

Climate Change Uncertainty Spillover in the Macroeconomy*

Mike Barnett[‡]

William Brock*

Lars Peter Hansen**

August 23, 2021

Abstract

The design and conduct of climate change policy necessarily confronts uncertainty along multiple fronts. We explore the consequences of ambiguity over various sources and configurations of models that impact how economic opportunities could be damaged in the future. We appeal to decision theory under risk, model ambiguity and misspecification concerns to provide an economically motivated approach to uncertainty quantification. We show how this approach reduces the many facets of uncertainty into a low dimensional characterization that depends on the uncertainty aversion of a decision-maker or fictitious social planner. In our computations, we take inventory of three alternative channels of uncertainty and provide a novel way to assess them. These include i) *carbon dynamics* that capture how carbon emissions impact atmospheric carbon in future time periods; ii) *temperature dynamics* that depict how atmospheric carbon alters temperature in future time periods; iii) *damage functions* that quantify how temperature changes diminish economic opportunities. We appeal to geoscientific modeling to quantify the first two channels. We show how these uncertainty sources interact for a social planner looking to design a prudent approach to the social pricing of carbon emissions.

*University of Wisconsin and University of Missouri, Columbia

**University of Chicago

‡Arizona State University

*We thank Shirui Chen, Han Xu and Jiaming Wang for the computational support on this research. Zhenhuan Xie, Samuel Zhao, and especially Diana Petrova provided valuable help in preparing this manuscript. An online notebook, which includes supplemental results and the code used to derive our model solutions, is available at <https://climateuncertaintyspillover.readthedocs.io/en/latest/>. Financial support for this project was provided by the Alfred P. Sloan Foundation [grant G-2018-11113]. We benefited from valuable feedback from Fernando Alvarez, Marty Eichenbaum, Michael Greenstone, Kevin Murphy, Tom Sargent and Chris Sims for helpful conversations while preparing this manuscript. Finally, Per Krusell, Ishan Nath, and Mar Reguant provided thoughtful discussions of earlier versions of the research that helped us in subsequent revisions of this manuscript.

1 Introduction

There are many calls for policy implementation to address climate change based on confidence in our knowledge of the adverse impact of economic activity on the climate, and conversely the negative effects of climate change on economic outcomes. Our view is that the knowledge base to support quantitative modeling in the realm of climate change and elsewhere remains incomplete. While there is a substantial body of evidence demonstrating the adverse human imprint on the environment, uncertainty comes into play when we build quantitative models aimed at capturing the dynamic transmission of human activity on the climate and on how adaptation to climate change will play out over time. In many arenas, it has been common practice in discussions of economic policy to shunt uncertainty to the background when building and using quantitative models. To truly engage in “evidence-based policy” requires that we are clear both about the quality of the evidence and the sensitivity to the modeling inputs used to interpret the evidence. While the importance of quantifying uncertainty has been stressed and implemented in a variety of scientific settings, the analysis of economic policy provides some unique challenges. Specifically, our aim is to explore ways to incorporate this uncertainty for the purposes of making quantitative assessments of alternative courses of action while exploring a broader conceptualization of uncertainty than is typical in econometric analyses. We see this challenge as much more than putting standard errors on econometric estimates or incorporating risk (uncertainty with known probabilities) into the analysis. We turn to developments in dynamic decision theory as a guide to how we confront uncertainty in policy analysis.

In climate economics, Weitzman (2012), Wagner and Weitzman (2015) and others have emphasized uncertainty in the climate system’s dynamics and how this uncertainty could create fat-tailed distributions of potential damages. Relatedly, Pindyck (2013) and Morgan et al. (2017) find existing integrated assessment models in climate economics to be of little value in the actual prudent policy. We are sympathetic to their skepticism, and are not offering simple repairs to the existing integrated assessment models in this area nor quick modifications to EPA postings for the social cost of carbon. Nevertheless, we find value in the use of models to engage in a form of “quantitative storytelling” and we explore the consequences for policy when multiple models or specifications are entertained.¹ Instead of proceeding with comparing policies model by model, our ambition is to incorporate at least some of the model uncertainty into the formal analysis. That is, our aim is to explore ways to assess policies with a more explicit accounting for the limits to our understanding. In the climate economics arena, not only is there substantial uncertainty about the economic modeling inputs, but also about the geoscientific inputs.

Drawing on insights from decision theory and asset valuation, Barnett et al. (2020) proposed a framework for assessing uncertainty, broadly-conceived, to include ambiguity over alternative

¹The term “model” is used differentially in statistical discussions of uncertainty. For us, a model conditions on any unknown parameters. Thus, we differentiate a model from a parameterized family of models.

models and the potential form of the misspecification of each. As is demonstrated in their paper, this broad notion of uncertainty is reflected in an endogenously determined adjustment to the probabilities used to depict meaningful economic values. This adjustment pushes well beyond the familiar discussions of social discount rates in the environmental economics literature. But the examples in Barnett et al. scratch the surface of the actual quantitative assessment of uncertainty pertinent to the economics of climate change, and they abstract from setups in which the uncertainty is at least partially resolved in the future.

In sections 3 and 4, this paper takes inventory of the alternative sources of uncertainty that are pertinent to climate change policy:

- *carbon dynamics* mapping carbon emissions into carbon in the atmosphere
- *temperature dynamics* mapping carbon in the atmosphere into temperature changes
- *economic damage functions* that depict the fraction of the productive capacity that is reduced by temperature changes

We necessarily adopt some stark simplifications to make this analysis tractable. Many of the climate models are both high-dimensional and nonlinear. Rather than using those models directly, we rely on outcomes of pulse experiments applied to the models. We then take the outcomes of these pulse experiments as inputs into our simplified specification of the climate dynamics inside our economic model. We follow much of the environmental macroeconomic modeling literature in the use of *ad hoc* static damage functions, and explore the consequences of changing the curvature in these damage functions. Even with these simplifications, our uncertainty analysis is sufficiently rich to show how uncertainty about the alternative channels by which emissions induce economic damages interact in important ways. Modeling extensions that confront heterogeneity in exposure to climate change across regions will also open the door to the inclusion of cross-sectional evidence for measuring potential environmental damages.

Decision theory provides tractable ways to explore a tradeoff between projecting the “best guess” consequences of alternative courses of action versus “worst possible” outcomes among a set of alternative models. Rather than focusing exclusively on these extremal points, we allow our decision-maker to take intermediate positions in accordance with parameters that govern aversions to model ambiguity and potential misspecification. We presume a decision-maker confronts many dimensions of uncertainty and engages in a sensitivity analysis. To simplify the policy analysis, we consider a world with a “fictitious social planner.” Thus, we put to the side important questions pertaining to heterogeneity in the exposure to climate change and to the consequent policy objectives by different decision-makers. Instead, we simplify the policy implementation to that of a Pigouvian tax that eliminates the wedge between market valuation and social valuation. We use this setup to illustrate how uncertainty can contribute to social valuation while recognizing the

need for further model richness in future research. Our planner confronts risk, model ambiguity, and model misspecification formally and deduces a socially efficient emissions trajectory. The planner’s decision problem adds structure to the sensitivity analysis and reduces a potentially high-dimensional sensitivity analysis to a two-dimensional characterization of sensitivity parameterized by aversion to model ambiguity and potential misspecification. We describe formally in section 5 some convenient continuous-time formulations of decision theory designed so that recursive methods familiar in economic dynamics can be applied with tractable modifications.

We use the social cost of carbon (SCC) as a barometer for investigating the consequences of uncertainty for climate policy. In settings with uncertainty, the SCC is the economic cost to the current and future uncertain environmental and economic damages induced by an incremental increase in emissions. In effect, it is the current period cost of an adverse social cash flow. Borrowing insights from asset pricing, this cash flow should be discounted stochastically in ways that account for uncertainty. This follows in part revealing discussions in Golosov et al. (2014) and Cai et al. (2017) who explore some of the risk consequences for the social cost of carbon. We extend this by taking a broader perspective on uncertainty. The common discussion in environmental economics about what “rate” should be used to discount future social costs is ill-posed for the model ambiguity that we feature. Rather than a single rate, we borrow and extend an idea from asset pricing by representing broadly based uncertainty adjustments as a change in probability over future outcomes for the macroeconomy. As we argue formally in section 6, when we incorporate uncertainty we are pushed away from the commonly employed modular approaches for measuring the SCC as the modular components to the SCC become much more intertwined.² Drawing on insights from recursive approaches to economic dynamics adds clarity to how best to rationalize and quantify the SCC in presence of decision-maker uncertainty.

Finally, this paper extends previous work by “opening the hood” of climate change uncertainty and exploring which components have the biggest impact on valuation. Rather than embrace a “one-model-fits-all-type-of-approaches” perspective, we give three computational examples designed to illustrate different points. The example presented in section 7 is by far the most ambitious and sets the stage for the other two. This first example explores what impact of future information about environmental and economic damages, triggered by temperature anomaly thresholds, should have on current policy. It adds a dynamic richness missing from other treatments of model uncertainty. The second example, presented in section 8, implements a novel decomposition of uncertainty assessing the relative importance of uncertainties in carbon dynamics, temperature dynamics and damage function uncertainty. The approach that is described and implemented in section 8 is more generally applicable to other economic environments. Finally, the third example investigates the interacting implications of the uncertainties in the develop-

²See, for instance, National Academies of Sciences, Engineering and Medicine (2017) for a discussion and a defense for the modular approach.

ment of green technologies and in environmental damages for prudent policy. This example is developed in Section 9.

In the next section, we elaborate on some the prior contributions that motivate our analysis.

2 Some motivating literature

Palmer and Stevens (2019) take inventory of what we know about climate change from basic principles while noting the limits to the existing efforts at quantitative modeling. They articulate the disconnect between arguments made to advance environmental policy and the state of knowledge coming from climate science. Moreover, they argue for the systematic inclusion and quantification of stochastic components in climate models as a way to make a substantive improvement in predictive models from climate science, even though the “big picture” is quite settled. Palmer and Stevens proposed modeling improvements that are well beyond the ambition of our work, but we have a shared appreciation for explicit stochastic modeling. It is important for our uncertainty quantification methods that we incorporate explicit randomness to partially disguise the model ambiguity and misspecification from a decision-maker. We aim to enrich the policy discussions by acknowledging rather than disguising uncertainty.

Our specification of damage function uncertainty can be motivated in part by “tipping points” in the climate system. Consistent with our formulation, Sharpe and Lenton (2021) and Lenton (2020), while noting that the “great majority of climate tipping points are damaging ones, and they may be closer than is often assumed,” also present evidence for tipping points that open the door to a far greener and less damaged economy. Lenton, then, highlights the need to comprehensively study the uncertainty in such complicated, non-linear settings so that we can effectively risk-manage positive and negative tipping points. Our example includes the possibility of good news with the delay in the Poisson event realization. On the other hand, while we are illustrating an important message for policy-making, our example is too simplistic to connect formally to tipping point specifications and the resulting uncertainties.³

Hausfather and Peters (2020) noted that policymakers and researchers have increasingly designated scenarios as “business-as-usual” without good justification. The formal use of decision theory allows for useful distinctions between adverse scenarios that are possible and best guesses for decision-makers to trade off considerations between such projections.

In prior work, Rudik (2020) explored damage function uncertainty. He reviewed the extensive literature on damage functions and developed a Bayesian learning framework about uncertain damage function parameters together with an analysis of the effects of robustness concerns caused by misspecification of the damage functions. Our formulation of damage function uncertainty

³See Lemoine and Traeger (2016) and Cai et al. (2016) for an example of an economic analysis with tipping point uncertainty.

differs substantially from his because of our Poisson event that governs damage function steepness. In our analysis, observations prior to this event are not informative about the damage function curvature beyond a threshold yet to be attained.

3 Uncertain climate dynamics

In this section, we first describe some very tractable characterizations of cross-model variation in the dynamic responses of temperature to emission pulses. To support our analysis, we then build a simplified stochastic specification of the pulse responses.

3.1 Simple approximations to climate dynamics

Recent contributions to the climate science literature have produced low-dimensional approximations, emulators, and pulse experiments that provide tractable alternatives to full-scale Atmospheric-Oceanic General Circulation Models (AOGCMs) used by climate scientists. These results allow for the inclusion of climate models within economic frameworks in ways that can be informative and revealing. We use the pulse experiment results of Joos et al. (2013) and Geoffroy et al. (2013) across various carbon and climate dynamics models to build the set of models we will use in our uncertainty analysis.⁴

Joos et al. (2013) report the responses of atmospheric carbon concentration to emission pulses of one hundred gigatons of carbon for several alternative Earth System models. The emission pulse experiments follow a standardized model intercomparison analysis so that outcomes are directly comparable. We use the responses for nine such models to capture the variation and uncertainty present in models of carbon cycle dynamics.

We feed these responses for carbon concentration into log-linear approximations of temperature dynamics constructed by Geoffroy et al. (2013). In accordance with the Arrhenius (1896) equation, these dynamics relate the logarithm of carbon in the atmosphere to future temperature. The parameters that Geoffroy et al. (2013) constructed using their simplified representation differ depending on the model being approximated. We use the 16 models listed in Appendix A. Thus, we take the nine different atmospheric carbon responses as inputs into the 16 temperature dynamics approximations, giving us a total of 144 different temperature responses to emissions.⁵

Figure 1 captures the resulting temperature responses across various sets of these 144 models. The top panel provides the results based on all 144 models, the middle panel provides the results based on variation in the carbon models, and the bottom panel provides the results based on variation in the temperature models. In each case, the maximal temperature response to an

⁴See Seshadri (2017), Eby et al. (2009), Matthews et al. (2009), and MacDougall et al. (2017) for additional examples of work in this area.

⁵Appendix A provides additional details on the emission pulse responses from Joos et al. (2013), the approximating model of Geoffroy et al. (2013), and lists the specific models we use from these two studies.




Figure 1: Percentiles for temperature responses to emission impulses. The emission pulse was 100 gigatons of carbon (GtC) spread over the first year. The temperature units for the vertical axis have been multiplied by ten to convert to degrees Celsius per teraton of carbon (TtC). The boundaries of the shaded regions are the upper and lower envelopes. Top panel: percentiles for impulse responses including both carbon and temperature dynamic uncertainty. Center panel: responses obtained for the different carbon responses for nine models each averaged over the 16 models of temperature dynamics. Bottom panel: percentiles for the 16 temperature responses using each averaged over the nine models of carbon concentration dynamics.

emission pulse occurs at about a decade and the subsequent response is very flat. These dynamics are consistent with the response patterns featured by Ricke and Caldeira (2014).

The top panel of Figure 1 also reports the percentiles for each horizon computed using the 144 different temperature response functions from all the different combinations of models of carbon and temperature dynamics. While there are similar patterns across the temperature response functions, there is considerable heterogeneity in the magnitudes of the responses. For a further characterization of this heterogeneity, we compute the exponentially weighted average of each of these response functions and use them in our computations. We report the resulting histogram as the top panel of Figure 2.




Figure 2: Histograms for the exponentially weighted average responses of temperature to an emissions impulse from 144 different models using a rate $\delta = .01$.

The eventually flat trajectories of the temperature response functions are consistent with model comparisons made using what is called the transient climate response (TCRE) to CO_2 emissions. The TCRE is the ratio of CO_2 -induced warming realized over an interval of time to the

cumulative carbon emissions over that same time interval. This linear characterization provides a simplification suggested by Matthews et al. (2009) and others by targeting the composite response of the carbon and temperature dynamics instead of the components that induce it. MacDougall et al. (2017) provide a pedagogical summary of this literature and report a histogram for the TCRE computed for 150 model variants. Their histogram looks very similar to what we report in Figure 2.

The middle and bottom panels of Figure 1 show the contribution of uncertainty in temperature and carbon dynamics to the temperature impulse responses. In generating the middle panel of Figure 1, we computed the implied temperature responses for nine alternative models of atmospheric CO_2 dynamics averaging over the 16 models of temperature dynamics. In generating the lower panel of Figure 1, we computed the 16 temperature responses for 16 temperature models while averaging over the nine models of atmospheric CO_2 dynamics. Consistent with the results reported by Ricke and Caldeira (2014), we find heterogeneity in the temperature responses to be more prominent than that coming from the atmospheric CO_2 dynamics.⁶

3.2 Stochastic climate pulses

To explore uncertainty, we introduce explicit stochasticity as a precursor to the study of uncertainty. We capture this randomness in part by an exogenous forcing process that evolves as:

$$dZ_t = \mu_z(Z_t)dt + \sigma_z(Z_t)dW_t$$

where $\{W_t : t \geq 0\}$ is a multivariate standard Brownian motion. We partition the vector Brownian motion into three subvectors as follows:

$$W_t = \begin{bmatrix} W_t^y \\ W_t^n \\ W_t^k \end{bmatrix}$$

where the first component consists of the climate change shocks, the second component contains damage shocks, and the third component captures the technology shocks. Consider an emissions “pulse” of the form:

$$(\iota_y \cdot Z_t) \mathcal{E}_t (\theta dt + \varsigma \cdot dW_t^y)$$

where $\mathcal{E}_t$ is fossil fuel emissions and $\iota_y \cdot Z = \{\iota_y \cdot Z_t : t \geq 0\}$ is a positive process which we normalize to have mean one. The $\iota_y \cdot Z$ -process captures “left out” components of the climate system’s reaction to an emission of $\mathcal{E}_t$ gigatons into the atmosphere while the $\varsigma \cdot dW_t^y$ process captures short time scale fluctuations. We will use a positive Feller square root process for the

⁶Ricke and Caldeira (2014) also consider separately two sources of temperature dynamics.

