s} = -\chi\, F_t'(\cdot) T_t'(\cdot) \frac{s}{c}$ and $\epsilon_{T,s} = \frac{dT}{ds}/\frac{T}{s}$.

Under Assumption A2 we have

$$
\begin{aligned}
T_{t+1} &= \gamma\; T_t + s\,\Gamma_1(M_{t+1}) + \Gamma_2(t) + \epsilon_t \\
&= \gamma\big(\gamma T_{t-1} + s\,\Gamma_1(M_t) + \Gamma_2(t-1) + \epsilon_{t-1}\big) \\
&\qquad + s\,\Gamma_1\big(M_{t+1}(M_t, E_t)\big) + \Gamma_2(t) + \epsilon_t \\
&= \gamma^{t+1} T_0 \;+ s \sum_{\tau=1}^{t+1} \gamma^{t+1-\tau} \Gamma_1(M_\tau(\cdot)) \;+ \sum_{\tau=0}^{t} \gamma^{t-\tau} \Gamma_2(\tau) \;+ \sum_{\tau=0}^{t} \gamma^{t-\tau}\; \tilde{\epsilon}_\tau.
\end{aligned}
\tag{20}
$$

Equation (27) in Appendix D.2 states the details of this relation for our numeric DICE implementation. Because the equation is linear in climate sensitivity, we find that $\frac{sT_t''(\cdot)}{T_t'(\cdot)} = 0$.

Moreover, from $A_t(s) = \sum_{\tau=1}^{t} \frac{\partial T_t}{\partial M_\tau}(\mathrm{s}) \frac{\partial M_\tau}{\partial E_0}$ we have

$$\frac{sA_t'(\cdot)}{A_t} = \frac{s\sum_{\tau=1}^{t} \frac{\partial^2 T_t}{\partial s \partial M_\tau}(\mathrm{s}) \frac{\partial M_\tau}{\partial E_0}}{\sum_{\tau=1}^{t} \frac{\partial T_t}{\partial M_\tau}(\mathrm{s}) \frac{\partial M_\tau}{\partial E_0}} = \frac{s\sum_{\tau=1}^{t} \frac{\partial \Gamma_1(M_\tau)}{\partial M_\tau} \frac{\partial M_\tau}{\partial E_0}}{\sum_{\tau=1}^{t} s\frac{\partial \Gamma_1(M_\tau)}{\partial M_\tau} \frac{\partial M_\tau}{\partial E_0}} = 1$$

$$\frac{sA_t''(\cdot)}{A_t'(\cdot)} = 0.$$

Thus, under Assumption A2, equation (19) simplifies to

$$\begin{aligned} &\frac{s^2\, z_t''(s)}{u_t'(\cdot)(-F'(\cdot))A_t(\cdot)} \\ &\quad = \underbrace{\mathrm{RRA}\, \epsilon_{c,s} \left[2 + \mathrm{Prud}\, \epsilon_{c,s} + 3\, \mathrm{Dam}_2\, \epsilon_{T,s}\right] + \mathrm{Dam}_2\, \epsilon_{T,s} \left[2 + \mathrm{Dam}_3\, \epsilon_{T,s}\right]}_{\equiv X_t(\cdot)}. \end{aligned} \tag{21}$$

Combining equations (18) and (21), and properly writing $\mathbb{E}\, s$ again for the expected climate sensitivity, we obtain the approximate risk premium to the SCC today in consumption equivalents as

$$\begin{aligned} \frac{\Delta SCC_0}{\mathbb{V}\mathrm{ar}(s)} &\approx \frac{1}{2u_0'(C_0)} \sum_{t=1}^{\infty} \beta^t z_t''([\mathbb{E}\, s]) = \frac{1}{2(\mathbb{E}\, s)^2} \sum_{t=1}^{\infty} \beta^t \frac{u_t'(C_t)}{u_0'(C_0)} (-F'(\cdot)) A_t(\cdot) X_t(\cdot) \\ &= \sum_{t=1}^{\infty} \sum_{\tau=1}^{t} \beta^t \frac{u_t'(c_t)}{u_0'(c_0)} \frac{\partial F_t}{\partial T_t} \frac{\partial T_t}{\partial M_\tau} \frac{\partial M_\tau}{\partial E_0} \frac{1}{2(\mathbb{E}\, s)^2} X_t(\cdot), \end{aligned}$$

which concludes the proof of the proposition.

Combining equations (18) and (19) generalizes the statement of Proposition 2 by abandoning Assumption 1.

Corollary 1. *Let $T_t(s)$ denote future temperature as a function of climate sensitivity (given emissions and today's state) and define $A_t(s) = \frac{\partial T_t}{\partial E_0}(s) \equiv \sum_{\tau=1}^{t} \frac{\partial T_t}{\partial M_\tau}(\mathrm{s}) \frac{\partial M_\tau}{\partial E_0}$. Under Assumption A2, uncertainty over climate sensitivity increases the social cost of carbon contribution from a given period if and only if*

$$\begin{aligned} X_t^*(\cdot) \equiv\, &\mathrm{RRA}\, \epsilon_{c,s} \left[2\frac{sA_t'(\cdot)}{A_t} + \mathrm{Prud}\, \epsilon_{c,s} + 3\, \mathrm{Dam}_2\, \epsilon_{T,s}\right] + \mathrm{Dam}_2\, \epsilon_{T,s} \left[2\frac{sA_t'(\cdot)}{A_t} + \mathrm{Dam}_3\, \epsilon_{T,s}\right] \\ &+ \frac{sT_t''(\cdot)}{T_t'(\cdot)} \left[\mathrm{RRA}\, \hat{\epsilon}_{c,s} + \mathrm{Dam}_2\, \hat{\epsilon}_{T,s}\right] + \frac{sA_t''(\cdot)}{A_t'(\cdot)} \frac{sA_t'(\cdot)}{A_t} \end{aligned}$$

is greater than zero. Arguments and their period-dependence are suppressed. Moreover, under a small risk approximation, the climate uncertainty premium is

$$\Delta SCC_0 \approx -\,\mathbb{E}_0 \sum_{t=1}^{\infty} \underbrace{\beta^t \frac{u_t'(C_t)}{u_0'(C_0)}}_{\substack{\text{consumption}\\ \text{discount}\\ \text{factor}}} \underbrace{\frac{\partial F_t}{\partial T_t}\frac{\partial T_t}{\partial E_0}}_{\substack{\text{marginal}\\ \text{emission}\\ \text{damage}}} \frac{\mathbb{V}\mathrm{ar}(s)}{2(\mathbb{E}\,s)^2}\, X_t^*(\cdot).$$

C Proof of Proposition 3

The proof runs analogously to that of Proposition 2 and we focus on the steps that are different. The relevant uncertainty is now the one generated by the shock process $(\epsilon_t)_{t\in\mathbb{N}}$ to temperature rather than the epistemological uncertainty of climate sensitivity s. We spell out equation (17) suppressing all dependencies but those on the stochastic shock process

$$SCC_0 = \frac{1}{u_0'(C_0)}\,\mathbb{E}\left[\sum_{t=1}^{\infty}\beta^t \underbrace{u_t'\big(C_t(T_t(\varepsilon_1,...,\varepsilon_{t-1}))\big)\Big[-\tfrac{\partial F_t}{\partial T_t}(T_t(\varepsilon_1,...,\varepsilon_{t-1}))\Big]\Big[\sum_{\tau=1}^{t}\tfrac{\partial T_t}{\partial M_\tau}\tfrac{\partial M_\tau}{\partial E_0}\Big]}_{\equiv v_t(\varepsilon_1,...,\varepsilon_{t-1})}\right].$$

The main difference to the case of climate sensitivity uncertainty is that the temperature response to CO_2, $\frac{\partial T_t}{\partial M_\tau}$, is independent of the temperature shock ε_t and that we face a sequence of stochastic shocks.

Because temperature behaves as an autoregressive process we have $\frac{\partial T_t}{\partial\,\varepsilon_\tau} = \gamma^{t-\tau-1}$ for $\tau < t$ and, therefore, $\frac{\partial v_t}{\partial\,\varepsilon_\tau} = \gamma^{t-\tau-1}\frac{\partial v_t}{\partial\,\varepsilon_{t-1}}$. We find the small risk approximation

$$\begin{aligned}
&\mathbb{E}\sum_t \beta^t v_t(\varepsilon_1,...,\varepsilon_{t-1})\\
&\quad\approx \sum_t \beta^t v_t(\mathbb{E}\,\varepsilon_1,...,\mathbb{E}\,\varepsilon_{t-1}) + \mathbb{E}\sum_t\sum_{\tau<t}\beta^t \frac{\partial v_t(\mathbb{E}\,\varepsilon_1,...,\mathbb{E}\,\varepsilon_{t-1})}{\partial\,\varepsilon_\tau}(\varepsilon_\tau - \mathbb{E}\,\varepsilon_\tau)\\
&\qquad + \mathbb{E}\,\frac{1}{2}\sum_t\sum_{\tau<t}\sum_{k\le t}\beta^t \frac{\partial^2 v_t(\mathbb{E}\,\varepsilon_1,...,\mathbb{E}\,\varepsilon_{t-1})}{\partial\,\varepsilon_\tau\,\partial\,\varepsilon_k}(\varepsilon_\tau - \mathbb{E}\,\varepsilon_\tau)(\varepsilon_k - \mathbb{E}\,\varepsilon_k)\\
&\quad= \sum_t \beta^t v_t(0,...,0) + \frac{1}{2}\sum_t\sum_{\tau<t}\beta^t\gamma^{2(t-\tau-1)}\frac{\partial^2 v_t(\mathbb{E}\,\varepsilon_1,...,\mathbb{E}\,\varepsilon_t)}{\partial{\varepsilon_{t-1}}^2}\,\mathbb{V}\mathrm{ar}(\varepsilon_\tau)\\
&\qquad + \sum_t\sum_{\tau<t}\sum_{k<\tau}\beta^t\gamma^{t-\tau-1}\gamma^{t-k-1}\frac{\partial^2 v_t(\mathbb{E}\,\varepsilon_1,...,\mathbb{E}\,\varepsilon_t)}{\partial{\varepsilon_{t-1}}^2}\,\mathrm{COV}(\varepsilon_\tau,\varepsilon_k).
\end{aligned}$$

Because the shocks are iid, the triple sum in the last line vanishes ($\mathrm{COV}(\varepsilon_\tau,\varepsilon_k)=0$) and $\mathbb{V}\mathrm{ar}\,\varepsilon \equiv \mathbb{V}\mathrm{ar}\,\varepsilon_\tau = \mathbb{V}\mathrm{ar}\,\varepsilon_t$. For small risk, the present day risk premium of the SCC per unit

variance is

$$\begin{aligned}\frac{\hat{\Delta}SCC_0}{\mathrm{VAR}(\varepsilon)} &= \frac{1}{u_0'(C_0)}\frac{\mathbb{E}\sum_{t=1}^{\infty}\beta^t v_t(\varepsilon_1,...,\varepsilon_t) - \sum_{t=1}^{\infty}\beta^t v_t(0,...,0)}{\mathrm{VAR}(\varepsilon)} \\ &\approx \frac{1}{2u_0'(C_0)}\sum_t\sum_{\tau\le t}\beta^t\gamma^{2(t-\tau)}\frac{\partial^2 v_t(0,...,0)}{\partial{\varepsilon_{t-1}}^2} \\ &= \frac{1}{2u_0'(C_0)}\sum_t\beta^t\frac{\partial^2 v_t(0,...,0)}{\partial{\varepsilon_{t-1}}^2}\sum_{\tau\le t}\gamma^{2(t-\tau-1)} \\ &= \frac{1}{2u_0'(C_0)}\sum_t\beta^t\frac{1-\gamma^{2t}}{1-\gamma^2}\frac{\partial^2 v_t(0,...,0)}{\partial{\varepsilon_{t-1}}^2}, \end{aligned} \tag{22}$$