$\iota_y \cdot Z$ process in our analysis.

Within this framework, we impose the “Matthews’ approximation” by making the consequence of the pulse permanent:

$$dY_t = \mu_y(Z_t, \mathcal{E}_t)dt + \sigma_y(Z_t, \mathcal{E}_t)dW_t^y$$

where:

$$\begin{aligned}\mu_y(z, e) &= e(\iota_y \cdot z)\theta \\ \sigma_y(z, e) &= e(\iota_y \cdot z)\varsigma'\end{aligned}$$

Throughout, we will use uppercase letters to denote random vector or stochastic processes and lower case letters to denote possible realizations. Armed with this “Matthews’ approximation,” we collapse the climate change uncertainty into the cross-model empirical distribution reported in Figure 2. We will eventually introduce uncertainty about θ .

This specification misses the initial build up in the temperature response and instead focuses exclusively on the flat trajectories depicted in the upper panel of Figure 1. We expect that this error might be small when the prudent social planner embraces preferences that have a low rate of discounting the future, but this requires further investigation. While others in climate sciences find linear approximations to be relevant, we recognize the need for subsequent efforts to explore systematically the potential importance of nonlinearities. Ghil and Lucarini (2020) is a thorough review of climate physics at a hierarchy of temporal and spatial scales that embraces the inherent complexity of the climate system.

Remark 3.1. *For a more general starting point, let Y_t be a vector used to represent temperature dynamics where the temperature impact on damages is the first component of Y_t . This state vector evolves according to:*

$$dY_t = \Lambda Y_t dt + (\iota_y \cdot Z_t) \mathcal{E}_t (\Theta dt + \Phi dW_t^y)$$

where Λ is a square matrix and Θ is a column vector. Given an initial condition Y_0 , the solution for Y_t satisfies:

$$Y_t = \exp(t\Lambda)Y_0 + \int_0^t \exp[(t-u)\Lambda](\iota_y \cdot Z_u) \mathcal{E}_u (\Theta du + \Phi dW_u^y)$$

Thus under this specification, the expected future response of Y to a pulse at date zero is:

$$\exp(u\Lambda)\Theta.$$

It is the first component of this function that determines the response dynamics. This generalization allows for multiple exponentials to approximate the pulse responses. Our introduction of a multiple exponential approximation adapts for example, Joos et al. (2013) and Pierrehumbert

(2014).⁷

As an example, we capture the initial rise in the emission responses by the following two-dimensional specification

$$\begin{aligned} dY_t^1 &= Y_t^2 dt \\ dY_t^2 &= -\lambda Y_t^2 dt + \lambda \theta \mathcal{E}_t dt \end{aligned}$$

which implies the response to a pulse is:

$$\theta [1 - \exp(-\lambda t)] \mathcal{E}_0$$

A high value of λ implies more rapid convergence to the limiting response $\theta \mathcal{E}_0$. This approximation is intended as a simple representation of the dynamics where the second state variable can be thought of as an exponentially weighted average of current and past emissions.⁸

Remark 3.2. The approximation in Geoffroy et al. (2013) includes the logarithm of carbon in the atmosphere as argued for by Arrhenius (1896) which is not directly reflected in the linear approximation to the temperature dynamics that we use. The pulse experiments from Joos et al. (2013) show a more than proportional change in atmospheric carbon when the pulse size is changed. It turns out that this is enough to approximately offset the logarithmic Arrhenius adjustment so that the long-term temperature response remains approximately proportional for small pulse sizes. See also Pierrehumbert (2014) who discusses the approximate offsetting impacts of nonlinearity in temperature and climate dynamics.

4 Uncertain environmental and economic damages

Discussions of climate change policy are often simplified to specifications of temperature anomaly targets. For instance, many such discussions used a two degree anomaly as an upper bound on the amount of climate change should tolerate. More recently, this target number has been reduced to the point where numbers as low as a 1.5 degree anomaly should be entertained as a target upper bound with warnings of potentially severe consequences once warming exceeds these

⁷See equation (5) of Joos et al. (2013) and equations (1)-(3) of Pierrehumbert (2014). Pierrehumbert puts the change in radiative forcing equal to a constant times the logarithm of the ratio of atmospheric CO_2 at date t to atmospheric CO_2 at baseline date zero. His Figures 1 and 2 illustrate how an approximation of the Earth System dynamics by three exponentials plus a constant tracks a radiative forcing induced by a pulse into the atmosphere at a baseline date from the atmosphere works quite well with half lives of approximately six, sixty five, and four hundred and fifty years.

⁸In independent work, Dietz and Venmans (2019) and Barnett et al. (2020) have used such simplified approximations within an explicit economic optimization framework. The former contribution includes the initial rapid upswing in the impulse response functions. The latter contribution abstracts from this. Barnett et al. instead explore ways to confront uncertainty, broadly-conceived, while using the Matthews approximation.

thresholds. There is considerable debate as to the scientific underpinnings of such thresholds. Moreover, economic analyses have often introduced so called damage functions whereby economic opportunities are reduced by global warming depending on the curvature of the damage function. While damage functions are *ad hoc* simplifications that simplify the model solution and analysis, there remains considerable uncertainty as to their steepness.

For purposes of illustration, we introduce explicitly stochastic models of damages. The specification includes an unknown threshold whereby the curvature becomes apparent. In some of our computations, this threshold occurs somewhere between 1.5 and 2 degrees Celsius, but we also explore what happens when this interval is shifted to the right. Under a baseline specification, damage function curvature is realized in accordance with a Poisson event and an intensity that depends on the temperature anomaly. The event is more likely to be revealed in the near future when the temperature anomaly is larger. While we adopt a probabilistic formulation as a baseline, we will entertain ambiguity over damage function curvature and potential misspecification of the Poisson intensity. We intend our specification of the damage function to reflect that the value of future empiricism in the near term will be limited as the climate-economic system is pushed into uncharted territory. On the other hand, we allow for the damage function steepness to be revealed in the future as the climate system moves potentially closer to an environmental tipping point.

We posit a damage process, $N_t = \{N_t : t \geq 0\}$ to capture negative externalities on society imposed by carbon emissions. The reciprocal of damages, $\frac{1}{N_t}$, diminishes the productive capacity of the economy because of the impact of climate change. We follow much of climate economics literature by presuming that the process N reflects, in part, the outcome of a damage function evaluated at the temperature anomaly process. Importantly, we use a family of damage functions in place of a single function. Our construction of the alternative damage functions is similar to Barnett et al. with specifications motivated in part by prior contributions. Importantly, we modify their damage specifications in three ways:

- we entertain more damage functions, including ones that are more extreme;
- we allow for damage function steepness to emerge at an *ex ante* unknown temperature anomaly threshold;
- we presume that *ex post* this uncertainty is resolved;

We consider a specification under which there is a temperature anomaly threshold after which the damage function could be much more curved. This curvature in the “tail” of the damage function is only revealed to decision-makers when a Poisson event is triggered. As our model is highly stylized, the damages captured by the Poisson event are meant to capture more than just the economic consequences of a narrowly defined temperature movements. Temperature

changes are allowed to trigger other forms of climate change that in turn can spill over into the macroeconomy.

In our computational implementation, we use a piecewise log-quadratic function for mapping how temperature changes induced by emissions alter economic opportunities. The Poisson intensity governing the jump probability is an increasing function of the temperature anomaly. We specify it so that the Poisson event is triggered prior to the anomaly hitting an upper threshold $\bar{y}$. Construct a process

$$\bar{Y}_t = \begin{cases} Y_t & t < \tau \\ Y_t - Y_\tau + \bar{y} & t \geq \tau \end{cases}$$

where τ is the date of a Poisson event. Notice that $\bar{Y}_\tau = \bar{y}$. The damages are given by

$$\log N_t = \Gamma(\bar{Y}_t) + \iota_n \cdot Z_t \quad (1)$$

where:

$$\Gamma(y) = \gamma_1 y + \frac{\gamma_2}{2} y^2 + \frac{\gamma_3^m}{2} \mathbf{1}_{y \geq \bar{y}} (y - \bar{y})^2$$

and the only component of dW pertinent for the evolution of $\iota_n \cdot Z_t$ is dW_t^n . Decision-makers do not know when the Poisson event will be triggered nor do they know *ex ante* what the value of γ_3^m is prior to the realization of that event. At the time of the Poisson event, one of M values of γ_3^m is realized. In our application the coefficients γ_3^m are specified so that the proportional damages are equally spaced after the threshold $\bar{y}$.

The intensity function, $\mathcal{J}$, determines the possibility of a jump over the next small increment in time. For $Y_t = y$, $\epsilon \mathcal{J}(y)$ is the approximate jump probability over small time increment ϵ . Equivalently, $\mathcal{J}$ is a local measure of probability per unit of time. In our computations, we use intensity function

$$\mathcal{J}(y) = \begin{cases} r_1 \left(\exp \left[\frac{r_2}{2} (y - \underline{y})^2 \right] - 1 \right) & y \geq \underline{y} \\ 0 & 0 \leq y < \underline{y} \end{cases}$$

as depicted in Figure 3. As $\mathcal{J}$ is increasing in y , jumps become all the more likely as Y_t approaches the upper threshold $\bar{y}$. This intensity depends on $\underline{y}$, which we refer to as the lower threshold. We set the values of (r_1, r_2) so that the probability that the Poisson event is realized prior to $Y_t = \bar{y}$ is essentially unity. Thus, the uncertainty is concentrated for state Y in the interval $[\underline{y}, \bar{y}]$. We use the intensity plotted in Figure 3 in computations that follow for $\underline{y} = 1.5$ and $\bar{y} = 2$.




Figure 3: Intensity function, $r_1 = 1.5$ and $r_2 = 2.5$. With this intensity function, the probability of a jump at an anomaly of 1.6 is approximately .02 per annum, increasing to about .08 per annum at an anomaly of 1.7, increasing further to approximately .18 per annum at an anomaly of 1.8 and then to about one third per annum when the anomaly is 1.9.

The probability that the process has not jumped over the time interval $[0, t)$ is:

$$\exp \left[- \int_0^t \mathcal{J}(Y_\tau) d\tau \right]$$

so that a larger intensity makes the jumps more likely.

Figure 4 shows the range of damage function uncertainty that we impose in our computations. Given our intensity specification, we expect the Poisson jumps to occur between 1.5 and 2 degrees celsius. The upper panel shows the potential damage function quantiles when the jump is delayed until a two degree temperature anomaly. These functions are all continuous extensions of initial damage functions beyond $\bar{y}$. When a jump occurs at anomalies less than two, the damage functions are steeper. This is illustrated in the lower panel, which shows the possible damage function quantiles for temperature anomalies beyond the anomaly just prior to the jump. Earlier jump dates imply steeper damage functions and are thus a form of “bad news.” In contrast, delayed jumps are “good news.”

Remark 4.1. *Our choice of $\underline{y} = 1.5$ and $\bar{y} = 2$ degree celsius thresholds for the temperature anomaly is motivated by discussions in the climate science literature. Thus these damage functions are more than economic responses to changes in temperature but also inclusive of potentially dramatic environmental changes triggered by so called “tipping points.” Drijfhout et al. (2015) provided a catalog of potential abrupt changes implied by projections from a variety of earth science models. Motivated by such concerns about such changes, Rogelj et al. (2018) and Rogelj et al. (2019) suggested a 1.5 degree temperature anomaly as a goal for limiting human damages of climate system. As this goal may be unachievable, they point to a 2 degree target in line with*




Figure 4: Range of possible damage functions for different jump thresholds. The shaded regions in these plots gives the range of possible values for $\exp(-n)$, which measures the proportional reduction of the productive capacity of the economy. The top panel shows the damage function curvature when the jump occurs at $Y_t = 2.0$ and the bottom panel shows the damage function curvature if that jump happens at $Y_t = 1.5$.

the 2015 Paris agreement, while noting the increased danger of severe damages. While there are concerns about temperature anomalies triggering these tipping points, Ritchie et al. (2021) used recent developments in dynamical systems theory to argue that “the point of no return” for climate thresholds is highly uncertain. As an external form of sensitivity analysis, we also report results with the temperature threshold shifted to the right with $\underline{y} = 1.75$ and $\bar{y} = 2.25$.

5 Model ambiguity and misspecification concerns

The model we have built so far is one in which uncertainty is captured by the stochastic specification of shocks as is typical when building dynamic stochastic models in macroeconomics. We think of this shock specification as characterizing *risk*. The presence of these shocks opens the door to a comprehensive assessment of uncertainty in which we entertain a broader notion of uncertainty. We include uncertainty over model specifications and parameters, which we refer to as ambiguity. In this discussion, we treat models and parameters as synonymous by thinking of each parameter as indexing an alternative model. As in our prior work, Barnett et al. (2020) and Berger and Marinacci (2020), we are led to depart from the Bayesian approach which starts with the specification of a subjective prior over the alternative models, but does not distinguish the role of subjective probabilities over models from the probabilities given a model. Instead, we use recent formalisms from decision theory under uncertainty to explore the impact of uncertainty over the subjective inputs. Within statistics, this gave rise to robust counterparts to Bayesian inferences in the study of prior sensitivity, often outside the realm of a specific decision problem. The decision theory framework formalizes the question of “sensitivity to what?” and the formal tradeoff between making best guesses versus possible bad outcomes as we look across models. While in some settings, data richness may diminish the role of prior sensitivity, we find the economics of climate change to be a problem whereby prior sensitivity remains an important question for the decision-maker. Of course, any model we write down is necessarily a simplification. We also incorporate concerns about the potential *misspecification* of the models under exploration using ideas from robust control theory extended to dynamic economic models.

Our use of decision theory gives rise to a form of uncertainty quantification. Uncertainty quantification in the sciences is typically done by researchers. For instance, we might ask how the social cost of carbon differs as we change the modeling ingredients. But decision-makers also confront this uncertainty, including ones inside the models that we build. Thus, model ambiguity or misspecification concerns by decisions makers should arguably be taken into account when determining the prudent course of action. This same uncertainty emerges as adjustments to the social cost of carbon as set by say a benevolent social planner. Just like risk aversion can induce caution in decision making, the same can be said of broader notions of uncertainty aversion. While the decision theory that we use does not determine the magnitude of what this aversion should

be, it reduces a potentially high dimensional sensitivity analysis to a much lower dimensional one captured by low dimensional representations of uncertainty aversion.

We analyze this uncertainty using the formalism of decision theory under uncertainty. We apply two versions of such theory, one comes under the heading of variational preferences and the other under smooth ambiguity preferences. We adapt both to continuous-time specifications, which facilitates their implementation and interpretation. We use this decision theory to reduce the sensitivity analysis to a one or two-dimensional parameterization that locates the potential misspecification that is most consequential to a decision-maker. Our aim is to provide a more complete uncertainty quantification within the setting of decision problems.

5.1 State dynamics

Posing our model in continuous-time leads to a simplified characterization of robustness. When constructing dynamic programming solutions, it is advantageous to represent the local dynamics

$$\begin{aligned} \lim_{\epsilon \downarrow 0} \frac{1}{\epsilon} \mathbb{E}[V(X_{t+\epsilon}) - V(X_t) \mid X_t = x, a] = & \mu(x, a) \cdot \frac{\partial V}{\partial x}(x) \mu(x, a) + \frac{1}{2} \text{trace} \left[\frac{\partial^2 V}{\partial x \partial x'}(x) \Sigma(x, a)' \right] \\ & + \mathcal{I}(x) \int [V(\tilde{x}) - V(x)] Q(d\tilde{x} \mid x) \end{aligned} \quad (2)$$

where Σ is the diffusion matrix:

$$\Sigma(x, a) = \sigma(x, a) \sigma(x, a)'$$

and local exposure of X_t to a multivariate standard Brownian motion is $\sigma(X_t, a) dW_t$, $\mathcal{I}$ is the state dependent jump intensity, and $Q(d\tilde{x} \mid x)$ is the jump distribution conditioned on the process jumping. Both the intensity and the jump distribution can depend on the current Markov state. The first two contributions to the drift of $V(X_t)$ are familiar from Ito's Lemma. We allow for both μ and σ to depend on a current action a . The baseline probabilities are the ones implied by this stochastic process specification.

Given our interest in recursive methods, in what follows, we will describe in turn implications for mis-specifying a Brownian motion, a jump process and an ambiguity adjustment for the local mean of the dynamical system.⁹ We introduce robustness by following the approach described in Anderson et al. (2003) and model ambiguity by following the approach in Hansen and Miao (2018).

⁹See Hansen and Sargent (2020b) and Cerreia-Vioglio et al. (2021) for decision theoretic discussions of the distinct roles for model ambiguity and misspecification concerns.