where we used the closed-form expression for the geometric sum $\sum_{\tau<t}\gamma^{2(t-\tau-1)} = \sum_{l=0}^{t-1}\gamma^{2l} = \frac{1-\gamma^{2t}}{1-\gamma^2}$ (note that index τ starts at zero whereas index t starts at unity).

The evaluation of $\frac{\partial^2 v_t(0,...,0)}{\partial{\varepsilon_{t-1}}^2}$ follows closely that of $z''(s)$ in the proof of Proposition 2. In contrast to climate sensitivity uncertainty, the temperature shocks ε_t leave the temperature response to CO_2, $\frac{\partial T_t}{\partial M_\tau}$, unchanged. As a consequence, $\frac{\partial^2 T_t}{\partial M_\tau \partial \epsilon_t} = 0$ and the contributions relating to $A'(s)$ in the case of climate sensitivity will be missing in the case of temperature stochasticity. A derivation analogous to equation (19) yields

$$\begin{aligned}\frac{\frac{\partial^2 v_t(0,...,0)}{\partial{\varepsilon_{t-1}}^2}}{u_t'(\cdot)(-F'(\cdot))A_t} &= \tfrac{1}{T_t^2}\left\{\mathrm{RRA}\,\epsilon_{c,T}\left[\mathrm{Prud}\,\epsilon_{c,T} + 3\,\mathrm{Dam}_2\right] + \mathrm{Dam}_2\,\mathrm{Dam}_3\right\} \\ &\quad + \tfrac{1}{T_t^2}\underbrace{\left\{\frac{T_t T_t''(\cdot)}{T_t'(\cdot)}\left[\mathrm{RRA}\,\epsilon_{c,T} + \mathrm{Dam}_2\right]\right\}}_{\equiv Z_t(\cdot)}, \end{aligned} \tag{23}$$

where $\epsilon_{c,T} = -\frac{dc}{dT}/\frac{c}{T} = F_t'(\cdot)\frac{T}{c}$ (and the other elasticity is unity). Under Assumption 2, we know from equation (20) that $T_t''(\cdot) = \frac{\partial^2 T_t(\cdot)}{\partial \epsilon_{t-1}^2} = 0$ and the second line vanishes.

Combining equations (22) and (23) (with $T_t''(\cdot) = 0$) delivers the stochasticity premium to the SCC today in consumption equivalents as

$$\begin{aligned}\frac{\hat{\Delta}SCC_0}{\mathbb{V}\mathrm{ar}(\epsilon)} &\approx \frac{1}{2u_0'(C_0)}\sum_{t=1}^{\infty}\beta^t\frac{1-\gamma^{2t}}{1-\gamma^2}\frac{\partial^2 v_t(0,...,0)}{\partial{\varepsilon_{t-1}}^2} \\ &= \frac{1}{2}\sum_{t=1}^{\infty}\beta^t\frac{1-\gamma^{2t}}{1-\gamma^2}\frac{u_t'(C_t)}{u_0'(C_0)}\frac{-F'(\cdot)A_t(\cdot)}{T_t}Z_t(\cdot) \\ &= \sum_{t=1}^{\infty}\sum_{\tau=1}^{t}\beta^t\frac{1-\gamma^{2t}}{1-\gamma^2}\frac{u_t'(C_t)}{u_0'(C_0)}\frac{\partial F_t}{\partial T_t}\frac{\partial T_t}{\partial M_\tau}\frac{\partial M_\tau}{\partial E_0}\frac{1}{2(T_t)^2}Z_t(\cdot), \end{aligned}$$

where the terms are evaluated at the expected shock values ($\epsilon_t = 0\ \forall t$). Emphasizing this

evaluation for temperatures we write $\mathbb{E}\,T_t$ instead of T_t in the Proposition. Having proven the second part of Proposition 3, the first part follows as in the proof of Proposition 2 from Jensen's inequality and equation (23).

For an additional corollary we relax Assumption 1 in the last argument.

Assumption 1'. *The temperature evolution is:*

$$T_{t+1} = \gamma\ T_t + s\,\Gamma_1(M_{t+1}) + \Gamma_2(t) + \Gamma_3(\epsilon_t)$$

where $\gamma \in (0,1)$, s is potentially uncertain, Γ_1, Γ_2, and Γ_2 are real functions, and $(\epsilon_t)_{t\in\mathbb{N}}$ is an iid process of mean zero temperature shocks.

Then $T_t''(\cdot) = \Gamma_3''(\epsilon_t)$ no longer vanishes. Note that we maintain the autoregressive structure of temperature to keep the simplified sum in equation (22) over shocks in different periods. Then, combining equations (22) and (23) delivers the following corollary.