5.2 Misspecified Brownian increments

Following James (1992), Hansen and Sargent (2001) and others, the potential misspecification of a Brownian motion has a particularly simple form. It is known from the famed Girsanov Theorem that a change in distribution represented by a likelihood ratio replaces the standard Brownian motion by a Brownian motion with a drift. Specifically, under such a change in probability distribution, dW_t is changed from a Brownian increment to a Brownian increment with a drift or local mean that can be state (or model) dependent, which we denote $H_t dt$. Thus, to explore the consequences of misspecification, we modify our (locally) normally distributed shocks by entertaining possible mean distortions. Allowing for arbitrary changes in the drift without a constraint or a penalty, leads to an uninteresting and inflexible decision problem. Here we follow one of the preference specifications in Hansen and Sargent whereby we use an expected log-likelihood ratio measure of discrepancy called relative entropy to restrain the search over alternative possible drift specifications.¹⁰ For Brownian motion models, the relative entropy penalty is $\frac{\xi_r}{2}|H_t|^2 dt$ where ξ_r is the penalty parameter that governs the decision-maker's concern for misspecification and $\frac{1}{2}|H_t|^2 dt$ is the local contribution to relative entropy.

Given our intention to use recursive methods, we characterize the impact that the Brownian motion has for the state dynamics in terms of the local means for functions of the Markov state. We allow for the probabilities implied by a multivariate standard Brownian motion to be modified to include possible drift distortions. Equivalently, we introduce potential model misspecification by replacing μ by

$$\min_h \frac{\partial V}{\partial x} \cdot (\mu + \sigma h) + \frac{\xi_r}{2} h' h = \frac{\partial V}{\partial x} \cdot \mu - \frac{1}{2\xi_r} \left(\frac{\partial V}{\partial x} \right)' \Sigma \left(\frac{\partial V}{\partial x} \right) \quad (3)$$

where the minimizing h is:

$$h^* = -\frac{1}{\xi_r} \sigma' \frac{\partial V}{\partial x}$$

h is a drift distortion in the Brownian increment. This distortion is state dependent as reflected by the partial derivative of the value function. The term $\frac{1}{2} h' h$ is the local measure of relative entropy implied the local normality of Brownian motion. Large values of the penalty parameter ξ_r induce small robustness adjustments.

5.3 Misspecified jump components

To specify a Markov jump process requires both a) a state-dependent intensity, $\mathcal{I}$, governing the probability of a jump and b) the distribution $Q(\cdot | x)$ over the post-jump state given the current jump state, x . Both of these could be mistaken. We capture potential misspecification

¹⁰This approach is a continuous-time version of the dynamic variational preferences of Maccheroni et al. (2006).

by introducing positive functions g of the post-jump state

$$\mathcal{I}(x) \int [V(\tilde{x}) - V(x)] g(\tilde{x}) Q(d\tilde{x} | x)$$

The implied probabilities resulting from the concern about misspecification are given by

$$\frac{1}{\bar{g}(x)} g(\tilde{x}) Q(d\tilde{x} | x)$$

and the implied intensity is given by $\bar{g}(x)\mathcal{I}(x)$ where

$$\bar{g}(x) \doteq \int g(\tilde{x}) Q(d\tilde{x} | x).$$

Thus the choice of g determines both the jump probabilities and the jump intensities in a manner that is mathematically tractable. The local relative entropy discrepancy for a jump process is:

$$\mathcal{I}(x) \int [1 - g(\tilde{x}) + g(\tilde{x}) \log g(\tilde{x})] Q(d\tilde{x} | x)$$

This measure is nonnegative because the convex function $g \log g$ exceeds its gradient $g-1$ evaluated at $g = 1$.

We make a robustness adjustment by solving:

$$\min_g \mathcal{I}(x) \int [V(\tilde{x}) - V(x)] g(\tilde{x}) Q(d\tilde{x} | x) + \xi_r \mathcal{I}(x) \int [1 - g(\tilde{x}) + g(\tilde{x}) \log g(\tilde{x})] Q(d\tilde{x} | x)$$

The minimizing g is given by

$$g^*(\tilde{x}, x) = \exp \left(-\frac{1}{\xi_r} [V(\tilde{x}) - V(x)] \right)$$

with a minimized objective given by

$$\xi_r \mathcal{I}(x) \int \left[1 - \exp \left(-\frac{1}{\xi_r} [V(\tilde{x}) - V(x)] \right) \right] Q(d\tilde{x} | x).$$

We use this outcome in place of the jump contribution

$$\int [V(\tilde{x}) - V(x)] Q(d\tilde{x} | x)$$

in the local mean contribution for a continuous time Hamilton-Jacobi-Bellman (HJB) equation.

5.4 Structured ambiguity

To assess the consequences of the heterogeneous responses from alternative climate models, we use what are called recursive smooth ambiguity preferences proposed by Klibanoff et al. (2009). For an important special case of these preferences, Hansen and Sargent (2007) provide a robust prior/posterior interpretation of these preferences. This alternative interpretation has advantages both in terms of calibration and representation of social valuation which we exploit in our characterization. In deploying such preferences, we use a robust prior interpretation in conjunction with the continuous-time formulation of smooth ambiguity proposed by Hansen and Miao (2018).

To assess the consequences of the heterogeneous responses from alternative climate models, we follow Hansen and Miao (2018). Suppose that $\mu(x, a | \theta)$ where θ is an unknown parameter. In our applications, θ is the climate sensitivity parameter pertinent for the Matthew's approximation. Let π denote a subjective probability distribution over θ . For instance, this could be the cross model distribution coming from the alternative pulse experiments reported in Figure 1. To justify this distribution as a meaningful distribution from a decision-theoretic perspective requires that all model outputs be treated as equally plausible. In Bayesian parlance, π would be a current period posterior dependent on an initial prior and likelihood. The decision-maker does not have in the full posterior and engages in a robust adjustment by solving

$$\min_{f, \int f(\theta)\pi(d\theta)=1} \frac{\partial V}{\partial x}(x) \cdot \int \mu(x, a | \theta) f(\theta) \pi(d\theta) + \xi_a \int f(\theta) \log f(\theta) \pi(d\theta) \quad (4)$$

where $f(\theta)$ is a density relative to a baseline probability π . The minimizing f is

$$f^*(\theta | x, a) \propto \exp \left[- \left(\frac{1}{\xi_a} \right) \frac{\partial V}{\partial x}(x) \cdot \mu(x, a | \theta) \right]$$

where the right side requires scaling to be a proper relative density.¹¹ The minimizing objective is

$$-\xi_a \log \left(\int \exp \left[- \left(\frac{1}{\xi_a} \right) \frac{\partial V}{\partial x}(x) \cdot \mu(x, a | \theta) \right] \pi(d\theta) \right)$$

which replaces

$$\frac{\partial V}{\partial x}(x) \cdot \int \mu(x, a | \theta) \pi(d\theta)$$

in an HJB equation.

5.5 A valuation adjustment for uncertainty

There is much discussion in the literature on environmental economics about what discount rate to use. In our analysis so far, there is a single discount rate used to define the preferences of a

¹¹In particular, the right-hand side needs to integrate to one over θ .

fictitious social planner. But the discussions in the literature usually refer to present discounted value formulas for marginal valuation. We represented the robust adjustments in terms of altered probabilities which we compute in conjunction with the HJB equations used for optimization. As Barnett et al. (2020) demonstrate, these same probabilities provide the uncertainty adjustments for social valuation. Thus, to account for uncertainty, broadly conceived, we are pushed beyond the question of what discount rate to use because the necessary adjustment is most conveniently depicted as an altered probability measure.

5.6 Other approaches to uncertainty quantification across models

We briefly discuss three prior forms of uncertainty quantification as it pertains to unknown parameters or models. We give these as illustrations, but the list is by no means exhaustive.

Olson et al. (2012) propose and implement a Bayesian method for making inferences about certain parameters of interest, including a climate sensitivity parameter coming from the UVic (University of Victoria) earth system climate model. They document posterior sensitivity of the climate sensitivity parameter to priors and other unknown modeling inputs. In particular, they show the need to use an informative prior for climate sensitivity to obtain reasonable results, therefore demonstrating the posterior uncertainty in their informative statistical investigation. While not the focal point of their analysis, there is additional uncertainty in the likelihood construction. These forms of uncertainty are pertinent not only to researchers presenting evidence, but also to decision or policy makers in their efforts. Thus, we move the uncertainty quantification “inside the decision problem,” including the sensitivity analysis. This allows us to explore the impact of model or parameter ambiguity for choosing socially prudent emissions trajectories and imputing the implied social cost of carbon.

In an alternative investigation of uncertainty within and across climate-economic models, Gillingham et al. (2018) and Nordhaus (2018) computed distributions of model outcomes given *a priori* distributions of parameters, specifications and model inputs, including emissions pathways. From their analysis, they are able to produce a set of outputs associated with each parameter or model configuration to demonstrate the role of uncertainty in their setting. Their static analysis occurs “outside the decision problem,” but it opens the door to exploring changes in the prior probability distribution without a systematic analysis of the sensitivity. Our framework uses recursive methods and decision theoretic tools to determine endogenously prudent choices of emissions over time and the implied social cost of carbon trajectories when the policy maker confronts prior ambiguity. Policy outcomes include endogenous feedbacks and dynamic impacts on the social cost of carbon, and, importantly, an adjustment for uncertainty that is either unresolved, or only resolved well into the future.¹²

¹²See Lemoine and Rudik (2017) who provided a related commentary, arguing why recursive methods from economic dynamic can open the door to important extensions in climate economics including parameter learning.

In a third approach, Hassler et al. (2018) conducted an analysis of uncertainty by comparing policy outcomes across two parameter intervals, one pertaining to damages and another to climate sensitivity. Instead of putting a probability distribution over parameters, they evaluate policy outcomes at the extreme points of the parameter space. Their analysis can be thought of as a simple illustration of robust decision making allowing for arbitrary probabilities over the unknown parameters and is a revealing starting point to the policy problem they investigate. Our analysis of the policy problems is explicitly dynamic and imposes probabilistic restraints on the probabilities that could be assigned over a potentially large set of alternative model configurations. The dynamic decision theory formulation we use collapses our resulting sensitivity analysis to a low dimensional representation in terms of ambiguity and misspecification aversion parameters.

6 Social cost of carbon

The social cost of carbon (SCC) is intended to be the expected discounted value of future environmental and economic damages induced by carbon emissions during the current period. In the special case in which socially efficient allocations are used, the wedge between the discounted social costs and market prices determines the Pigouvian taxes that support the efficient allocation in the presence of the environmental externality provided that the socially efficient emissions are positive.¹³ More generally, the SCC is a decidedly local measure applicable to marginal changes in policy. Since the SCC can be represented as an expected discounted value, we draw on insights from asset pricing to explore the implications of uncertainty.

To place our research in a broader context of environmental economics, we draw on a recent report from the National Academies of Sciences, Engineering and Medicine (2017), which featured a four-step modular approach to measuring the SCC. We quote from the executive summary:

The committee specifies criteria for future updates to the $SC - CO_2$. It also recommends an integrated modular approach for $SC - CO_2$ estimation to better satisfy the specified criteria and to draw more readily on expertise from the wide range of scientific disciplines relevant to $SC - CO_2$ estimation. Under this approach, each step in $SC - CO_2$ estimation is developed as a module – socioeconomic, climate, damages, and discounting – that reflects the state of scientific knowledge in the current, peer-reviewed literature.¹⁴

Nordhaus (2018) noted the inability of his framework to address such endogenous feedbacks and unresolved uncertainty, and also pointed out the potential value to using the type of recursive methods we employ in our analysis as a way to address such issues.

¹³As Cai and Lontzek (2019) noted and encountered in some of their simulations, when emissions hit a zero constraint, the SCC may reflect a desire for negative emissions while Pigouvian taxes are needed only to reach the zero emissions outcome.

¹⁴We use units of carbon as opposed to CO_2 in our computation, which is in effect a different choice of units.

The report goes on to argue:

In addition, the committee details longer-term research that could improve each module and incorporate interactions within and feedbacks across modules.

The modularization has the advantage of compartmentalizing different ingredients to the SCC estimation. Perhaps this is a valuable shortcut, but as we will argue, it misses or disguises contributions to valuation that we take to be central. This is especially true once we seek to formally integrate uncertainty into the analysis. In what follows, we will use recursive methods from control theory and economic dynamics to expose some of the considerations missed or treated inconsistently by the modular approach.

6.1 Impacts of marginal changes in emissions

For pedagogical (and computational) simplicity, we consider a highly stylized problem in order to pose some important conceptual challenges. We suppose a single policy maker with a dynamic stochastic objective, but, importantly, we will not suppose that the emissions are set in a socially optimal way. We suppose that the control A_t includes emissions $\mathcal{E}_t$ as its last entry and that this control can be expressed as an invariant function of the state vector X_t . The invariance is imposed for simplicity, but the dependence on the state vector will be important in what follows. Even if we are dubious about whether the emissions will approximate a socially efficient outcome, we wish to allow for some policy responses whereby future emissions depend on, say, the magnitude of the temperature anomaly or on economic damages that emerge in the future. We allow both damages and the temperature anomaly to be components of the state vector X_t . An efficient social planner would certainly do this, but the feedback of emissions onto endogenous state variables can occur much more generally. A more complete analysis would impose a specific market structure and impose explicit constraints on the policy maker recognizing the implied market responses. The plausible assessments of exogenous scenarios arguably have such considerations in the background, but as we will see their formal presence can have an important impact on the construction of the SCC.

Consider the following recursive formulation of social valuation posed in the absence of model ambiguity or misspecification aversion:

$$0 = -\delta V(x) + U(x, a) + \mu(x, a) \cdot \frac{\partial V}{\partial x}(x) + \frac{1}{2} \text{trace} \left[\frac{\partial^2 V}{\partial x \partial x'}(x) \Sigma(x, a)' \right] + \mathcal{I}(x) \int [V(\tilde{x}) - V(x)] Q(d\tilde{x} | x) \quad (5)$$

where δ is the subjective rate of discount used by a fictitious social planner, the control is defined as $a = \psi(x)$, and U is the instantaneous contribution to the planner's continuation value function

V. This relation is a version of a Feynman-Kac equation with a jump component. Intuitively, it says that the instantaneous utility contribution, U , should exactly offset the local mean of the discounted continuation value process.

6.2 A discrete-time representation

For a measure of the net benefits (or minus net costs), we deduce the marginal impact of an additional unit of emissions into the atmosphere. Although we will derive continuous-time formulas, we start with a discrete time approximation with ϵ as the gap between time points. Heuristically, we aim to evaluate the ratio of

$$MV_t = \epsilon \frac{d}{d\mathcal{E}_t} \mathbb{E} \left[\sum_{j=0}^{\infty} \exp(-\delta j \epsilon) U(X_{t+j\epsilon}, A_{t+j\epsilon}) \mid \mathfrak{F}_t \right] \quad (6)$$

to the marginal utility of a numeraire consumption good at a date $t = \bar{j}\epsilon$ for some integer $\bar{j}$. We include the ϵ on the right side of (6.2) to view this expression as a Reimann sum approximation of an integral over time. Express

$$\begin{aligned} MV_t &= \epsilon \frac{d}{d\mathcal{E}_t} \mathbb{E} \left[\sum_{j=0}^{\infty} \exp(-\delta j \epsilon) U(X_{t+j\epsilon}, A_{t+j\epsilon}) \mid \mathfrak{F}_t \right] \\ &= \epsilon \mathbb{E} \left(\sum_{j=0}^{\infty} \exp(-\delta j \epsilon) \left[\frac{\partial U}{\partial a}(X_{t+j\epsilon}, A_{t+j\epsilon}) \frac{\partial \psi}{\partial x}(X_{t+j\epsilon}) + \frac{\partial U}{\partial x}(X_{t+j\epsilon}, A_{t+j\epsilon}) \right] R_{t+j\epsilon, t} \mid \mathfrak{F}_t \right) \end{aligned}$$

and

$$R_{t,t+j\epsilon} = \frac{dX_{t+j\epsilon}}{d\mathcal{E}_t}$$

is the random response of the future state vector $X_{t+j\epsilon}$ to an emissions pulse $\mathcal{E}_t$ at date t . This discounted expected value is expressed in utility units and becomes the social cost of carbon once we divide by the marginal utility of the consumption numeraire. For linear dynamics, as is well-known from impulse response theory, $R_{t,t+j\epsilon}, j = 0, 1, 2, \dots$ will not be random. More generally, there are well known methods for characterization of nonlinear impulse responses for diffusions.¹⁵

6.3 Two forms of impulse responses

As we have seen, impulse responses are inputs into the SCC computations. We consider two alternative formulations of these responses, including one that is common in scenario analysis and another that is familiar in dynamic stochastic equilibrium theory.

¹⁵See, for instance, Borovička et al. (2014) for a pedagogical treatment of nonlinear impulse response functions for diffusions and related computations pertinent for valuation. The calculations relate closely to two well-known mathematical tools, the method of characteristics and Malliavin differentiation.

It is common to run “scenarios” through climate-economic models. Prominent examples in climate science are the Representative Concentration Pathway (RCP) scenarios which are typically specified as exogenous paths of atmospheric carbon over time. See, for example, Zickfeld et al. (2013) (Figure 1), for the four main RCP’s for carbon concentration. In generating these and other scenarios, the emissions or the atmospheric carbon trajectories are treated as an exogenous input and not as something that is determined endogenously by the model. Such scenarios are helpful in understanding temperature dynamics and making cross-model comparisons without attempting to complete or close the dynamical system. We call these functions “scenario response functions” (SceRF’s).