Corollary 2. *Under Assumptions A1' and A2, temperature stochasticity increases the social cost of carbon contribution from a given period if and only if*

$$Z_t^*(\cdot) \equiv \mathrm{RRA}\,\epsilon_{c,T}\,[\mathrm{Prud}\,\epsilon_{c,T} + 3\,\mathrm{Dam}_2] + \mathrm{Dam}_2\,\mathrm{Dam}_3 + \frac{T_t\Gamma_3''(0)}{\Gamma_3'(0)}\,[\mathrm{RRA}\,\epsilon_{c,T} + \mathrm{Dam}_2]$$

is greater than zero. Arguments and their period-dependence are suppressed.
Moreover, under a small risk approximation, the temperature stochasticity premium is

$$\hat{\Delta}SCC_0 \approx \sum_{t=1}^{\infty}\sum_{\tau=1}^{t} \underbrace{\frac{1-\gamma^{2t}}{1-\gamma^2}}_{\substack{\textit{intertemporal}\\ \textit{temperature}\\ \textit{correlation}}} \underbrace{\beta^t \frac{u_t'(c_t)}{u_0'(c_0)}}_{\substack{\textit{consumption}\\ \textit{discount}\\ \textit{factor}}} \underbrace{\frac{\partial F_t}{\partial T_t}\frac{\partial T_t}{\partial M_\tau}\frac{\partial M_\tau}{\partial E_0}}_{\substack{\textit{marginal}\\ \textit{emission}\\ \textit{damage}}} \frac{\mathbb{V}\mathrm{ar}(\epsilon)}{2(\mathbb{E}\,T_t)^2}\,Z_t^*(\cdot).$$

D The Quantitative Model

This section introduces the quantitative model. It is based on the DICE-2013 model with updated modifications derived in Traeger (2014). These modification are helpful for a stochastic dynamic programming implementation, reducing the state space and interpolating exogenous processes in closed-form. Moreover, they allow us to fit the modified DICE in our generic IAM description presented in Section 2. As a co-benefit, these modifications also reduce the slightly exaggerated sluggishness of DICE's climate model (Traeger, 2014; Van der Ploeg et al., 2020).

D.1 Details of the quantitative model

The following model emulates DICE-2013. The three most notable differences are the annual time step (DICE-2013 features a five year time step), the infinite time horizon, and the replacement of the ocean feedbacks by exogenous processes. We adopt the latter simplification because the ocean carbon sink and ocean temperature would each require an own state variable in a recursive framework, which is computationally costly and would further complicate the analytic formulas. Conveniently, the "double-delayed" climate dynamics of DICE was recently found to be a bit sluggish compared to models used in the sciences Van der Ploeg et al. (2020). Instead, we calibrate a decay rate for atmospheric carbon and a temperature difference between atmosphere and ocean, which closely matches the behavior of a carbon cycle. For a detailed description of the procedure and its (very good) performance relative to the original DICE model see Traeger (2014), who compares both to the MAGICC model used in the IPCC reports. That paper also shows how to reformulate the decision problem expressing capital stock and consumption in efficient labor units, which is better suited for the numeric implementation. All parameters are characterized and quantified in Table 2 on page 61. These values are for DICE's starting period and iterated forward to the year 2020.[24]

Climate. Global average temperatures respond with a delay to the forcing of atmospheric carbon stocks M_t (relative to preindustrial level M_{pre}) and other non-CO2 forcing. Equation (29), marking conditional period t uncertainty with a tilde, becomes

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

In equation (29) we use the definitions

$$\begin{aligned} \chi_t(M_{t+1},t) &= \sigma_{forc}\frac{\log \frac{M_{t+1}}{M_{pre}}}{\log 2} + \frac{EF_{t+1}}{\eta_{forc}}\ , \quad \text{and} \\ \xi(T_t,t) \qquad &= (1-\sigma_{forc})T_t - \sigma_{ocean}\Delta T_t\ . \end{aligned} \tag{24}$$

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

$$\Delta T_t = \max\{0.74 + 0.028.*t - 0.00014.*t.^2\ ,\ 0.6\}\ .$$

Exogenous forcing EF_t from non-CO2 greenhouse gases, aerosols and other processes is

[24] In particular, prescribing DICE climate evolution to 2020 suggests correctly a present warming around 1.04°C. The economic values are updated in our alternative calibration to PWT 2019 values, see footnote **??** for details

assumed to follow the process

$$EF_t = EF_0 + 0.01(EF_{100} - EF_0) \times \max\{10 + t, 90\} \ .$$

We note that it starts out slightly negatively and that the 10 results from Nordhaus' start date 2000 for EF (instead of 2010 as for the other variables). Carbon in the atmosphere accumulates according to

$$M_{t+1} = M_{pre} + (M_t - M_{pre})\,(1 - \delta_M(t)) + E_t \quad \text{with} \tag{25}$$
$$\delta_{M,t} = \frac{\delta_{M,1}}{1 + \exp(\delta_{M,2} + \delta_{M,3}\ t)} \ .$$

The stock of CO_2 (M_t) exceeding preindustrial levels (M_{pre}) decays exponentially at the rate $\delta_M(M, t)$. This decay rate falls exogenously over time to replicate the carbon cycle, mimicking that the ocean reservoirs reduce their uptake rate as they fill up (see Traeger, 2014).

Emissions. Yearly CO_2 emissions E_t, consisting of industrial emissions and emissions from land use change an forestry B_t are

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

Emissions from land use change and forestry fall exponentially over time

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

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

$$\sigma_t = \sigma_0 \exp\left[\frac{g_\sigma((1 + \delta_\sigma)^t - 1)}{\log(1 + \delta_\sigma)}\right]$$

Economy. We denote global net output in per effective labor units by $y_t = \frac{Y_t}{A_t L_t}$. Population grows exogenously as

$$L_{t+1} = L_0 \left(\frac{L_{inf}}{L_0}\right)^{1-(1-g_L^*)^{t/5}}$$

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

as

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

The economy accumulates capital according to

$$k_{t+1} = \left[(1-\delta_k)\, k_t + y_t - c_t\right] \exp[-(g_{A,t} + g_{L,t})] \ , \tag{26}$$

where δ_K denotes the depreciation rate and c_t denotes aggregate global consumption of produced commodities per effective unit of labor. The growth rates in equation (26) result from the effective unit per labor normalization.

Global net output results from the (reduced-form) Cobb-Douglas production (k_t^κ) less abatement costs and less damages

$$y_t = \big(1 - \Lambda(\mu_t)\big)\big(1 - D(T_t)\big)k_t^\kappa$$

changing to $y_t = \frac{1-\Lambda(\mu_t)}{1+D(T_t)} k_t^\kappa$ for our DICE 2007 damage specification. Abatement costs are

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

and characterized as percent of output. They depend on the emission control rate $\mu_t \in [0,1]$. The coefficient of the abatement cost function Ψ_t follows

$$\Psi_t = \frac{a_0}{a_2} \sigma_t \exp[-g_\Psi \ t]$$

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

Climate damage as a fraction of world output depends on the temperature difference T_t of current to preindustrial temperature and is

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

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

We note that our DICE implementation indeed fits our generic integrated assessment model where emissions are treated as a production factor. For this purpose, we re-write the emission control rate $\mu_t = 1 - \frac{E_t}{E_t^{BAU}}$ as a function of actual and business as usual emissions and the DICE production function in non-effective-labor units as

$$F(K_t, T_t, E_t, t) = (A_t L_t)^{1-\kappa} K_t^\kappa \left(1 - D(T_t) - \Psi_t \Big[1 - \frac{E_t}{E_t^{BAU}(K_t)}\Big]^{2.8}\right)$$

where $E_t^{BAU}(K_t)$ is DICE's business as usual emission forecast as a function of the (endogenous) capital stock.

D.2 Temperature is linear in climate sensitivity

This brief subsection shows that temperature is linear in climate sensitivity, irrespectively of how many periods we look ahead into the future, in line with Assumption 1. Advancing the temperature equation by one period and inserting for the previous period's temperature gives

$$\begin{aligned}
\tilde{T}_{t+1} &= (1-\sigma_{forc})T_t + \tilde{s}\ \frac{\sigma_{forc}}{\ln 2}\ \ln\left(\frac{M_{t+1}}{M_{pre}}\right) + \frac{\sigma_{forc}}{\lambda}\ EF_t - \sigma_{ocean}\Delta T_t + \tilde{\epsilon}_t\ , \\
\tilde{T}_{t+2} &= (1-\sigma_{forc})\left((1-\sigma_{forc})T_t + \tilde{s}\ \frac{\sigma_{forc}}{\ln 2}\ \ln\left(\frac{M_{t+1}}{M_{pre}}\right) + \frac{\sigma_{forc}}{\lambda}\ EF_t - \sigma_{ocean}\Delta T_t + \tilde{\epsilon}_t\right) \\
&\quad + \tilde{s}\ \frac{\sigma_{forc}}{\ln 2}\ \ln\left(\frac{M_{t+2}}{M_{pre}}\right) + \frac{\sigma_{forc}}{\lambda}\ EF_{t+1} - \sigma_{ocean}\Delta T_{t+1} + \tilde{\epsilon}_{t+1}\ , \\
&= (1-\sigma_{forc})^2 T_t + (1-\sigma_{forc})\tilde{\epsilon}_t + \tilde{\epsilon}_{t+1} \\
&\quad + (1-\sigma_{forc})\left(\tilde{s}\ \frac{\sigma_{forc}}{\ln 2}\ \ln\left(\frac{M_{t+1}}{M_{pre}}\right)\right) + \tilde{s}\ \frac{\sigma_{forc}}{\ln 2}\ \ln\left(\frac{M_{t+2}}{M_{pre}}\right) \\
&\quad + (1-\sigma_{forc})\left(\frac{\sigma_{forc}}{\lambda}\ EF_t - \sigma_{ocean}\Delta T_t\right) + \frac{\sigma_{forc}}{\lambda}\ EF_{t+1} - \sigma_{ocean}\Delta T_{t+1}\ .
\end{aligned}$$

Iteratively we arrive at the formula

$$\begin{aligned}
\tilde{T}_{t+\tau} = \tilde{s}\ \frac{\sigma_{forc}}{\ln 2}\ \sum_{i=1}^{\tau}(1-\sigma_{forc})^{\tau-i}\ln\frac{M_{t+i}}{M_{pre}}\ + \sum_{i=0}^{\tau-1}(1-\sigma_{forc})^{\tau-i-1}\ \tilde{\epsilon}_{t+i} \\
+ \sum_{i=0}^{\tau-1}(1-\sigma_{forc})^{i}\left(\frac{\sigma_{forc}}{\lambda}EF_{t+i} - \sigma_{ocean}\Delta T_{t+i}\right) + (1-\sigma_{forc})^{\tau}T_t\ .
\end{aligned} \tag{27}$$

In particular, we observe that temperature in any period $t+\tau$ depends linearly on climate sensitivity.

D.3 Useful derivatives

By equation (27), we find for $t, \tau \geq 1$

$$\begin{aligned}
\frac{\partial T_{t+\tau}}{\partial M_t} &= \tilde{s}\ \frac{\sigma_{forc}}{\ln 2}\ (1-\sigma_{forc})^{\tau}\ \frac{1}{M_t} \\
\frac{\partial T_{t+\tau}}{\partial T_t} &= (1-\sigma_{forc})^{\tau} \\
\frac{\partial T_{t+\tau}}{\partial\, \varepsilon_t} &= (1-\sigma_{forc})^{\tau-1}. \\
\frac{\partial T_{t+\tau}}{\partial s} &= \frac{\sigma_{forc}}{\ln 2}\ \sum_{i=1}^{\tau}(1-\sigma_{forc})^{\tau-i} \ln \frac{M_{t+i}}{M_{pre}}
\end{aligned}$$