In contrast, for a completely specified dynamical system, emissions and atmospheric carbon trajectories are determined endogenously. As these variables may feedback onto the state of the dynamical system, there are alternative impulse responses that take into account this endogeneity. We call these “system response functions” (SysRF’s). These impulse responses are also the ones pertinent to represent the SCC when the full climate-economic model has a recursive representation of emissions as a function of the Markov state. These latter impulses are commonly used when depicting the implications of dynamic stochastic equilibrium models. In what follows, we will represent both forms of impulse responses.

For simplicity, we will abstract from the jump components, although there are direct extensions to the more general case in which they are included in the analysis.

6.3.1 Exogenously specified scenarios

Consider first a SceRF. We suppose a block recursive structure for the dynamics:

$$\begin{aligned} dX_t &= \mu(X_t, A_t)dt + \sigma(X_t, A_t)dW_t \\ d\bar{X}_t &= \bar{\mu}(\bar{X}_t)dt + \bar{\sigma}(\bar{X}_t)dW_t \end{aligned} \tag{7}$$

and

$$A_t = \bar{\psi}(\bar{X}_t)$$

where X_t includes states used to represent climate dynamics and environmental or economic damages. We rewrite the first equation block in this system as:

$$dX_t = \mu[X_t, \bar{\psi}(\bar{X}_t)]dt + \sigma[X_t, \bar{\psi}(\bar{X}_t)]dW_t$$

to represent the Markov dynamics in the composite state vector $(X_t, \bar{X}_t)$. Under this specification, the response $\bar{X}_{t+\tau}$ to an emissions pulse is zero.

In this setting, alternative scenarios are alternative specifications of $(\bar{\mu}, \bar{\sigma}, \bar{\psi})$. We construct the response functions by introducing emissions “pulses” at an initial time period.

6.3.2 Endogenous emissions

Suppose that the full state dynamics are captured by:

$$dX_t = \mu(X_t, A_t)dt + \sigma(X_t, A_t)dW_t$$

where

$$A_t = \psi(X_t).$$

A special case of this is when ψ is computed as the solution to a fictitious social planner's problem. We represent the state dynamics as:

$$dX_t = \mu[X_t, \psi(X_t)]dt + \sigma[X_t, \psi(X_t)]dW_t$$

This specification allows for emissions to depend on the Markov state vector, including endogenous states that capture temperature anomalies or other forms of climate change and economic damages. A marginal change in emissions at date t induces marginal changes in emissions in future time periods because of some form of policy or market response. Suppose, for instance, that economic policies that influence carbon emissions adapt to temperature anomalies in the future. This will consequently alter the marginal impact of emissions today on future temperature anomalies in ways that are missed by treating emissions as fully exogenous. For instance, when we study a social planner's problem, emissions feed back on the temperature anomaly and mute the response of the temperature anomaly relative to the flat response featured in the Matthew's approximation.

This endogeneity illustrates the “interactions within and feedbacks across modules” mentioned in the National Academies of Sciences, Engineering and Medicine (2017) report that should be embraced in a “longer-term research” agenda.

6.3.3 Hybrid approach

A third approach is a hybrid of the first two. It borrows a well-known idea from macroeconomic dynamics sometimes referred to as “big K little k.” Consider a pre-specified ψ giving the feedback of $a = \psi(x)$. Suppose that $\bar{X}_t = X_t$ as an equilibrium outcome. Write the dynamical system as in (7) where

$$\begin{aligned}\bar{\mu}(x) &= \mu[x, \psi(x)] \\ \bar{\sigma}(x) &= \sigma[x, \psi(x)]\end{aligned}$$

We impose the initialization: $X_0 = \bar{X}_0$. A hypothetically small emissions pulse alters marginally the trajectory of $X_{t+\tau}$ but not that of $\bar{X}_{t+\tau}$. While this approach imposes internal consistency

along an equilibrium path, this consistency is not imposed in the perturbation leading to a different computation of the SCC. On the other hand, the exogenous scenario is constructed in a manner that is consistent with an endogenous policy response. As part of the equilibrium, this approach takes account of future market or policy responses when assessing plausible scenarios in a manner that is internally consistent while preserving the exogenous perspective of an emissions scenario.

6.4 Markov solutions

We explore Markov solutions for the second two approaches as the first one, in an unconstrained manner, seems hard to defend from a dynamical systems perspective. For the marginal change in emissions, we are led to differentiate the right-hand side with respect to emissions. e , and to evaluate the derivative at $e = \psi(y)$:

$$MV(x) = \frac{\partial U}{\partial e}[x, \psi(x)] + \frac{\partial V}{\partial x}(x) \cdot \frac{\partial \mu}{\partial e}[x, \psi(x)] + \frac{1}{2} \text{trace} \left[\frac{\partial^2 V}{\partial x \partial x'}(x) \frac{\partial}{\partial e} \Sigma[x, \psi(x)] \right]. \quad (8)$$

where V solves equation (5). Recall that e is the last entry of the current period action a . Given our continuous-time specification, $MV(X_t)$ is a per unit time measure of the marginal impact of a change in emissions. Under a social planner's solution $MV(x) = 0$, unless the emissions solution is at a boundary.

This same logic can be modified when we proceed under the prospective of an exogenous emissions scenario. Now, the value function $\bar{V}$ depends on both x and $\bar{x}$ and satisfies a Feynman-Kac equation that is the analog to (5). A marginal change in e is presumed to only impact x . The formula for the net benefit in the current state is now:

$$\begin{aligned} \overline{MV}(x, \bar{x}) = & \frac{\partial U}{\partial e}[x, \bar{\psi}(\bar{x})] + \frac{\partial \bar{V}}{\partial x}(x, \bar{x}) \cdot \frac{\partial \mu}{\partial e}[x, \bar{\psi}(\bar{x})] \\ & + \frac{1}{2} \text{trace} \left(\left[\begin{array}{cc} \frac{\partial^2 \bar{V}}{\partial x \partial x'} & \frac{\partial^2 \bar{V}}{\partial x \partial \bar{x}'} \\ \frac{\partial^2 \bar{V}}{\partial \bar{x} \partial x'} & \frac{\partial^2 \bar{V}}{\partial \bar{x} \partial \bar{x}'} \end{array} \right] (x, \bar{x}) \frac{\partial}{\partial e} \bar{\Sigma}[x, \bar{x}, \bar{\psi}(\bar{x})] \right). \end{aligned}$$

where $a = \bar{\psi}(\bar{x})$ and

$$\bar{\Sigma}(x, \bar{x}, a) \doteq \begin{bmatrix} \sigma(x, a) \\ \bar{\sigma}(\bar{x}) \end{bmatrix} \begin{bmatrix} \sigma(x, a)' & \bar{\sigma}(\bar{x})' \end{bmatrix}$$

While under the hybrid approach $X_t = \bar{X}_t$, it will typically be the case that:

$$\overline{MV}(x, x) \neq MV(x).$$

This follows because of the differences in the implied responses to emissions pulses.

While MV or $\overline{MV}$ are the net benefits (benefits minus costs), the separation of the two components depends on the details of the planner's problem, as we will illustrate in some example

economies.

6.5 Uncertainty adjustment

To accommodate ambiguity and misspecification aversion, we introduce minimization into the construction of the value function using the approaches described in Section 5. The outcome of this minimization gives an uncertainty adjustment for valuation, which is conveniently represented as a probability distribution. This probabilistic adjustment emerges because of our choice to include uncertainty “inside” the decision problem. Just as external researchers confront uncertainty, so do policymakers. The discounting module in the National Academies of Sciences, Engineering and Medicine (2017) misses this uncertainty adjustment. Since the outcome of the minimization depends on uncertainty in both the climate dynamics and the damage function specification, this identifies important linkages between the different modules that are missed by the simplistic discount rate sensitivity analyses often featured in the environmental economics literature. While the choice of δ is important, so is the *decision-maker’s aversion to ambiguity and potential model misspecification*.

Remark 6.1. *It is well-known in the macroeconomic literature when using aggregate consumption as a numeraire for the purposes of marginal valuation that the subjective rate of discount is augmented with a growth adjustment that depends in part on the elasticity of intertemporal substitution. Relatedly, in the macroeconomic-asset pricing literature, investor risk aversion alters discounting by making it stochastic. The discounting depends on the exposure of the cash flow being discounted to aggregate risk. Cai et al. (2017) include both of these considerations in their treatment of the SCC. Our calculations extend these insights by constructing probabilistic adjustments for model ambiguity for social valuation.*

Next, we show how to make the uncertainty adjustment for valuation that incorporates robustness to model ambiguity and potential misspecification. We modify the construction of MV , but the adjustment to $\overline{MV}$ is entirely analogous. To incorporate uncertainty, we modify the Feynman Kac equation (5) by incorporating a minimization problem. Formally, the equation (5) now becomes a Hamilton-Jacobi-Bellman equation pertinent for continuous-time discounted

dynamic programming solutions with minimization:

$$\begin{aligned}
0 = & \min_{f, \int f(\theta)\pi(d\theta)=1} \min_{g(\tilde{x})} \min_h -\delta V(x) + U(x, a) + \frac{\partial V}{\partial x}(x) \cdot [\sigma(x, a)h] + \frac{1}{2} \text{trace} \left[\frac{\partial^2 V}{\partial x \partial x'}(x) \Sigma(x, a) \right] \\
& + \frac{\partial V}{\partial x}(x) \cdot \int \mu(x, a \mid \theta) f(\theta) \pi(d\theta) + \xi_a \int f(\theta) \log f(\theta) \pi(d\theta) \\
& + \mathcal{I}(x) \int [V(\tilde{x}) - V(x)] g(\tilde{x}) Q(d\tilde{x} \mid x) \\
& + \xi_r \left[\frac{1}{2} h' h + \mathcal{I}(x) \int [1 - g(\tilde{x}) + g(\tilde{x}) \log g(\tilde{x})] Q(d\tilde{x} \mid x) \right]
\end{aligned}$$

where $a = \psi(x)$. We then use this value function solution in equation (8) to construct an uncertainty adjusted version of MV .¹⁶

We next consider three example economies: the first features damage function uncertainty and its resolution; the second features a novel uncertainty decomposition that incorporates robustness to model ambiguity and misspecification; the third investigates the impact of uncertain advances in the availability of less carbon-intensive technologies. While our methods for making uncertainty adjustments for the SCC are more generally applicable, in all three cases, we feature emission trajectories that are outcomes of social planning problems. While highly stylized, each of the three examples illustrates novel impacts of uncertainty.

7 Uncertain damages

Our first example features how the perceived unraveling of uncertainty about economic and environmental damages influences prudent decisions.

7.1 Capital evolution

We consider an AK technology for which output is proportional to capital and can be allocated between investment and consumption. Capital in this specification should be broadly conceived to include human capital or intangible capital. Suppose that there are adjustment costs to capital that are represented as the product of capital times a quadratic function of the investment-capital ratio.

As a modeling construct, we first consider “undamaged” counterparts to consumption, output and capital. Abstracting from damages, capital evolves as:

$$dK_t = K_t \left[\mu_k(Z_t) dt + \left(\frac{I_t}{K_t} \right) dt - \frac{\kappa}{2} \left(\frac{I_t}{K_t} \right)^2 dt + \sigma_k(Z_t) dW_t^k \right],$$

¹⁶Including parameter learning requires additional state variables that serve as sufficient statistics for the unknown parameter vector θ under the base probability specification.

where K_t is the capital stock and I_t is investment. The capital evolution expressed in logarithms is:

$$d \log K_t = \left[\mu_k(Z_t) + \left(\frac{I_t}{K_t} \right) - \frac{\kappa}{2} \left(\frac{I_t}{K_t} \right)^2 \right] dt - \frac{|\sigma_k(Z_t)|^2}{2} dt + \sigma_k(Z_t) dW_t^k.$$

We let consumption, C_t , and investment, I_t , sum up to output, which is proportional to capital:

$$C_t + I_t = \alpha K_t.$$

Next, we consider environmental damages. We suppose that the damage process that we depicted previously shifts proportionately consumption and capital by a multiplicative factor. For instance, the damage adjusted consumption is $\tilde{C}_t = \frac{C_t}{N_t}$, and the damage adjusted capital is $\tilde{K}_t = \frac{K_t}{N_t}$.

Without uncertainty aversion, preferences for the planner are time-separable with a unitary elasticity of substitution. The planner's instantaneous utility from “damaged consumption” and emissions is given by:

$$\begin{aligned} & (1 - \eta) \log \tilde{C}_t + \eta \log \mathcal{E}_t \\ &= (1 - \eta)(\log C_t - \log N_t) + \eta(\log K_t - \log N_t) + \eta \log \mathcal{E}_t \end{aligned}$$

where we will denote the subjective rate of discount used in preferences as δ . We can think of emissions and consumption as distinct goods, or we can think of $\tilde{C}_t$ as an intermediate good that when combined with emissions determines final consumption.

Given this formulation of the model, there are two noteworthy simplifications that we exploit in both characterizing a solution to the planner's problem and solving it numerically.

Remark 7.1. *The model, as posed, has a solution that conveniently separates. We may solve two separate control problems: i) determines “undamaged” consumption, investment and capital ii) determines emissions, the temperature anomaly and damages. It is the latter one that is of particular interest. Undamaged consumption, investment and capital are merely convenient constructs that allow us to simplify the model solution.*

Remark 7.2. *We obtain a further simplification by letting:*

$$\tilde{\mathcal{E}}_t = \mathcal{E}_t (\iota_y \cdot Z_t).$$

We use $\tilde{\mathcal{E}}_t$ as the control variable and then deduce the implications for $\mathcal{E}_t$.

Remark 7.3. *In this illustration, the costs of emissions are given solely by the environmental and economic damages. In our previous research, we followed Bornstein et al. (2017) and Casassus et al. (2018) by including reserves as a state variable that can be augmented by an investment in*

new discoveries. See Barnett et al. (2020). While this richer specification has more substantive interest, the emissions costs implied by this technology were quite small relative to social costs induced by damages. For this example, we dropped this additional state variable to simplify further our characterization of the uncertainty implications.

Remark 7.4. As Anderson et al. (2003) note, there is a preference equivalence between a concern for model misspecification and risk aversion in the recursive utility formulation of Kreps and Porteus (1978) and Epstein and Zin (1989) when there is a unitary elasticity of substitution as we have assumed here. Our macroeconomic model, by design, can capture what is called “long-run risk” in the macro-finance literature in the absence of climate change. See Bansal and Yaron (2004). The long-run risk literature explores the valuation consequence of growth rate uncertainty using a recursive utility model of investor preferences. The preference specification presumes a full commitment to the baseline probabilities, but the rationale for this commitment appears to be weak when confronting specific forms of growth rate uncertainty. The long-run risk literature often imposes a seemingly large risk aversion parameter that arguably can look more plausible when reinterpreted as a concern of model misspecification. In what follows, we will impose $\xi_r = 5$ as our largest value. The implied risk aversion under recursive utility is twenty one, which is certainly large but typically not dismissed as too large in the empirical literature on long-run risk. The Brownian drift induced by a robustness concern is quite sizable for the counterpart adjustment to the consumption/capital dynamics. Borrowing and updating a specification of growth rate uncertainty of Hansen et al. (2008), Hansen and Sargent (2020a) fit a simple consumption/capital model to aggregate data designed to measure macroeconomic growth-rate uncertainty. Their model is the undamaged version of the model we pose here, with two shocks. One shock is to the stochastic process for growth rate productivity, and the other is an independent shock to only the capital productivity. These shocks imply two of the consumption shocks in Bansal and Yaron (2004).¹⁷ The implied drift distortions for the stochastic capital evolution are:

$$h = \begin{bmatrix} -.715 \\ -.170 \end{bmatrix} \quad \begin{array}{l} \text{productivity growth rate shock distortion} \\ \text{capital productivity shock distortion.} \end{array}$$

Given the relative magnitudes of the adjustments, it is the growth uncertainty channel that is of particular importance to the decision-maker. Since .715 is large in comparison to the unit standard deviation of the shock, one might dismiss our choices of ξ_r as being too extreme.

While we find the comparison between misspecification and risk aversion in the presence of growth rate uncertainty to be revealing, an uncomfortable feature of the long-run risk formulation

¹⁷These shocks imply two of the consumption shocks in Bansal and Yaron (2004). Bansal and Yaron also include a shock to stochastic volatility that we abstract from here and consider implications for changing the intertemporal elasticity of substitution.

in the macro asset pricing literature is that the long-run risks are a “black box.” In what follows, we will step back from the dual interpretation of risk aversion and abstract from misspecification implications for the undamaged consumption or capital evolution. The implied Brownian motion distributions for the climate dynamics will turn out to be considerably smaller than the growth rate shock distortion reported here.

7.2 HJB equations and robustness

We now describe our approach to solving the model and incorporating concerns about robustness and ambiguity aversion. The uncertainty that we consider has a single jump point after which the damage function uncertainty is revealed. This leads us to compute continuation value functions conditioned on each of the damage function specifications. These continuation value functions then are used to summarize post-jump outcomes when we compute the initial value function. We describe the Hamilton-Jacobi-Bellman (HJB) equations for each of these steps in what follows. Some further details about the computations and parameter settings are provided in Appendix B.

7.2.1 Post-jump continuation value functions

We first compute the value functions pertinent after the Poisson event that reveals that the damage function curvature is realized. While the damage specification uncertainty is resolved with this event, there remains climate model uncertainty.