Moreover, by equation (25), we have

$$\begin{aligned}
\frac{\partial M_{t+\tau}}{\partial M_{t+\tau-1}} &= 1-\delta_M(t+\tau-1) \\
\frac{\partial M_{t+\tau}}{\partial M_t} &= \prod_{i=t}^{t+\tau-1}\left(1-\delta_M(i)\right).
\end{aligned}$$

In particular it follows that

$$\sum_{\tau=1}^{t}\frac{\partial T_t}{\partial M_\tau}\frac{\partial M_\tau}{\partial E_0}\ =\ \sum_{\tau=1}^{t}\tilde{s}\ \frac{\sigma_{forc}}{\ln 2}\ (1-\sigma_{forc})^{t-\tau}\ \frac{1}{M_\tau}\prod_{i=\tau}^{t-1}(1-\delta_M(i))$$

The partial derivative $\frac{\partial T_{t+\tau}}{\partial s_t}$ above only captures the direct impact of a change in climate sensitivity onto temperature. Yet, policy response will partly crowd out the direct effect. We denote the joint impact on temperature by Including as well the impact of a change in climate sensitivity on the optimally controlled emission we write the joined impact as the total derivative

$$\begin{aligned}
\frac{dT_{t+\tau}}{ds_t} &\equiv \frac{\partial T_{t+\tau}}{\partial s_t}+\sum_{j=0}^{i-1}\frac{\partial T_{t+\tau}}{\partial E_{t+j}}\frac{dE_{t+j}}{ds} \\
&= \frac{\sigma_{forc}}{\ln 2}\ \sum_{i=1}^{\tau}(1-\sigma_{forc})^{\tau-i}\left(\ln\frac{M_{t+i}}{M_{pre}}+\frac{s}{M_{t+i}}\ \sum_{j=0}^{i-1}\frac{dE_{t+j}}{ds}\prod_{l=t+j}^{t+i-1}\left(1-\delta_M(l)\right)\right). \qquad (28)
\end{aligned}$$

Note, somewhat unclear whether sum over j should start at 0 or 1. I picked 0 because that way next period emission stock is affected by this period emissions that can respond to present period's update on climate sensitivity. Note that at the beginning of the Appendix when summarizing the DICE equation we have T_{t+1} affected by M_t, whereas here in all our equations we use that T_{t+1} is a function of M_{t+1}, which is then affected by M_t.

By equation (12) we obtain an approximation to the optimal emissions response to a change in the SCC. Let $\Delta E_t \equiv E_t^{unc} - E_t^{det}$ denote the emission response to introducing uncertainty over the climate sensitivity, which we use to approximate $\frac{dE_t}{ds}$ in equation (28) above. Similarly, let $\Delta\,\mathrm{SCC} = \mathrm{SCC}^{unc} - \mathrm{SCC}^{det}$ denote the response of the SCC to uncertainty, i.e., the risk premium. Then by equation (12)

$$-\frac{\partial^2 F(E_t^{det},\cdot)}{\partial E_t^2}\Delta E_t \approx \frac{-\frac{\partial F(E_t^{det}+\Delta E_t,\cdot)}{\partial E_t^{unc}} + \frac{\partial F(E_t^{det},\cdot)}{\partial E_t}}{\Delta E_t}\Delta E_t = \mathrm{SCC}_t^{unc} - \mathrm{SCC}_t^{det} = \Delta\,\mathrm{SCC}_t$$

$$\Rightarrow \Delta E_t = -\frac{\Delta\,\mathrm{SCC}_t}{\frac{\partial^2 F(E_t^{det},\cdot)}{\partial E_t^2}}.$$

For $\Delta\,\mathrm{SCC}_t$ we use our original (somewhat too high) climate risk premium. We can then also iterate a second time in order to use the already (once) adjusted risk premium.

bl Note 2: I am not sure we have actually numerically used the "$F(E,\cdot)$" formulation of DICE before that we use in the text? Is it already encoded? If I am not mistaken it should be that marginal befits in DICE are

$$MB_t(E_t) = \frac{\partial F}{\partial E_t} = (A_t L_t)^{1-\kappa} K_t^{\kappa} D(T_t)\Psi_t\Big[1 - \frac{E_t}{E_t^{BAU}}\Big]^{1.8}\frac{2.8}{E_t^{BAU}}$$

so that

$$-\frac{\partial^2 F}{\partial E_t^2} = (A_t L_t)^{1-\kappa} K_t^{\kappa} D(T_t)\Psi_t\Big[1 - \frac{E_t}{E_t^{BAU}}\Big]^{0.8}\frac{5.04}{{E_t^{BAU}}^2}.$$

Would be great if you could double check that this last one is correct O-:

E Full Stochastic Dynamic Programming Model

The full stochastic model takes into account that future decision makers have more information on climate dynamics than present decision makers. Section E.1 specifies the informational dynamics (Bayesian updating) and the Bellman equation.