The state variables are the temperature anomaly, and the exogenous Brownian uncertainty. Recall that after the Poisson event, the argument of the function Γ is $\bar{Y}_t = Y_t - Y_\tau + \bar{y}$, where τ is the date of the Poisson event. We abuse notation a little bit by letting y denote a potential realization of $\bar{Y}_t$. Since $-Y_\tau + \bar{y}$ is time invariant post-jump, both Y_t and $\bar{Y}_t$ share the same increment. Moreover, $\bar{Y}_\tau = \bar{y}$, which is a pertinent boundary condition for the post-jump value functions.

We solve the optimization problems for the continuation value functions ϕ_m conditioned on each of the damage functions, $m = 1, 2, \dots, M$. The ϕ_m ’s are then computed conditioned on each of the damage functions, $m = 1, 2, \dots, M$. Because the optimization problem is well-posed whenever $y \geq 0$, we only care about the range in which $y \geq \bar{y}$. In our computations we use twenty equally spaced values for γ_3^m .

In formulating the HJB equation, we include robustness considerations as described in Section

5 where ξ_a and ξ_r are penalty parameters. The HJB equation conditioned on a ϕ_m is given by:

$$\begin{aligned}
0 = \max_{\tilde{e}} \min_h \min_{\omega_j, \sum_{\ell=1}^L \omega_\ell = 1} & -\delta \phi_m(y) + \eta \log \tilde{e} \\
& + \frac{d\phi_m(y)}{dy} \tilde{e} \varsigma \cdot h + \frac{(\eta-1)}{\delta} [\gamma_1 + \gamma_2 y + \gamma_3^m(y - \bar{y})] \tilde{e} \varsigma \cdot h + \frac{\xi_r}{2} h' h \\
& + \frac{d\phi_m(y)}{dy} \sum_{\ell=1}^L \omega_\ell \theta_\ell \tilde{e} + \frac{1}{2} \frac{d^2 \phi_m(y)}{(dy)^2} |\varsigma|^2 \tilde{e}^2 \\
& + \frac{(\eta-1)}{\delta} \left([\gamma_1 + \gamma_2 y + \gamma_3^m(y - \bar{y})] \sum_{\ell=1}^L \omega_\ell \theta_\ell \tilde{e} + \frac{1}{2} (\gamma_2 + \gamma_3^m) |\varsigma|^2 \tilde{e}^2 \right) \\
& + \xi_a \sum_{\ell=1}^L \omega_\ell (\log \omega_\ell - \log \pi_\ell). \tag{9}
\end{aligned}$$

In this calculation, the π_ℓ 's are the climate model probabilities, and the ω_ℓ 's are the alternative probabilities that we consider when making a robust adjustment. A jump happening in the model is equivalent to an increase in y to $\bar{y}$. After a jump, the model follows one of the m damage specifications with y reinitialized at the threshold. For computational purposes, the $\phi_m(\bar{y})$'s are a fixed set of numbers imposed in the HJB computation. For a given ω and h , the minimization over $\tilde{e}$ has a quadratic objective and can be solved quasi-analytically. For a given $\tilde{e}$, the minimization problem for h is quadratic and can also be solved analytically. Similarly, the minimizing ω_ℓ 's satisfy:

$$\omega_\ell^* \propto \pi_\ell^a \exp \left(-\frac{1}{\xi_a} \left[\frac{d\phi_m(y)}{dy} \theta_\ell \tilde{e} + \frac{(\eta-1)}{\delta} [\gamma_1 + \gamma_2 y + \gamma_3^m(y - \bar{y})] \theta_\ell \tilde{e} \right] \right).$$

The expression on the right side requires scaling in order that the resulting ω_ℓ^* 's sum to one. We exploit the resulting formulas in our recursive computations. We solve (9) for the M continuation value functions that we use in our analysis.

7.2.2 Pre-jump value function

The pre-jump value function has a similar structure with two exceptions: i) we include the intensity function discussed earlier and ii) we introduce robustness concerns for both the intensity and distribution over the alternative γ_3^m coefficients. Given these modifications, we include:

$$\mathcal{J}(y) \sum_{m=1}^M g_m \pi_m [\phi_m(\bar{y}) - \phi(y)] + \xi_r \mathcal{J}(y) \sum_{m=1}^M \pi_m (1 - g_m + g_m \log g_m) \pi_m$$

in the HJB equation and to minimize with respect to the non-negative g_m 's. The continuation value functions ϕ_m are all evaluated at $y = \bar{y}$ in this HJB equation pertinent for the pre-jump analysis. This occurs because immediately after a Poisson event is triggered, the damage function

Γ is evaluated at the threshold point $\bar{y}$.

We now illustrate the impact of uncertainty about the magnitude and timing of damages from climate change along with uncertainty in the carbon-temperature dynamics.

Our specifications of aversion to climate model ambiguity and damage function misspecification are expressed in terms of the two penalty parameters, ξ_a and ξ_r . We use these parameters to restrain the search over alternative probabilities. The outcome of this search is an uncertainty adjusted probability measure that is of interest for two reasons. First, it shows implied probabilities that are most problematic to the decision-maker. Should these appear to be too extreme, then the penalization used in the decision problem is not severe enough. This type of inspection follows common practices for robust Bayesian methods following the proposal of Good (1952). Second, these altered probabilities provide an adjustment for social valuation implied by model ambiguity and misspecification uncertainty.

7.3 Robust adjustments to climate model uncertainty

For the 144 carbon-climate dynamic models, we take as our baseline probabilities an equal weighting of all of the models. While it is straightforward to explore a whole family of values for ξ_a , in the calculations that follow, we set $\xi_a = .01$. To determine whether or not this is a reasonable choice of ξ_a to use in our analysis, we examine the implied distortion to the probability distribution of θ_ℓ values resulting from our choice as compared to the baseline prior probability distribution. Both the original prior probability distribution (red histogram) and the distorted probability distribution (blue histogram) of θ_ℓ values are given in Figure 2. The increased concern about uncertainty over the geo-scientific inputs leads to a shift to the right in the θ_ℓ probability distribution, highlighting increased concerns about worst-case climate dynamics, while still maintaining a spread in the weights on the values of θ_ℓ and not loading all the weight on the far right tail. We, therefore, view this shift in the distribution as reasonable to entertain. The implied mean distortion is about 0.26 for the unknown parameter θ . While the concerns about geo-scientific uncertainty are state-dependent, the distortion in the probability distribution for θ remains roughly constant over the course of our simulations.

There is an additional mean shift in this temperature distribution that is induced by misspecification concerns. This shift is negligible for $\xi_r = 5$ and only about .07 for $\xi_r = 1$. Its impact is much for substantial for the low penalty, $\xi_r = 0.3$ with an implied mean shift of about 0.23. These distortions show very little sensitivity to the state dynamics.




Figure 5: Histograms of climate sensitivity parameters. The red histogram is the outcome of equally weighting all 144 climate models. The blue histogram is the outcome of the minimization in the recursive formulations of our social planners problem.

7.4 Robust adjustments to damage function uncertainty

We next consider the penalty parameter ξ_r that governs concerns about misspecifying the Poisson jump process, including both the jump intensity and the probability distribution conditioned on a jump. Recall that we use this process to capture uncertainty of the steepness in the damage function and timing of when this steepness becomes known to the decision-maker. This uncertainty is only pertinent prior to the realization of the Poisson event. We report results for three different values of this parameter $\xi_r = 5, \xi_r = 1, \xi_r = .3$ in Figure 6. The distorted histogram for the lowest value, $\xi_r = .3$, is arguably extreme, although the other two choices seem considerably harder to dismiss.




Figure 6: Distorted probabilities of damage functions. Baseline probabilities for damage functions are $1/20$ (red bars), and the blue bars are robust adjustments to the probabilities induced by model misspecification concerns (left panel: $\xi_r = 5$, center panel $\xi_r = 1$, right panel: $\xi_r = 0.3$). These histograms are the outcome of recursive minimizations. These distortions are close to being constant as the temperature anomaly increases up to the Poisson jump date.

Finally, in Figure 7, we display the probabilities that a jump will occur prior to the specified dates along the socially efficient trajectory for emissions. Again, we impose $\xi_r = 1$. The jump is pretty much assured to happen by about one hundred years out, at which point the temperature anomaly is 2 degrees Celsius. On so-called “business as usual” trajectories, the jump probabilities will converge to one much more quickly than what is displayed in this figure:




Figure 7: The left panel shows the probabilities that a jump will occur prior to the date given on the horizontal axis. The right panel shows the simulated pathway for the temperature anomaly and the points where the anomaly reaches $\underline{y} = 1.5$ and $\bar{y} = 2.0$ (black dashed lines).

7.5 Emission and anomaly trajectories

To show the effects of concerns about damage function uncertainty on policy decisions of the planner, we explore the behavior of emissions for different amounts of aversion to uncertainty. In

Figure 8, we report the pre-jump control laws for $\tilde{e}$ as a function of the temperature anomaly, and $\underline{y}$, for three values of ξ_r in Figure 8. While the e , in contrast to $\tilde{e}$, also depends multiplicatively on the exogenous state vector, it is the dependence of the temperature anomaly that is of particular interest. For comparison, this figure also includes the control law for $\tilde{e}$ when the planner has full commitment to the baseline probabilities. We confine the domain of the control laws to temperature anomalies between 1.1 and 1.5 degrees Celsius. After the temperature anomaly reaches 1.5, the probability of a jump occurring becomes non-zero.

As we see, even in advance of gaining more information about damage function curvature, the fictitious planner embraces a substantial level of precaution due to the concerns about the unknown future damage state. In light of uncertainty concerns, the control law for emissions is about twenty percent lower when $\xi_r = 1$ than the control law based solely on the baseline probabilities. We also see that, as the value of ξ_r is decreased, the caution is amplified and the choice of emissions is lowered even further. It follows that the emission trajectories for the lower control laws necessarily reach the $\underline{y} = 1.5$ threshold later starting from a common initial condition.




Figure 8: Emissions as a function of the temperature anomaly for different penalty configurations. The thresholds are $\underline{y} = 1.5$ and $\bar{y} = 2.0$. We limit the domain of the function to be 1.1 – 1.5 because for larger temperature anomalies the Poisson event may be realized.

While the 1.5 and 2 degree thresholds have dominated much of the policy discussion, there is debate as to the extent to which these are firmly backed up by evidence. For this reason, in Figure 9, we report the consequences of shifting the thresholds we use in our computations to $\underline{y} = 1.75$ and $\bar{y} = 2.25$. The results for emissions are very similar, except that in comparison to

Figure 8, the control laws are shifted to the right as should be expected because of delay in when the more extreme damage function curvature is manifested.




Figure 9: Emissions as a function of the temperature anomaly for different penalty configurations. The thresholds are $\underline{y} = 1.75$ and $\bar{y} = 2.25$. We limit the domain of the function to be $1.1 - 1.75$ because for larger temperature anomalies the Poisson event may be realized.

Returning to our original specification, when the temperature anomaly reaches values of $y \in [\underline{y}, \bar{y}] = [1.5, 2]$, the Poisson event revealing the full damage function will at some point be realized. Once it is revealed, the emissions trajectories will jump to either a higher or a lower level depending on how much damage function curvature is realized. We report the initial emissions, post-jump, in Figure 10 as a function of γ_3 governing the curvature of the damage function for large temperature anomalies. Importantly, this function is highly convex. The realization of a very low damage function curvature is good news for the planner, resulting in an increase in emissions in contrast to many of the other damage function specifications that could be realized. For the damage functions with even a little more curvature, there is a large reduction in emissions as reflected in the steep slope of the function of optimal emissions choices for small values of γ_3 . The emissions choices are increasingly more concentrated at similar values for higher curvature, as seen in the much flatter slope for the larger values of γ_3 .




Figure 10: Emissions choices, conditioned on a jump having occurred, for different realized damage function parameters γ_3 upon realization of the jump.

7.6 Temperature anomalies

Given the probabilistic nature of the Poisson event, the emissions and resulting temperature anomalies behave probabilistically, even abstracting from Brownian motion risk. While the planner has uncertainty about these probabilities, we find it revealing to report the distributions under the baseline specification. We show the implied temperature anomaly distributions conditioned on the Poisson event being realized in Figure 11.¹⁸ We extend the simulation out one hundred years to ensure that the no jump probability is essentially zero. The vast majority of temperature anomaly values are less than the 2.0 degree threshold, though there is a small right tail going beyond that and a small right peak leading up to it. This relatively constrained distribution of temperature anomalies is driven largely by the initial caution exercised by the planner and the continued caution for most of the damage function realizations. A small fraction of realizations of the damage function curvature parameter γ_3 result in higher temperature anomalies than 2.0 degrees because the planner is willing to increase emissions after the Poisson event reveals “good news.”

¹⁸The number of outcomes in the histograms is determined by the number of values of γ_3^m , which is twenty, and the time discretization used in the simulation.




Figure 11: Histogram of possible temperature anomaly values for the scenario where $\xi_a = 0.01$, $\xi_r = 1$. The temperature anomaly values are for year 100. The simulation is done under the baseline probabilities and abstracts from the Brownian motion shocks.

7.7 Social cost of carbon

We use the SCC of carbon evaluated at the socially efficient trajectory as the barometer for the economic externality induced by climate change. The planner equates marginal social costs and benefits of emissions. We represent the marginal benefits in units of damaged consumption so that:

$$SCC_t = \frac{\eta(\tilde{C}_t)}{(1-\eta)(\tilde{\mathcal{E}}_t)} = \frac{\eta(C_t)}{(1-\eta)(N_t)\tilde{\mathcal{E}}_t}$$

where the right-hand side variables are evaluated along the socially efficient trajectory. Taking logarithms, we get:

$$\log SCC_t = \log \eta - \log(1-\eta) + (\log C_t - \log N_t) - \log \tilde{\mathcal{E}}_t$$

As we noted previously, “undamaged consumption” evolves in a manner consistent with a long-run risk model familiar from macro asset pricing. The logarithm of consumption grows stochastically along a linear trajectory with variation increasing approximately linearly over the growth horizon. Our focus instead will be on the behavior of:

$$\log \eta - \log(1-\eta) - (\log \tilde{N}_t) - \log \tilde{\mathcal{E}}_t. \quad (10)$$

where $\log \tilde{N}_t$ excludes the exogenous stochastic contribution to $\log N_t$, which is common across our specifications of the robustness parameter ξ_r . The variation of this measure over time depends entirely on the temperature anomaly trajectory prior to reaching the lower threshold $\underline{y}$. In the figures that follow we report this dependence.

Since emissions depend on the temperature anomaly, under the planner's solution there is an important distinction between the SceRF and the SysRF discussed in Section 6.3. The SceRF for the temperature anomaly behaves in accordance with the Matthew's approximation whereby emissions today have a permanent impact on the future temperature anomaly. The SysRF incorporates the dependence of emissions on the temperature anomaly, and it is the SysRF that is embedded in the SCC computation for the social planner's problem. While this dependence is a direct outcome of the planner's problem, more generally the plausibility of future emissions trajectories should be tied to the potential policy responses as we experience the impact of climate change.

To deduce the emissions contribution, we differentiate the HJB equation (9) with respect to $\tilde{e}$ and solve for $\frac{\eta}{\tilde{e}}$:

$$\frac{\eta}{\tilde{e}} = -\frac{d\phi(y)}{dy} \sum_{\ell=1}^L \omega_{\ell} \theta_{\ell} - \frac{d^2\phi(y)}{(dy)^2} |\varsigma|^2 \tilde{e} + \frac{(1-\eta)}{\delta} \left[(\gamma_1 + \gamma_2 y) \sum_{\ell=1}^L \omega_{\ell} \theta_{\ell} + \gamma_2 |\varsigma|^2 \tilde{e} \right] \quad (11)$$

for $0 < \underline{y}$. For the planner's problem, since marginal benefits are equated to marginal costs, we may use either side of this equation to measure the emissions contribution to the SCC.

Figure 12 shows the log SCC for the baseline case of $\xi_r = \infty, \xi_a = \infty$ (red line), and three cases of increasing concerns about damage function uncertainty: $\xi_r = 5$ (green line), $\xi_r = 1$ (yellow line), and $\xi_r = 0.3$ (blue line) where $\xi_a = .01$ for all three cases. The results use the socially efficient emissions trajectories for the different values of ξ_r . The log SCC values are calculated using formulas (1) and (11). In Figure 12, we see substantial values of the log SCC in each case. The magnitudes are amplified as we increase concerns about damage function misspecification (by decreasing the value of ξ_r). In particular, as a function of the temperature anomaly, the SCC for $\xi_r = 1$ is between and twenty and thirty percent higher than when we abstract from robustness concerns. Although not reported here, changing the thresholds to be $\underline{y} = 1.75$ and $\bar{y} = 2.25$ effectively shifts the curves in Figure 12 degrees to the right with a corresponding smaller SCC at the initial temperature anomaly of $y = 1.1$.

7.8 Summary

In this example economy, the social planner adopts an emissions policy that is cautious at the outset even though considerably more information about potential damages will be available in the future. The damage function uncertainty is resolved by a single Poisson event that becomes more




Figure 12: $\log(SCC)$ as functions of temperature anomaly under different penalty configurations. The SCC is state-dependent, and we focus on the domain for which the anomaly is less than 1.5 degrees. The values in logarithms are translated by the initial period logarithm of consumption.

likely the larger the temperature anomaly. Once this event is realized, there is an asymmetric response. For a small fraction of the damage functions with the most modest curvature, emissions immediately increase. These are “good news events” and are determined endogenously within our model. For a much larger fraction of damage function specifications, the emissions responses continue to be modest, although the magnitude of these responses depend on the curvature of the damage function that is revealed. The implied social costs of carbon increase by about twenty percent due to the combined damage function and carbon-climate model uncertainty prior to the realization of the Poisson event. This impact can be larger or smaller depending on the social planner’s aversion to ambiguity and model misspecification. While acknowledging the simplified nature of the model used for our computations, our results demonstrate the importance of accounting not only for different uncertainty channels, but also for the information dynamics when designing optimal climate policy.

8 Uncertainty decomposition

An advantage to the more structured approach implemented as smooth ambiguity is that it allows us to “open the hood” so to speak on uncertainty. We build on the work of Ricke and Caldeira (2014) by exploring the relative contributions of uncertainty in the carbon dynamics versus uncertainty in the temperature dynamics. We depart from their analysis by studying the

relative contributions in the context of a decision problem and we include robustness to model misspecification as a third source of uncertainty. This latter adjustment applies primarily to the damage function specification. We continue to use the social cost of carbon as a benchmark for assessing these contributions. We perform these computations using the model developed in the previous section, although the approach we describe is applicable more generally. For the uncertainty decomposition, we hold fixed the control law for emissions, and hence also the implied state evolution for damages, and explore the consequences of imposing constraints on minimization over the probabilities across the different models.

Recall that we use climate sensitivity parameters from combinations of 16 models of temperature dynamics and 9 models of carbon dynamics. A parameter θ corresponds to climate-temperature model pair. Let Θ denote the full set of $L = 144$ pairs, and let P_j for $j = 1, 2, \dots, J$ be a partition of the positive integers up to L . The integer J is set to 9 or 16 depending on whether we target the temperature models or the carbon. For any given such partition, we solve a constrained version of the minimization problem (4) by targeting the probabilities assigned to partitions while imposing the benchmark probabilities conditioned on each partition.

$$\begin{aligned} \min_{\bar{\omega}_j, j=1,2,\dots,J} & \left(\frac{\partial V}{\partial x} \right) \cdot \sum_{j=1}^J \bar{\omega}_j \sum_{\ell \in P_j} \left(\frac{\pi_\ell}{\sum_{\ell \in P_j} \pi_\ell} \right) \mu(x, a \mid \theta_\ell) \\ & + \xi_a \sum_{j=1}^J \bar{\omega}_j (\log \bar{\omega}_j - \log \bar{\pi}_j) \end{aligned}$$

where: $\bar{\pi}_j = \sum_{\ell \in P_j} \pi_\ell$ and

$$\frac{\pi_\ell}{\bar{\pi}_j} \quad \ell \in P_j$$

are the baseline conditional probabilities for partition j . We only minimize the probabilities across partitions while imposing the baseline conditional probabilities within a partition.

We impose $\xi_r = \infty$ when performing this minimization and let $\xi_a = .01$ as in section 7. We perform additional calculations where we let $\xi_r = 1$ and $\xi_a = \infty$ in order to target damage function uncertainty rather than temperature or climate dynamics uncertainty.¹⁹ The two states in our problem are $x = (y, n)$, and we look for a value function of the form $V(y, n) = \phi(y) + \frac{(\eta-1)}{\delta} n$ while imposing that $\tilde{e} = \epsilon(y)$. For each partition of interest, we construct the corresponding HJB equation that supports this minimization.

Since we are imposing the control law for emissions but constraining the minimization, the first-order conditions for emissions will no longer be satisfied. Recall formula (8) from section 6 with adjustments for uncertainty. In the absence of optimality, the net benefit measure $MV(x)$ is not zero with the minimization constraints imposed. Consistent with the SCC computation from

¹⁹While the robustness adjustment also applies to the climate dynamics, as we saw in the previous section, this adjustment was small relative to the ambiguity adjustment.

the previous section, we use

$$-\frac{\partial V}{\partial x}(x) \cdot \frac{\partial \mu}{\partial e}[x, \phi(x)] - \frac{1}{2} \text{trace} \left[\frac{\partial^2 V}{\partial x \partial x'}(x) \frac{\partial}{\partial e} \Sigma[x, \phi(x)] \right].$$

for our cost contributions in the SCC decomposition.

We obtain the smallest cost measure when we preclude minimization altogether while solving for the value function and the largest one when we allow for full minimization with $\xi_r = 1$ and $\xi_a = .01$. We have three intermediate cases corresponding to temperature dynamic uncertainty, climate dynamic uncertainty and damage function uncertainty. The smallest of these measures corresponds to a full commitment to the baseline probabilities. We form ratios with respect to the smallest measure, take logarithms and multiply by 100 to convert the numbers to percentages. Importantly, we change both probabilities and value functions in this computation.

We report the results in Figure 13. From this figure, we see that the uncertainty adjustments in valuation account for twenty to thirty percent of the social cost of carbon. The contributions from temperature and carbon are essentially constant over time with the temperature uncertainty contribution being substantially larger. The damage contribution is initially below half the total uncertainty, but this changes to more than half by the time the temperature anomaly reaches the lower threshold of 1.5 degrees Celsius.




Figure 13: Uncertainty decomposition for the logarithm of the marginal value of emissions (scaled by 100). These computations impose $\xi_a = .01$ and $\xi_r = 1$. The figures report log differences between marginal values of the different components relative to baseline probability counterparts. The uncertainty partitions account separately for temperature dynamics ambiguity, carbon dynamics ambiguity and robustness to damage function misspecification.

Remark 8.1. *The uncertainty decomposition we implement depends on the underlying emissions trajectory we impose. For the reported computations, we used the planner’s solution for when all uncertainty components are considered. Since our planner cares about uncertainty, robustness considerations lead our planner to avoid excessive exposure to uncertainty when possible. In our particular setting, with uncertainty aversion, the planner will prefer to avoid being vulnerable to damage function uncertainty, which can be achieved in part by delaying when the potentially steep slope of the damage function becomes operative. Yet, the exposure components of uncertainty can look very different for, say, business-as-usual trajectories of emissions or even socially optimal trajectories of emissions that do not incorporate concerns about uncertainty. Thus, our decompositions are of potential interest for emissions trajectories other than those chosen as part of a solution to an uncertainty averse planner’s problem.*

9 Carbon abatement technology

While the model posed in section 7 illustrated how the unfolding of damages should alter policy, the economic model was not designed to confront transitions to fully carbon-neutral economies.

There have been several calls for such transitions with little regard for the role or impact of uncertainty. We now modify the model to allow for green technology in decades to come.

We next consider a technology that is close to the Dynamic Integrated Climate-Economy (DICE) model of Nordhaus (2017). See also Cai et al. (2017) and Cai and Lontzek (2019) for a stochastic extension (DSICE) of the DICE model.²⁰ For our setting, we alter the output equation from our previous specification as follows:

$$\frac{I_t}{K_t} + \frac{C_t}{K_t} + \frac{J_t}{K_t} = \alpha$$

where:

$$\frac{J_t}{K_t} = \begin{cases} \alpha \vartheta_t \left[1 - \left(\frac{\mathcal{E}_t}{\alpha \lambda_t K_t} \right) \right]^\theta & \left(\frac{\mathcal{E}_t}{\alpha K_t} \right) \leq \lambda_t \\ 0 & \left(\frac{\mathcal{E}_t}{\alpha K_t} \right) \geq \lambda_t \end{cases} \quad (12)$$

To motivate the term J_t , express the emissions to capital ratio as:

$$\frac{\mathcal{E}_t}{K_t} = \alpha \lambda_t (1 - \iota_t)$$

where $0 \leq \iota_t \leq 1$ is *abatement* at date t . The exogenously specified process λ gives the emissions to output ratio in the absence of any abatement. By investing in ι_t , this ratio can be reduced, but there is a corresponding reduction in output. Specifically, the output loss is given by:

$$J_t = \alpha K_t \vartheta_t (\iota_t)^\theta$$

Equation (12) follows by solving for abatement ι_t in terms of emissions.²¹ The planner's preferences are logarithmic over damaged consumption:

$$\log \tilde{C}_t = \log C_t - \log N_t = (\log C_t - \log K_t) - \log N_t + \log K_t.$$

In contrast to the previous specification, the planner's value function for this model is no longer additively separable in (y, k) , although it remains additively separable in log damages, n .

For the purposes of illustration, we consider two "Poisson events" that reduce both (ϑ_t, λ_t) , rep-

²⁰Among other stochastic components, the DSICE incorporates tipping elements and characterizes the SCC as a stochastic process. From a decision theory perspective, DSICE focuses on risk aversion and intertemporal substitution under an assumption of rational expectations.

²¹The link to the specification used in Cai and Lontzek (2019) is then:

$$\begin{aligned} \sigma_t &= \lambda_t \\ \vartheta_t &= \theta_{1,t} \\ \theta &= \theta_2 \\ \mu_t &= \iota_t \end{aligned}$$

representing technological innovations that decrease the cost of abatement and improve the emissions-to-output ratio. The first jump cuts (ϑ_t, λ_t) in half, and the second jump sets $(\vartheta_t, \lambda_t) = (0, 0)$, indicating a transition to a purely carbon-neutral economy. Both events have the same constant intensity, which we set so that the expected arrival time is twenty years. The stochastic specification of damages remains the same as in the previous models. Not surprisingly, these two new Poisson events change substantially our calculations. In our discussion that follows, we highlight a few of the important differences.

The penalty parameters ξ_r and ξ_a are not necessarily transportable across the different models. Instead, it is sensible to loosen the penalty settings for more complicated economic environments to achieve probability distortions of comparable magnitudes. In our calculations, we increased $\xi_a = .02$, making the implied distorted distribution for the climate sensitivity parameter similar to the one we computed for section 7. We again explore three settings for the robustness parameter ($\xi_r = 2.5, 5, 7.5$), and we explore which of the three Poisson events is of most concern to the social planner. Appendix C provides more detail about the parameter values we use and the approach to computation.

Figure 14 shows the baseline jump probabilities and along with the implied distorted probabilities for the three robustness settings. Since the second technological jump necessarily follows the first, the probabilities reported in the middle panel take as a starting point, the date at which the first technology advance is realized. Comparing the probabilities across the different plots, it is clear that the probabilistic specification of the first green technology advance is of the most concern to the planner. In particular, when we set $\xi_r = 2.5$, the probability slanting is arguably extreme implying probabilities that are about seventy percent lower than the baseline probabilities. But even when $\xi_r = 5$, there is about a twenty-five percent reduction in the probabilities relative to the baseline. The reductions are notably smaller for the other two Poisson events. Given the prospect of advances in green technology, damage function uncertainty concerns are less concerning to the planner than the technological uncertainty. Notice that in this economy, the technological advancements make it much more plausible to avoid the more severe damages.²²

²²To provide further confirmation of this interpretation, the initial emissions are a little higher for this model than the ones from section 7 that are depicted in Figure 9. They now range between 6.9 and 8.1, depending on the value of ξ_r .




Figure 14: Distorted probability of the Poisson events for technology changes and damages under different penalty configurations. The simulation uses the planner's optimal solution. The left panel shows the distorted jump probabilities for the first technology jump. The middle panel shows the distorted jump probabilities for the second technology jump. The right panel shows the distorted jump probabilities for the damage function curvature jump. The baseline probabilities for the right panel are computed using the state dependent intensities when we set $\xi_a = \xi_r = \infty$.

Figure 15 reports the probability distortions for the damage function and climate sensitivity models. Here, we have imposed $\xi_a = .02$ and $\xi_r = 5.0$. Note that the damage function probability distortions are relatively modest, consistent with our previous discussion. The climate model distortions, by design, are of similar magnitude as those reported previously in Figure 5.




Figure 15: Distorted probabilities of damage functions and climate models. These computations impose $\xi_a = .02$ and $\xi_r = 5.0$. Baseline probabilities are given by the red bars, and the blue bars are robust adjustments to the probabilities induced by model uncertainty concerns. The left panel shows the damage function probabilities and the right panel shows the climate model probabilities. The histograms are the outcome of recursive minimizations with the distortions calculated at year 40 of the simulation.

To summarize, it is the probability of the first technological advance that is of the biggest concern to the social planner. In particular, uncertainty in the environmental and economic damages induced by climate change is now less problematic given the potential advances in green technology.

The possibilities of technological improvements appearing sometime in the future together with the extra ability to mitigate emissions by paying a cost dampens the strong precautionary behavior of the fictitious planner that we found before. This finding points out the importance of adding an R & D sector for the development of more productive green technologies and to open the door to policies that subsidize these sectors as has been explored in other studies.²³ Since R & D investments can be highly speculative, the methods in this paper would allow for an investigation of the consequent uncertainties in addition to the ones explored here.

10 Denouement

We draw our paper to end by taking inventory on what we see as the valuable messages. During the course of writing and presenting this paper, we have received a variety of comments, some of which have helped us set the stage for future research, and others that are based on a potential misunderstanding about where our research is positioned relative to other discussions of prospective emissions and the social cost of carbon.

10.1 Quantitative storytelling

We have used highly stylized models to help us better understand how uncertainty should impact prudent policymaking. There are two interrelated reasons for the stylized nature of our exercise: one is to preserve tractability, and the other is to ensure that mapping from model inputs to conclusions is transparent. The models we use in this paper have obvious flaws and limitations, but the resulting analyses are intended to point to where more ambitious models might contribute to the knowledge base.

Calibrating such highly stylized models is a bit like walking on a tightrope. We want the findings to have credibility, with results that are not driven by perverse or extreme parameter settings. But the models we are parameterizing deliberately omit elements of both the climate and economic dynamics. We view this research as an initial step in what we hope will be a research agenda that will become even more substantively ambitious as we explore increasingly complex dynamic models and the associated uncertainties.

10.2 SCC

Our choice in this paper to feature solutions to planners' problems has important ramifications for interpreting our measurement of the social cost of carbon relative to the modular approaches that we discussed previously. As we noted, others have also featured planners' approaches in part motivated by the Pigouvian approach for taxing externalities. The planner's solution pins

²³See, for instance, Acemoglu et al. (2016).

down the contingency rule for emissions expressed in terms of the Markov states. Since the planner internalizes the impact of carbon emissions on damages, this has important ramifications for emissions trajectories and the SCC measured as a shadow price of emissions relative to a consumption numeraire. For instance, even the lower bound temperature anomaly $\underline{y}$ where the damage function jumps begin to occur in our analysis will not be crossed for somewhere between fifty and seventy-five years when simulating this path going forward for our first model. This pushes the potentially extreme damage realizations far into the future relative to say, “business-as-usual” trajectories. The SCCs are substantial relative to other computations based on solutions to planners’ problems and made larger because of the explicit adjustments for model ambiguity and robustness to potential misspecification. Recall that our approach to uncertainty quantification results in a two-parameter characterization depending on the two forms of aversion. While we reported sensitivity to misspecification aversion in the paper, we show the impact of additional concerns about model ambiguity in our accompanied notebook. Increasing ambiguity aversion in ways that are arguably plausible increases further the uncertainty contribution to about thirty-five percent. But these are for an illustrative model with no backup or green technology, the inclusion of which would alter substantially both the emissions trajectory and the SCC. In fact, this is what happens in the second model we consider. An interested reader can see these additional results and others in an online notebook we constructed to supplement this paper.²⁴ We mention such numbers to convey that the planner in our examples is responding aggressively to climate change, but such numbers are not good candidates for numerical values to post on EPA webpages because of the preliminary and stylized nature of the models including the uncertainty inputs into the computations.²⁵

While there is path uncertainty in our model because of shocks that impinge on the dynamical economic system, our approach does not include the multiple exogenous scenarios and accompanying uncertainty across the scenarios that are part of the modular approaches. The practical connection of the computation of the SCC computed from a social optimum to taxation is limited because it ignores the presence of other taxes and other constraints on policymaking including cross-country coordination. Nevertheless, in our view, the socially efficient allocation remains interesting as a benchmark to compare with existing policy outcomes. Within this context, we find the SCC to be a revealing barometer for the importance of uncertainty.

As explained in NYSERDA and RFF (2021) and elsewhere, the modular approach advocated by some for the purposes of regulation is local in nature. It is computed as a local or a small change around an existing allocation. While perhaps suggestive, this approach does not quantify

²⁴Our online notebook, which includes these supplemental results and the code used to derive our model solutions, is available at <https://climateuncertaintyspillover.readthedocs.io/en/latest/>.

²⁵As readers of Koonin (2021) and of the challenges by climate scientists to some of its claims such those noted in Mills et al. (2021), we expect rather heterogeneous views about plausible uncertainty bounds to use in quantitative investigations like ours.

the impact of global changes in policy. As this local construct of an SCC is also forward-looking, it depends on the prospective future emissions and the damages that they might induce. The actual computations depend on emissions scenarios. These are alternative emission pathways, typically imposed exogenously. The differences in the socially efficient emissions process and the array of emission scenarios can have an important impact on the SCC calculation. The treatment of uncertainty for this local approach typically a) is static in nature, b) presumes, at least implicitly, a priori probabilities across a small array of scenarios, and c) is outside the decision problem. The outcome is a histogram of SCCs instead of an SCC that incorporates uncertainty. Thus, this local construct of the SCC is very different from the one used in our paper, all the more so given common implementations.²⁶

The methods we describe could be applied to local measures of the SCC, provided that emissions scenarios be specified recursively in ways amenable to the application of dynamic programming methods. Importantly, there is more to discounting than the choice of a single constant discount rate. As featured in Barnett et al. (2020), these uncertainty adjustments can be most conveniently represented as a change in the probability measure used for evaluation. These types of adjustments lead naturally to important interactions in the components of the commonly employed modular approach. SCC measurements designed to incorporate uncertainty within a decision-making or policy-design framework cannot be easily modularized.