E.1 Uncertainty and Bellman Equation

A dynamic model of climate sensitivity uncertainty naturally invites a Bayesian learning model; in contrast to our analytic formula, a full dynamic model has to explicitly specify how uncertainty resolves over time.[25] Our model follows Kelly and Kolstad (1999) and represents

[25] Our formula merely approximates today's risk premium and we only have to model future uncertainty as viewed from today. A full stochastic recursive dynamic programming model derives the uncertainty as

uncertainty over climate sensitivity as a Bayesian prior. The decision maker updates the prior annually based on stochastic temperature observations. Assuming a normal distribution for both the climate sensitivity prior and the temperature shocks (likelihood) results in conjugate priors; the updated posterior will be of the same class of distributions as the prior. This feature is extremely helpful as it allows us to track information merely through the mean and the variance of climate sensitivity, which fully characterize the normal distribution. Limiting the informational states to two helps substantially when solving the Bellman equation for an already state-rich integrated assessment model.
[26]

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

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

Better estimates of the climate sensitivity result in a skewed distribution, but to solve the stochastic dynamic programming problem we rely on conjugate priors using the normal distribution. Given this limitation, $\sigma^2_{s,0} = 3$ is a rounded-up empirical approximation to the set of distributions found in Stocker et al. (2013).[27]

We can learn the value of climate sensitivity from observing the CO_2 stock and temperatures over time. Every period the decision maker foresees what a future realization of the temperature teaches her about climate sensitivity distribution and updates her prior accordingly, where temperature follows

$$\tilde{T}_{t+1} = \chi_t(M_{t+1})\tilde{s}_t + \xi_t(T_t) + \tilde{\epsilon}_t \ . \tag{29}$$

The factor χ_t captures the forcing from atmospheric CO_2 and other greenhouse gases. The term ξ_t reflects that atmospheric warming is a slow process and approximately governed by an AR(1) type process with a moving target. See equation (24) for the precise functional forms of χ_t and ξ_t. The structural parameter $\tilde{s}$ is the unknown climate sensitivity over which the decision forms her prior; the time index merely relates to the decision maker's subjective

viewed from today from the actual uncertainty prevailing in future periods conditional on future information. It thereby requires a more explicit formulation of the stochastic processes.

[26] Ongoing work by **?** uses a lower order approximation on a thinned approximation grid (Smolyak grid) to track a more realistic climate sensitivity distribution. Kelly and Tan (2013) use the normal distribution on the inverse of climate sensitivity to generate fat-tails. Cutting them off to avoid infinite temperatures (including negatively infinite temperatures) they use an approximate Bayesian updating with a highly skewed climate sensitivity distribution. Clearly climate sensitivity is slightly skewed, yet the fear of fat-tails seems to be increasingly tamed by work climate constraints **?**

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

prior, the climate sensitivity itself does not change over time.

Each year random events $\tilde{\epsilon}$ shock temperature. These "weather fluctuations" are normally distributed with mean zero. For a *given* value of climate sensitivity, the next period's temperature is then normally distributed

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

The variance σ_T^2 is exogenous. We estimate the annual volatility of global mean temperature as $\sigma_T^2 = 0.042$ using GISTEMP (2018).[28] For analytic purposes we will also use considerably larger values. In particular, very high stochasticity ($\sigma_T^2 = 0.7$) inhibits learning almost completely, which allows us to numerically isolate the effect of uncertainty about climate sensitivity from the learning effect in section **??**.

The policy maker's posterior in period t is the prior conditional on historic temperature realizations $\Pi(s|\hat{T}_1, ..., \hat{T}_t)$. This posterior also depends on the historic CO_2 stock which we suppress for notational convenience. Given the current stock M_t, a realization of temperature $\hat{T}_{t+1}$ in the subsequent period results in the updated posterior $\Pi(s|\hat{T}_1, ..., \hat{T}_{t+1})$. In Appendix E.2 we show that the updated posteriors are again normally distributed so that at all times $\Pi(s|\hat{T}_1, ..., \hat{T}_t) = \mathcal{N}(\mu_{s,t}, \sigma_{s,t}^2)$ for some $\mu_{s,t}$ and $\sigma_{s,t}^2$. Appendix E.2 also derives the following updating rules for expected climate sensitivity and its variance

$$\mu_{s,t+1} = \frac{\chi_t^2 \sigma_{s,t}^2 \frac{\hat{T}_{t+1} - \xi_t}{\chi_t} + \sigma_T^2 \mu_{s,t}}{\chi_t^2 \sigma_{s,t}^2 + \sigma_T^2} \qquad \text{and} \qquad \sigma_{s,t+1}^2 = \frac{\sigma_T^2 \sigma_{s,t}^2}{\chi_t^2 \sigma_{s,t}^2 + \sigma_T^2}\ , \tag{30}$$

where $\hat{T}_{t+1}$ characterizes observed temperature, σ_T^2 is the variance of temperature stochasticity (ϵ_t-shocks), and χ_t and $\xi_t(M_t)$ abbreviate the relevant endogenous quantities of the DICE model affecting the updating equation. The updated expected value of climate sensitivity s is a weighted mean of the previous expected value and the inferred "climate sensitivity observation", $\frac{\hat{T}_{t+1} - \xi_t}{\chi_t}$. The weight on the new observation is proportional to the precision (the inverse of the variance) of the temperature and the magnitude of the multiplicative factor $\chi_t(M_t)$, which increases in the carbon stock. This stock dependence of the updates differs from the common normal-normal updating formula and introduces active learning: by increasing the carbon stock the decision maker can increase $\chi_t(M_t)$ and, thus, the learning speed.

The social planner maximizes her value function subject to the exogenous and endogenous equations of motion for the economy and the climate summarized in Appendix D.1 as well as

[28] Kelly and Kolstad (1999) and Leach (2007) both use $\sigma_T^2 = 0.1$. Averaging temperatures over 174 countries and estimating yearly fluctuations with respect to a common trend over 109 years results in our lower $\sigma_T^2 = 0.042$.

the informational updating equation (30). The physical state variables describing the system are capital k, CO_2 stock M, and temperature T, and the informational state variables $\mu_{s,t}$ and $\sigma^2_{s,t}$ characterizing the climate sensitivity prior. We use the additional state time t to capture exogenously evolving processes, such as population, technology, and time changing climate processes. For numerical reasons, we express capital (and production and consumption) in effective labor units, i.e. $k_t = {}^{K_t}/_{A_t L_t}$ where A_t is technology level and L_t population at time t. The Bellman equation reads

$$V(k_t, M_t, t, T_t, \mu_{s,t}, \sigma^2_{s,t}) = \max_{c_t, \mu_t} \frac{(c_t)^{1-\hat{\eta}}}{1-\hat{\eta}} \tag{31}$$
$$+\beta_t \, \mathbb{E}\left[V(k_{t+1}, M_{t+1}, t+1, \tilde{T}_{t+1}, \tilde{\mu}_{s,t+1}, \tilde{\sigma}^2_{s,t+1})\right],$$

where we marked one-step-ahead uncertainty by a tilde and the time index on the discount factor β once again incorporates time-varying population. See Appendix D.1 for the transformation of population weighted average utility to the above form.

We solve the dynamic programming equation (31) by function iteration, using the collocation method to approximate the value function. As basis functions we choose Chebychev polynomials with $35,200$ Chebychev nodes and coefficients. The normal distributions for temperature stochasticity and the climate sensitivity prior are approximated by Gauss-Legendre quadrature with 3 nodes.[29] Our convergence criterion is a change in the value function coefficients of less than 10^{-4}. The code is written in Matlab, we use the CompEcon toolbox by Miranda and Fackler (2002) to generate and evaluate the Chebychev polynomials, and we use KNITRO for the optimization at a given node.

E.2 Derivation of Updating Rules

This appendix derives the updating rules for the climate sensitivity prior and the predictive distribution for temperature. Let $l_t(x_{t+1}|s) = \mathcal{N}(\mu_{x,t+1}, \sigma^2_T|s, x_t, h_t)$ denote the likelihood function in period t. Then[30]

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

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

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

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

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

$$\propto \bar{\Pi} \cdot \exp\left(\frac{\frac{1}{2}\cancel{\left(\frac{(x-\xi_t)\chi_t}{\sigma_T^2}\right)^2} + 2\frac{(x-\xi_t)\chi_t}{\sigma_T^2}\frac{\mu_{s,t}}{\sigma_{s,t}^2} + \cancel{\left(\frac{\mu_{s,t}}{\sigma_{s,t}^2}\right)^2} - \cancel{\frac{(x-\xi_t)^2}{\sigma_T^2}\frac{\chi_t^2}{\sigma_T^2}} - \frac{\mu_{s,t}^2}{\sigma_{s,t}^2}\frac{\chi_t^2}{\sigma_T^2} - \frac{(x-\xi_t)^2}{\sigma_T^2}\frac{1}{\sigma_{s,t}^2} - \cancel{\frac{\mu_{s,t}^2}{\sigma_{s,t}^2}\frac{1}{\sigma_{s,t}^2}}}{\frac{\chi_t^2}{\sigma_T^2} + \frac{1}{\sigma_{s,t}^2}} \right)$$

$$\propto \bar{\Pi} \cdot \exp\left(-\frac{1}{2\sigma_T^2\sigma_{s,t}^2} \frac{(x-\xi_t)^2 - 2(x-\xi_t)\chi_t\mu_{s,t} + \mu_{s,t}^2\chi_t^2}{\frac{\chi_t^2}{\sigma_T^2} + \frac{1}{\sigma_{s,t}^2}} \right)$$

$$\propto \bar{\Pi} \cdot \exp\left(-\frac{1}{2}\frac{(x-\xi_t-\chi_t\mu_{s,t})^2}{\chi_t^2\sigma_{s,t}^2 + \sigma_T^2} \right) .$$

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

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

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

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

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

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

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

and variance

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

Table 2: Parameters of the model

Economic Parameters		
η	2	intertemporal consumption smoothing and risk aversion
$a = b_1$	0.00284	damage coefficient
b_2	2	damage exponent
δ_u	1.34%	pure rate of time preference
L_0	6838	in millions, population in 2010
L_∞	10500	in millions, asymptotic population
g_L^*	0.035	rate of convergence to asymptotic population
K_0	135	in trillion 2005-USD, global capital stock in 2010
δ_K	10%	depreciation rate of capital
κ	0.3	capital elasticity in production
A_0	0.0067	initial labor productivity; corresponds to total factor productivity of 0.02722 used in DICE
$g_{A,0}$	2.35%	initial growth rate of labor productivity; corresponds to total factor productivity of 1.65% used in DICE
δ_A	0.6%	rate of decline of productivity growth rate
σ_0	0.133	CO_2 emission in GtC per unit of GDP in 2010
$g_{\sigma,0}$	-1%	initial rate of decarbonization
δ_σ	-0.1%	rate of decline of the rate of decarbonization
a_0	1261	cost of backstop in 2010
a_1	2	ratio of initial over final backstop cost
a_2	2.8	cost exponent
g_Ψ	0.5%	rate of convergence from initial to final backstop cost
Climatic Parameters		
T_0	0.83	in °C, temperature increase of preindustrial in 2005
σ_T^2	0.011	temperature stochasticity
M_{pre}	588	in GtC, preindustiral stock of CO2 in the atmosphere
M_0	819	in GtC, stock of atmospheric CO_2 in 2010
$\delta_{M,1}$	1.0415%	initial rate of decay of CO2 in atmosphere
$\delta_{M,2}$	-5	timing correction for initial rate of decay
$\delta_{M,3}$	4%	decline of decay rate of CO2
B_0	0.42	in GtC, initial CO2 emissions from LUCF
g_B	-4%	growth rate of CO2 emisison from LUCF
$\mu_{s,0}$	3	climate sensitivity prior mean in $t = 0$
$\sigma_{s,0}^2$	3	climate sensitivity prior variance in $t = 0$
EF_0	-0.06	initial external forcing
EF_{100}	0.62	external forcing in year 2100 and beyond
σ_{forc}	2.6%	warming delay, heat capacity atmosphere
σ_{ocean}	0.18%	warming delay, ocean related