10.3 Damage functions

Our damage function specifications, like many other ones in the integrated assessment literature, should be taken with a “grain of salt” so-to-speak. They are *ad hoc* primarily to support tractable illustration. As we mentioned previously, Pindyck (2013) and Morgan et al. (2017) raise valid points about the quantitative implications of integrated assessment models (IAM’s) in part because of the *ad hoc* nature of the damage functions. Most recently, Carleton and Greenstone (2021) have a valuable discussion of what is missed by this approach and the need to incorporate a broader range of empirical evidence. To incorporate such evidence in an internally consistent manner would require much more elaborate models than the ones we have used here. Such efforts would also be accompanied by a rather substantial degree of model uncertainty.

Our specification of damage possibilities included thresholds that could trigger more extreme damages. This formulation is intended to include contributions from environmental tipping points as well as reductions in economic opportunities, which adds to the challenge of how to bound uncertainties in meaningful ways. Productive discussions of our work, while still in progress,

²⁶The National Academies of Sciences, Engineering and Medicine (2017) report suggests investigating sensitivity to discount rates. With discount rates as low as two percent, some of the scenario paths must be projected far out into the future to compute present discounted values. It is well known that present-value calculations can be highly sensitive to assumptions about the distant future. Of course, recursive methods do not escape this challenge but address it by positing dynamic evolutions rather than paths.

included important reminders that the announced thresholds such as a 2 degree or 1.5 degree anomaly have been set as rather arbitrary policy guides and are not firmly grounded with scientific evidence. We find our probabilistic use of such thresholds coupled with robustness concerns to be revealing, but there is little doubt that further modeling and measurement would strengthen the analysis.

Many recent policy discussions have focused on “carbon budgets,” or formal limits on the cumulative amount of CO_2 emissions that should be allowed in the future based on specific temperature thresholds. While carbon budgets allow for simplicity in communication, they are hard to defend in the presence of uncertainty as to the timing and magnitude of temperature responses to carbon emissions. A temperature threshold can be viewed as a very special type of damage function whereby losses are negligible prior to reaching the threshold and very severe (infinite) after the threshold is crossed. But, again, uncertainty about the transition dynamics for future temperature changes as they respond to current emissions requires a commitment to capture carbon in the atmosphere to avoid exceeding the threshold. Though we have already noted the somewhat arbitrary nature of a single temperature threshold, we could approximate socially efficient responses to such a threshold within the framework of our models. This could be done by restricting the upper and lower temperature thresholds to be the same and embracing a steeply curved damage function. The embracing of a single steeply curved damage function could emerge formally as part of a robustness response to damage function uncertainty. Specifically, it could be viewed as the outcome of an extreme form of robustness where degenerate bounds have been placed on the probabilities over the alternative specifications for damage function curvature. Thus, a simple extension to our analysis could approximate a single temperature threshold decision criterion.

Previous commentaries have expressed valid modeling concerns about how to capture adaptation to climate change. As adaptation is inherently a dynamic process, to accommodate it in a meaningful way requires both evidence and a model to interpret the evidence while opening the door to an additional source of model uncertainty. Private sector adaptation to climate change is an important extension of ours in other analyses, but we should also expect policies to adapt to observed changes in the environment that partially resolve the damage function uncertainty. A modular approach that treats uncertainty in emissions trajectories to be independent of damage function uncertainty seems ill-conceived, once we take a more dynamic interactive perspective. Such considerations suggest that sensitivity analyses in the specification of baseline probabilities of the type that we propose would be a revealing way to addressing this potential dependence.

Our model of damage function “learning” is purposefully simplified to illustrate some important points. The stark way in which information is revealed is no doubt extreme, but one needs only to read the recent IPCC Sixth Assessment Report to see claims that observed damages to climate are speeding up our ability to learn about climate change. In our example economy,

in spite of the one-shot nature of the learning, the social planner adopts caution in advance of becoming more fully informed about damage functions. It is too costly to delay action until after the uncertainty is resolved. Also, there is an asymmetric response to the revelation of the damage function curvature. As with any form of rational learning, there is a potential for good news. In this case, the good news is that there might be less damage function curvature than was feared. In our computations, this news response is only really notable for a small portion of the good news realizations. Our purpose in this example is to illustrate the potential importance of going beyond what some in environmental economics, such as Nordhaus (2018), call “learn, then act,” and instead to incorporate informational dynamics into the analysis.

10.4 Technology

One concern from our model is the use of a Cobb-Douglas specification and the energy input share of fossil fuels that we use with that specification. As in Barnett et al. (2020), we use empirical estimates of the energy input share from Finn (1995), adjusted by the approximate proportion of fossil fuel energy consumption, for our value of the energy input share of fossil fuels. This parameter value allows us to match the optimal emissions choice for the solution to the model without climate impacts to the annual emissions of 10 GtC/yr measured by Figueres et al. (2018). In addition, the value is similar to the mean value for the fossil fuel income share found recently by Hassler et al. (2019). Importantly, even with this energy share value, uncertainty is still of first-order importance. Larger energy input share values would serve to amplify the importance of uncertainty for the SCC.

While the Cobb-Douglas specification allows us to simplify the computations and exploit numerically value function separability, a flaw in this specification is that it allows for the planner to reduce energy usages freely without incurring any abatement costs. Our second model follows the DICE literature by introducing explicitly abatement into the production technology whereby output is reduced when deviating from a fixed proportions technology. This allows us to further evaluate the impacts of uncertainty when adjusting carbon emissions is somewhat costly. This model opens the door to more interesting discussions of technological change and exploring the impact of greener technologies that might arrive in the future. As there is uncertainty about the arrival of cleaner technologies, this model includes an additional channel whereby uncertainty has an important impact on the analysis.²⁷

Each of the models we presented was designed to “tell an interesting quantitative story.” While we explored the uncertainty within each setup, we could have had our decision maker entertain both models simultaneously as possibilities. Instead, we find it more productive to push further

²⁷In contrast to our earlier work, Barnett et al. (2020), and some other contributions to the climate-economics literature, we also abstracted from production or resource extraction costs for fossil fuels. This was done for pedagogical simplicity since these costs are typically internalized in the productive process.

the computational boundaries in studying a more ambitious model of technology with interesting special cases.

11 Conclusion

In many dynamic settings, our understanding of the true underlying model relevant for economic decision-making is limited because existing evidence is weak along some important dimensions. As a result, the design and conduct of policy occurs in settings in which policy outcomes are uncertain. We offer the economics of climate change as an example, but there are many others. We turned to decision theory under uncertainty to serve as a guide for how we conduct uncertainty quantification as it contributes to the design of policy. Furthermore, we showed how different forms of uncertainty impact our quantification, how information about environmental and economic damages revealed in the future influence current policies, and how different sources of uncertainty contribute to the SCC in the presence of model and ambiguity and misspecification concerns.

Our analysis in this paper is made simpler here by posing the resource allocation problem as one faced by a single policy maker or social planner. To push closer to a realistic policy setting, multiple decision-makers come into play, including alternative policy makers as well as private sector consumers and investors. Since these different agents confront uncertainty from different perspectives, their uncertainty concerns are expressed in different ways. Moreover, in more realistic policy settings, political constraints prevent first-best solutions. While we fully appreciate the need to extend our analysis of uncertainty to address these modeling challenges, we have little reason to doubt that the uncertainty considerations should remain as first-order concerns and not be shunted to the background as they often are in policy discussions.

References

- Acemoglu, Daron, Ufuk Akcigit, Douglas Hanley, and William Kerr. 2016. Transition to clean technology. *Journal of Political Economy* 124 (1):52–104.
- Anderson, Evan W, Lars Peter Hansen, and Thomas J Sargent. 2003. A quartet of semigroups for model specification, robustness, prices of risk, and model detection. *Journal of the European Economic Association* 1 (1):68–123.
- Arrhenius, Svante. 1896. On the influence of Carbonic Acid in the Air upon Temperatue of the Ground. *Philosophical Magazine and Journal of Science Series* 5 41:237–276.
- Bansal, Ravi and Amir Yaron. 2004. Risks for the Long Run: A Potential Resolution of Asset Pricing Puzzles. *Journal of Finance* 59 (4):1481–1509.
- Barnett, Michael, William A. Brock, and Lars Peter Hansen. 2020. Pricing Uncertainty Induced by Climate Change. *Review of Financial Studies* 33 (3):1024–1066.
- Berger, Loic and Massimo Marinacci. 2020. Model Uncertainty in Climate Change Economics: A Review and Proposed Framework for Future Research. *Environmental and Resource Economics* 1–27.
- Bornstein, Gideon, Per Krusell, and Sergio Rebelo. 2017. Lags, Costs, and Shocks: An Equilibrium Model of the Oil Industry. NBER Working Paper 23423, National Bureau of Economic Research, Inc.
- Borovička, Jaroslav, Lars Peter Hansen, and Jose A. Scheinkman. 2014. Shock Elasticities and Impulse Responses. *Mathematics and Financial Economics* 8 (4).
- Cai, Yongyang and Thomas S Lontzek. 2019. The social cost of carbon with economic and climate risks. *Journal of Political Economy* 127 (6):2684–2734.
- Cai, Yongyang, Timothy M Lenton, and Thomas S Lontzek. 2016. Risk of multiple interacting tipping points should encourage rapid CO₂ emission reduction. *Nature Climate Change* 6 (5):520–525.
- Cai, Yongyang, Kenneth L. Judd, and Thomas S. Lontzek. 2017. The Social Cost of Carbon with Climate Risk. Tech. rep., Hoover Institution, Stanford, CA.
- Carleton, Tamma and Michael Greenstone. 2021. Updating the United States Government’s Social Cost of Carbon. SSRN Working Paper 2021-04, University of Chicago, Becker Friedman Institute for Economics.
- Casassus, Jaime, Pierre Collin-Dufresne, and Bryan R. Routledge. 2018. Equilibrium Commodity Prices With Irreversible Investment and Non-Linear Technologies. *Journal of Banking and Finance* 95:128–147.
- Cerreia-Vioglio, Simone, Lars Peter Hansen, Fabio Maccheroni, and Massimo Marinacci. 2021. Making Decisions under Model Misspecification.
- Dietz, Simon and Frank Venmans. 2019. Cumulative carbon emissions and economic policy: In search of general principles. *Journal of Environmental Economics and Management* 96:108–129.

- Drijfhout, Sybren, Sebastian Bathiany, Claudie Beaulieu, Victor Brovkin, Martin Claussen, Chris Huntingford, Marten Scheffer, Giovanni Sgubin, and Didier Swingedouw. 2015. Catalogue of abrupt shifts in Intergovernmental Panel on Climate Change climate models. *Proceedings of the National Academy of Sciences* 112 (43):E5777–E5786.
- Eby, M., K. Zickfeld, A. Montenegro, D. Archer, K. J. Meissner, and A. J. Weaver. 2009. Lifetime of anthropogenic climate change: Millennial time scales of potential CO₂ and surface temperature perturbations. *Journal of Climate* 22 (10):2501–2511.
- Epstein, Larry G. and Stanley E. Zin. 1989. Substitution, Risk Aversion and the Temporal Behavior of Consumption and Asset Returns: A Theoretical Framework. *Econometrica* 57 (4):937–969.
- Figueres, Christiana, Corinne Le Quéré, Anand Mahindra, Oliver Bäte, Gail Whiteman, Glen Peters, and Dabo Guan. 2018. Emissions Are Still Rising: Ramp Up the Cuts. *Nature* 564:27–30.
- Finn, Mary G. 1995. Variance Properties of Solow’s Productivity Residual and Their Cyclical Implications. *Journal of Economic Dynamics and Control* 19 (5-7):1249–1281.
- Geoffroy, O, D Saint-Martin, D J L Olivié, A Voldoire, G Bellon, and S Tytéca. 2013. Transient Climate Response in a Two-Layer Energy-Balance Model. Part {I}: Analytical Solution and Parameter Calibration Using CMIP5 AOGCM Experiments. *Journal of Climate* 26 (6):1841–1857.
- Ghil, Michael and Valerio Lucarini. 2020. The Physics of Climate Variability and Climate Change. *Reviews of Modern Physics* in press.
- Gillingham, Kenneth, William Nordhaus, David Anthoff, Geoffrey Blanford, Valentina Bosetti, Peter Christensen, Haewon McJeon, and John Reilly. 2018. Modeling Uncertainty in Integrated Assessment of Climate Change: A Multimodel Comparison. *Journal of the Association of Environmental and Resource Economists* 5 (4):791–826.
- Golosov, Mikhail, John Hassler, Per Krusell, and Aleh Tsyvinski. 2014. Optimal Taxes on Fossil Fuel in General Equilibrium. *Econometrica* 82 (1):41–88.
- Good, Irving J. 1952. Rational Decisions. *Journal of the Royal Statistical Society. Series B (Methodological)* 14 (1).
- Hansen, Lars Peter and Jianjun Miao. 2018. Aversion to Ambiguity and Model Misspecification in Dynamic Stochastic Environments. *Proceedings of the National Academy of Sciences* 115 (37):9163–9168.
- Hansen, Lars Peter and Thomas J. Sargent. 2001. Robust Control and Model Uncertainty. *The American Economic Review* 91 (2):60–66.
- . 2007. Recursive Robust Estimation and Control without Commitment. *Journal of Economic Theory* 136 (1):1–27.
- . 2020a. Macroeconomic Uncertainty Prices when Beliefs are Tenuous. *Journal of Econometrics* published online.

- . 2020b. Structured Ambiguity and Model Misspecification. *Journal of Economic Theory* published:105165.
- Hansen, Lars Peter, John C. Heaton, and Nan Li. 2008. Consumption Strikes Back? Measuring Long-Run Risk. *Journal of Political Economy* 116 (2):260–302.
- Hassler, John, Per Krusell, and Conny Olovsson. 2018. The Consequences of Uncertainty: Climate Sensitivity and Economic Sensitivity to the Climate. *Annual Review of Economics* 10:189–205.
- . 2019. Directed technical change as a response to natural-resource scarcity .
- Hausfather, Z and GP Peters. 2020. Emissions – the ‘business as usual’ story is misleading. *Nature* 577 (7792):618–620.
- James, Matthew R. 1992. Asymptotic Analysis of Nonlinear Stochastic Risk-Sensitive Control and Differential Games. *Mathematics of Control, Signals and Systems* 5 (4):401–417.
- Joos, F., R. Roth, J. S. Fuglestad, G. P. Peters, I. G. Enting, W. Von Bloh, V. Brovkin, E. J. Burke, M. Eby, N. R. Edwards, T. Friedrich, T. L. Frölicher, P. R. Halloran, P. B. Holden, C. Jones, T. Kleinen, F. T. Mackenzie, K. Matsumoto, M. Meinshausen, G. K. Plattner, A. Reisinger, J. Segschneider, G. Shaffer, M. Steinacher, K. Strassmann, K. Tanaka, A. Timmermann, and A. J. Weaver. 2013. Carbon Dioxide and Climate Impulse Response Functions for the Computation of Greenhouse Gas Metrics: A Multi-Model Analysis. *Atmospheric Chemistry and Physics* 13 (5):2793–2825.
- Klibanoff, P, M Marinacci, and S Mukerji. 2009. Recursive {Smooth} {Ambiguity} {Preferences}. *Journal of Economic Theory* 144:930–976.
- Koonin, Steven E. 2021. *Unsettled: What Climate Science Tells Us, What It Doesn’t, and Why It Matters*. Dallas, TX: BenBella Books.
- Kreps, David M. and Evan L. Porteus. 1978. Temporal Resolution of Uncertainty and Dynamic Choice. *Econometrica* 46 (1):185–200.
- Lemoine, Derek and Ivan Rudik. 2017. Managing Climate Change Under Uncertainty: Recursive Integrated Assessment at an Inflection Point. *Annual Review of Resource Economics* 9:117–142.
- Lemoine, Derek and Christian P Traeger. 2016. Ambiguous tipping points. *Journal of Economic Behavior & Organization* 132:5–18.
- Lenton, Timothy M. 2020. Tipping positive change. *Philosophical Transactions of the Royal Society B* 375 (1794):20190123.
- Maccheroni, Fabio, Massimo Marinacci, and Aldo Rustichini. 2006. Dynamic Variational Preferences. *Journal of Economic Theory* 128 (1):4–44.
- MacDougall, Andrew H., Neil C. Swart, and Reto Knutti. 2017. The Uncertainty in the Transient Climate Response to Cumulative CO₂ Emissions Arising from the Uncertainty in Physical Climate Parameters. *Journal of Climate* 30 (2):813–827.

- Matthews, H Damon, Nathan P Gillett, Peter A Stott, and Kirsten Zickfeld. 2009. The proportionality of global warming to cumulative carbon emissions. *Nature* 459 (7248):829–832.
- Mills, Mark, Kerry Emanuel, Thomas Frederikse, Zeke Hausfather, Lee Kump, David Lobell, Twila Moon, Frances Moore, Dana Royer, Danel Swain, Scott Wing, and Gary Yohe. 2021. Wall Street Journal article repeats multiple incorrect and misleading claims made in Steven Koonin’s new book ‘Unsettled’. *Climate Feedback* .
- Morgan, M. Granger, Parth Vaishnav, Hadi Dowlatabadi, and Ines L. Azevedo. 2017. Rethinking the Social Cost of Carbon Dioxide. *Issues in Science and Technology* 33 (4).
- National Academies of Sciences, Engineering and Medicine. 2017. *Valuing Climate Damages: Updating Estimation of the Social Cost of Carbon Dioxide*. National Academies Press.
- Nordhaus, William. 2018. Projections and uncertainties about climate change in an era of minimal climate policies. *American Economic Journal: Economic Policy* 10 (3):333–360.
- Nordhaus, William D. 2017. Revisiting the social cost of carbon. *Proceedings of the National Academy of Sciences* 114 (7):1518–1523.
- NYSERDA and RFF. 2021. Estimating the Value of Carbon: Two Approaches. Tech. rep., New York State Energy and Resource Development Authority and Resources for the Future.
- Olson, Roman, Ryan Sriver, Marlos Goes, Nathan M. Urban, H. Damon Matthews, Murali Haran, and Klaus Keller. 2012. A Climate Sensitivity Estimate using Bayesian Fusion of Instrumental Observations and an Earth System Model. *Journal of Geophysical Research Atmospheres* 117 (D04103):1–11.
- Palmer, Tim and Bjorn Stevens. 2019. The scientific challenge of understanding and estimating climate change. *Proceedings of the National Academy of Sciences* 116 (49):24390–24395.
- Pierrehumbert, Ramond T. 2014. Short-Lived Climate Pollution. *Annual Review of Earth and Planetary Science* 42:341–379.
- Pindyck, Robert S. 2013. Climate Change Policy: What Do the Models Tell Us? *Journal of Economic Literature* 51 (3):860–872.
- Ricke, Katharine L. and Ken Caldeira. 2014. Maximum Warming Occurs about One Decade After a Carbon Dioxide Emission. *Environmental Research Letters* 9 (12):1–8.
- Ritchie, Paul DL, Joseph J Clarke, Peter M Cox, and Chris Huntingford. 2021. Overshooting tipping point thresholds in a changing climate. *Nature* 592 (7855):517–523.
- Rogelj, Joeri, Alexander Popp, Katherine V Calvin, Gunnar Luderer, Johannes Emmerling, David Gernaat, Shinichiro Fujimori, Jessica Streffler, Tomoko Hasegawa, Giacomo Marangoni, et al. 2018. Scenarios towards limiting global mean temperature increase below 1.5 C. *Nature Climate Change* 8 (4):325–332.
- Rogelj, Joeri, Daniel Huppmann, Volker Krey, Keywan Riahi, Leon Clarke, Matthew Gidden, Zebedee Nicholls, and Malte Meinshausen. 2019. A new scenario logic for the Paris Agreement long-term temperature goal. *Nature* 573 (7774):357–363.

- Rudik, Ivan. 2020. Optimal Climate Policy When Damages Are Unknown. *American Economic Journal: Economic Policy* 12 (2):340–73.
- Seshadri, Ashwin K. 2017. Fast–slow climate dynamics and peak global warming. *Climate Dynamics* 48 (7-8):2235–2253.
- Sharpe, Simon and Timothy M Lenton. 2021. Upward-scaling tipping cascades to meet climate goals: plausible grounds for hope. *Climate Policy* 1–13.
- Wagner, Gernot and Martin Weitzman. 2015. *Climate Shock*. Princeton University Press.
- Weitzman, Martin L. 2012. GHG Targets as Insurance Against Catastrophic Climate Damages. *Journal of Public Economic Theory* 14 (2):221–244.
- Zickfeld, Kirsten, Michael Eby, Andrew J. Weaver, Kaitlin Alexander, Elisabeth Crespin, Neil R. Edwards, Alexey V. Eliseev, Georg Feulner, Thierry Fichet, Chris E. Forest, Pierre Friedlingstein, Hugues Goosse, Philip B. Holden, Fortunat Joos, Michio Kawamiya, David Kicklighter, Hendrik Kienert, Katsumi Matsumoto, Igor I. Mokhov, Erwan Monier, Steffen M. Olsen, Jens O.P. Pedersen, Mahe Perrette, Gwenaëlle Philippon-Berthier, Andy Ridgwell, Adam Schlosser, Thomas Schneider Von Deimling, Gary Shaffer, Andrei Sokolov, Renato Spahni, Marco Steinacher, Kaoru Tachiiri, Kathy S. Tokos, Masakazu Yoshimori, Ning Zeng, and Fang Zhao. 2013. Long-Term Climate Change Commitment and Reversibility: An EMIC Intercomparison. *Journal of Climate* 26 (16):5782–5809.

Appendix A Carbon and Temperature Model Sets

As mentioned previously, we use 16 models of temperature dynamics from Geoffroy et al. (2013) and nine models of carbon dynamics models from Joos et al. (2013). We briefly describe the model experiments used in these papers, list the models we include in our analysis, and provide details for the reader to find additional information about these models and model experiments.

Geoffroy et al. (2013) approximates the temperature dynamics of 16 different models using a two-layer energy-balance model (EBM) to study properties of atmosphere-ocean general circulation models (AOGCMs). Table 1 lists the model name for each of the 16 models used in their and our analysis and direct the reader to Geoffroy et al. (2013) and Seshadri (2017) for additional details about each of the models.

The Geoffroy et al. (2013) EBM model uses the following specification:

$$\begin{aligned} c_s \frac{dT^s}{dt} &= F - \gamma T^s - \epsilon \chi (T^s - T^o) \\ c_o \frac{dT^o}{dt} &= -\chi (T^o - T^s) \\ F &= 5.35 (\log CO_2 - \log \underline{CO_2}) \end{aligned}$$

where T^s is the surface temperature, T^o is the ocean temperature, CO_2 is atmospheric carbon dioxide, and $\underline{CO_2}$ is the preindustrial benchmark. The construction of F comes from the “Arrhenius” equation (Arrhenius (1896)). The EBM model is solved for explicit solutions, calibrated to fit the responses of 16 AOGCMs that participated in the CMIP5, and then validated by using the AOGCM responses to the linear forcing experiments of one percent of CO_2 per year. The parameters they estimate in this simplified representation differ depending on the model used in the calibration of the approximation, providing a measure of the heterogeneity and uncertainty present in models of temperature dynamics. We use this specification along with Geoffroy et al.’s estimates of the 16 temperature dynamics models in our simulations to capture the carbon-to-temperature component of climate model uncertainty.

Joos et al. (2013) use a carbon cycle-climate model intercomparison analysis to study the impulse response time-scales of Earth System models. From their analysis, we use the impulse response functions of nine models based on a 100GtC emission pulse added to a constant CO_2 concentration of 389 ppm.²⁸ All of the models we use are Earth System Models of Intermediate Complexity, except for the reduced form model Bern-SAR. We list the model name for each of the models used in our analysis in Table 2 below. We direct the reader to Appendix A in Joos et al. (2013) for detailed descriptions of these and other models used in their intercomparison analysis.

²⁸We thank Fortunat Joos for graciously providing the data for these and other response experiments on his website: https://climatehomes.unibe.ch/~joos/IRF_intercomparison/results.html.

Temperature Dynamics Models
BCC-CSM1-1
BNU-ESM
CanESM2
CCSM
CNRM-CM5
CSIRO-Mk3.6.0
FGOALS-s2
GFDL-ESM2M
GISS-E2-R
HadGEM2-ES
INM-CM4
IPSL-CM5A-LR
MIROC
MPI-ESM-LR
MRI-CGCM3
NorESM1-M

Table 1: List of temperature dynamics models from Geoffroy et al. (2013) and Seshadri (2017) used in our analysis.

Carbon Dynamics Models
Bern3D-LPJ (reference)
Bern2.5D-LPJ
CLIMBER2-LPJ
DCESS
GENIE (ensemble median)
LOVECLIM
MESMO
UVic2.9
Bern-SAR

Table 2: List of carbon dynamics models from Joos et al. (2013) used in our analysis.

Appendix B Value Function Components for Section 7 Model

In Section 7, we discussed a climate economics HJB equation in the state variable y . This HJB equation uses a quasi-analytical simplification for the damages state n of the form $\phi(y) - \frac{1-\eta}{\delta}n$, which derive using the “guess and verify” method. This is part of a larger system that can be solved with two additional subsystems of equations. The three subsystem solutions, when combined, give a solution to the composite HJB equation of the planner.

B.1 Climate economics system parameters

Parameter	Value
ς'	[2.23 0 0]
γ_1	.000177
γ_2	.0044
γ_3^m	$\frac{.333(m-1)}{19}$, $m=1,2,\dots,20$
η	.032
δ	.01

Table 3: Climate economics system parameters

To understand better the implications of the ς specification note that for a constant emissions path, the implied standard deviation associated with the coefficient of the Matthew’s approximation is .446 at twenty five years, .315 for fifty years and .223 for one hundred years.

B.2 Consumption-capital dynamics

The undamaged version of the consumption capital model, by design, has a straightforward solution. We use the “guess and verify” method to derive a solution for this subsystem, guessing a value function of $v_k \log k + \zeta(z)$. The HJB equation for this component is:

$$\begin{aligned}
0 = & \max_i \min_h -\delta [v_k \log k + \zeta(z)] + (1 - \eta) [\log (\alpha - i) + \log k] + \frac{\xi_r}{2} |h|^2 \\
& + v_k \left[\mu_k(z) + i - \frac{\kappa}{2} (i)^2 + \sigma_k(z)'h - \frac{|\sigma_k(z)|^2}{2} \right] \\
& + \frac{\partial \zeta}{\partial z}(z) [\mu_z(z) + \sigma(z)'h] + \frac{1}{2} \text{trace} \left[\sigma(z)' \frac{\partial^2 \zeta}{\partial z \partial z'}(z) \sigma(z) \right]
\end{aligned}$$

From this equation, we derive the constant scaling the capital component of the value function v_k and can see that it must be

$$v_k = \frac{1 - \eta}{\delta}.$$

Solving for the first-order conditions, we see that the first-order condition for h is

$$\xi_r h + \sigma_k v_k + \sigma_z \frac{\partial \zeta}{\partial z} = 0,$$

and the first-order conditions for the investment-capital ratio is

$$-(1 - \eta) \left(\frac{1}{\alpha - i} \right) + v_k (1 - \kappa i) = 0.$$

Notice that the equation for the optimal h is therefore

$$h = -\frac{1}{\xi_r} \left[\sigma_k v_k + \sigma_z \frac{\partial \zeta}{\partial z} \right]$$

and that the investment-capital ratio is constant. While there are two solutions for the first-order conditions for i , only one is positive. In our illustration we set $\alpha = .115$ and $\kappa = 6.667$.

The solution for h will be state dependent if we allow for σ_k or σ_z to depend on z or if there is nonlinearity in the drift specifications. Such dependence is common in the macro finance literature as a form of stochastic volatility. In the computations that follow, we will abstract from this dependence and impose linear dynamics for z . We impose that

$$\mu_k(z) = -.043 + .04(\iota_k \cdot z)$$

and

$$\sigma_k = .01 \begin{bmatrix} .87 & .38 \end{bmatrix} dW_t^k$$

where dW^k is a two dimensional subvector of the Brownian increment vector dW . The evolution for the process $\iota_k \cdot Z$ is given by a continuous-time autoregression:

$$d(\iota_k \cdot Z_t) = -.056(\iota_k \cdot Z_t) dt + \begin{bmatrix} 0 & .055 \end{bmatrix} dW_t^k$$

In this case $\zeta(z) = \zeta_0 + \zeta_1 \iota_k \cdot z$ where ζ_1 satisfies

$$-\delta \zeta_1 + v_k(.04) + \zeta_1(-.056) = 0.$$

The implied solution for h is constant and equal to:

$$h^* = -\frac{1}{\xi_r} \begin{bmatrix} .85 \\ 3.58 \end{bmatrix}$$

The implied consumption dynamics in this setting are consistent with the ones given in Hansen

and Sargent (2020a):²⁹

$$d \log C_t = .0194 + .04 Z_t dt + .01 \begin{bmatrix} .87 & .38 \end{bmatrix} \cdot dW_t^k.$$

B.3 Contribution of $\iota_y \cdot z$

There is one remaining contribution to the planner's HJB equation for each of our models. Note that while $\log \iota_y \cdot z$ is included in the objective of the planner, this term has not been accounted for in our solution so far. Thus there is a third contribution, $\tilde{\zeta}$ to the value function that solves:

$$\begin{aligned} \min_h -\delta \tilde{\zeta}(z) - \eta \log(\iota_y \cdot z) + \left[\frac{\partial \tilde{\zeta}}{\partial z}(z) \right] \cdot [\mu_z(z) + \sigma_z(z)h] + \frac{\xi_r}{2} h' h \\ + \frac{1}{2} \text{trace} \left[\sigma_z(z)' \frac{\partial^2 \tilde{\zeta}}{\partial z \partial z'}(z) \sigma_z(z) \right] = 0. \end{aligned} \quad (13)$$

To support this value function separation we impose that $\iota_y \cdot Z$ and $\iota_k \cdot Z$ are independent processes with $\iota_y \cdot Z$ constructed as a function of the dW^y increments and $\iota_k \cdot Z$ constructed in terms of the dW^k increments. Moreover, we assume that

$$\varsigma' \sigma_z(z)' \left[\frac{\partial \tilde{\zeta}}{\partial z}(z) \right] = 0 \quad (14)$$

where $\tilde{\zeta}$ is the solution to HJB equation $\left[\frac{\partial \tilde{\zeta}}{\partial z}(z) \right]$.

As a special case suppose that $\iota_y \cdot Z_t$ evolves as Feller square root process with mean one:

$$d(\iota_y \cdot Z_t) = -\chi(\iota_y \cdot Z_t - 1)dt + \sqrt{\iota_y \cdot Z_t} \tilde{\varsigma} \cdot dW_t^y.$$

where $\tilde{\varsigma} \cdot \varsigma = 0$. Then the solution of interest to equation (13) can be expressed as a functional equation in the scalar argument $\iota_y \cdot z$. Given the separability, this value function contribution is used for the figures that we produce.

As part of a “guess and verify” solution method we add the three value function components and the three components for the minimizing h together along with the proposed solutions for the investment capital ratio i and for scaled emissions $\tilde{e}$. In fact there may be good reasons to relax assumption (14) and combine the climate economics HJB contribution and that coming from (13) into a single HJB equation to be solved instead of two lower-dimensional functional equations.

²⁹Hansen and Sargent (2020a) represent the dynamics in terms of a time unit of one quarter instead of one year, and they report a different but observationally equivalent orthogonalization of the Brownian increments.

Appendix C Value Function Components for Section 9 Model

The HJB equation for the model in Section 9 depends on the state variables y and k . As with the previous model, this HJB equation uses a quasi-analytical simplification for the damages state n , as well as a separable subsystem for the exogenous forcing state z . The value function that solves the HJB equation is of the form $\varphi(y, k) - \frac{1}{\delta}n + \zeta(z)$, which is derived using the “guess and verify” method. The two subsystem solutions, one for $\varphi(y, k) - \frac{1}{\delta}n$ and one for $\zeta(z)$, when combined, give a solution to the composite HJB equation of the planner.

C.1 HJB equations details

As was the case for the model in Section 7, this model has pre- and post-jump values functions. These HJB equations are similar in structure to those shown in the previous model. However, in this case there are additional layers for two reasons: (i) the potential for two different technology jumps related to the abatement technology; (ii) the lack of separability between y and k due to the functional form of abatement technology. As a result, we must compute numerous continuation value functions based on post-jump outcomes across multiple dimensions, as well as a pre-jump value function. We denote the pre-damage jump value functions as $\varphi_i(y, k)$ and the post-damage jump value functions as $\varphi_{i,m}(y, k)$ where i denotes the number of technology jumps that have already occurred, and thus the values for (λ_t, ϑ_t) , so that $i \in \{0, 1, 2\}$. We denote the intensity rate for each technology jump as $\mathcal{H}$, given that it is constant and equal for each jump scenario.

C.2 Additional parameters values and initial conditions

We provide a table of the consumption-capital parameters and the abatement technology parameters here. Except for the parameters pertaining to abatement, which were not included in the previous model, the parameters for this model match those given in Appendix B.

Parameter	Value
θ	3.0
(ϑ_0, λ_0)	(.0453, .1206)
$\mathcal{H}$	.05

Table 4: Abatement parameters

The initial values for the abatement technology (ϑ_0, λ_0) are based on the implied values for 2020 from Cai and Lontzek (2019). We set the initial value of capital so that our initial GDP matches the 2020 World GDP value of \$85 trillion in the World Bank National Accounts data.

Therefore, $K_0 = 739.13$. We set the initial value of atmospheric temperature anomaly to match recent estimates provided by the IPCC. Therefore, $Y_0 = 1.1$ degrees Celsius.