Confronting Uncertainty in the Climate Change Dynamics*

Mike Barnett‡ William Brock* Lars Peter Hansen**

May 27, 2021

Abstract

In designing policy to combat climate change, there are many calls for immediate action. However, there has also been some apprehension expressed about acting prematurely based on existing limits to our understanding of the timing and the quantitative magnitude of climate change and its impact on economic and social outcomes. We approach this problem from the perspective of a decision maker who confronts uncertainty in a setting in which there will be future information about damage severity. The value of further empiricism in the near term will be limited as the climate-economic system moves into uncharted territory. The decision maker's concerns about uncertainty go beyond the standard risk considerations. There is ambiguity over how much credibility to assign to alternative models of climate change and economic damages and all such models are potentially misspecified. These uncertainties add to the challenge of projecting how economic damages in the future will be enhanced with current and future fossil fuel emissions. The solution to the decision problem we pose shows some initial caution until future damage possibilities are more fully revealed. With the more refined information, the decision maker may be either more wary of climate change, or more bullish, and act accordingly depending on what is revealed. We provide a quantitative illustration drawing in part on cross-model ambiguity revealed by pulse experiments conducted by geoscientists.

*University of Wisconsin and University of Missouri, Columbia

**University of Chicago

‡Arizona State University

*We thank Shirui Chen and Han Xu for the computational support on this research, and Diana Petrova and Zhenhuan Xie for the help in preparing this manuscript. Financial support for this project was provided by the Alfred P. Sloan Foundation [grant G-2018-11113].

1 Introduction

Climate change has attracted the attention of many players in the policy arena, as well as academics. While there is much agreement about the human impact on the climate system and the potential that it has for hindering economic opportunities and social well-being, the evidence from geosciences and from economics points to substantial uncertainty. There is uncertainty in both the quantitative impacts and in the time horizons over which these impacts become prominent. The uncertainties are in effect "multiplicative" working through both the carbon-climate dynamics and the damages to the economy induced by climate change. We are still only beginning to understand the many facets of uncertainty and how they should influence the economic policy aimed at confronting climate change. This paper features one such facet by studying the implications for policy making when we allow for this uncertainty to play out in an explicitly dynamic fashion. Specifically, as we increase fossil fuel emissions, the long-term consequences of climate change will become all the more evident. While more information will be revealed in the future, delaying action can be costly because one of the potentially bad scenarios with large social consequences might become a reality. It is this dynamic aspect of uncertainty that we feature in our analysis.

We study the implications of uncertainty in the highly stylized setting of a fictitious social planner confronting climate change on a global scale. There is, of course, a sizable gap between this setting and realistic policy-making in multiple political arenas. We are adopting this as a starting point for a better understanding of the explicit role of uncertainty for climate change policy. Like elsewhere in the climate-economics literature, we focus on the social cost of carbon (SCC), the marginal net social cost of emissions implied by the solution to the planner's problem. Since there is a climate change externality, social and private costs can differ in a substantial way. We use the SCC to assess the importance of uncertainty in designing prudent climate policy. We view the SCC through the lens of asset pricing along the lines of our earlier work in Barnett et al. (2020), but with a more dynamic perspective on how the uncertainty becomes realized. Asset pricing tools are pertinent because emissions today have consequences to economic and social well-being well into the future. We also embrace a broader notion of uncertainty as advocated by Barnett et al. by pushing outside the straightjacket of risk aversion and incorporating ambiguity aversion over alternative geo-scientific predictions for climate and aversion to the misspecification of damages that result from climate change.

In our example economy, we include four sources of uncertainty:

- *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 re-

duced by temperature changes

- *economic damage thresholds* that determine the temperature value at which potentially more severe damages begin.

The damage uncertainties in the third and fourth bullet points are resolved as a Poisson event whose likelihood increases depending upon the temperature anomaly. This formulation allows us to connect to some of the important climate change policy discussions that have arisen as researchers speculate about tipping points and the possibly severe damages that might emerge. In our setup, earlier realizations of the Poisson event are bad news about damages and later ones are, relatively speaking, good news. Emissions adjust discontinuously to the Poisson event. While the planner has a baseline specification for the Poisson arrival and the potential damages that might follow, he or she only treats this baseline as an approximation and makes a robustness adjustment when designing a decision rule for emissions.

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 are articulating 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 the 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.[1]

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 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 adopt a simplified stochastic specification of the pulse responses developed in Barnett et al. (2021)

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.[2]

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,

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

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

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 the appendix. 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.

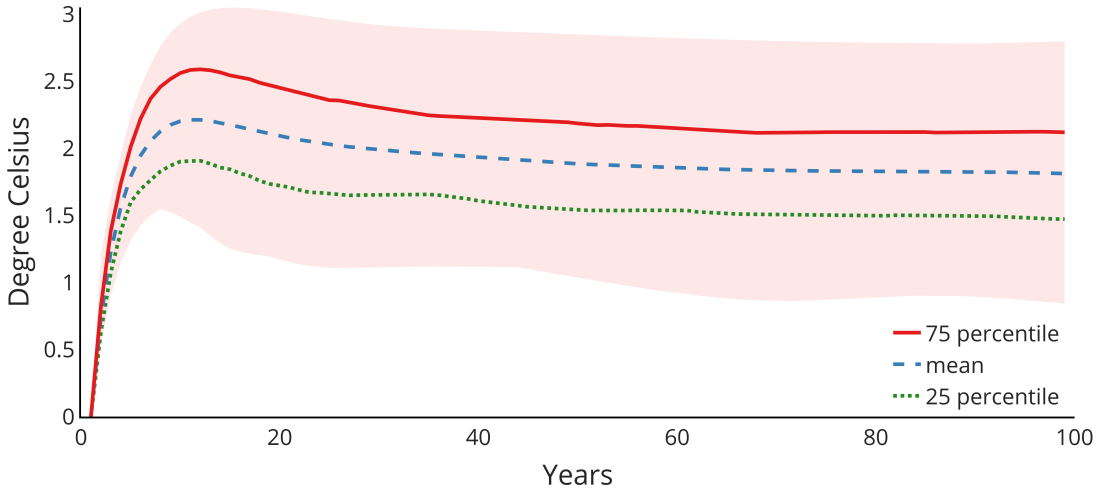

Figure 1: Percentiles for temperature responses to emission impulses. The emission pulse was 100 gigatons of carbon (GtC) spread over the first year. The temperature units for the vertical axis have been multiplied by ten to convert to degrees Celsius per teraton of carbon (TtC). The boundaries of the shaded regions are the upper and lower envelopes.

Figure 1 captures the resulting temperature responses across various sets of these 144 models. In each case, the maximal temperature response to an 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). 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. 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.

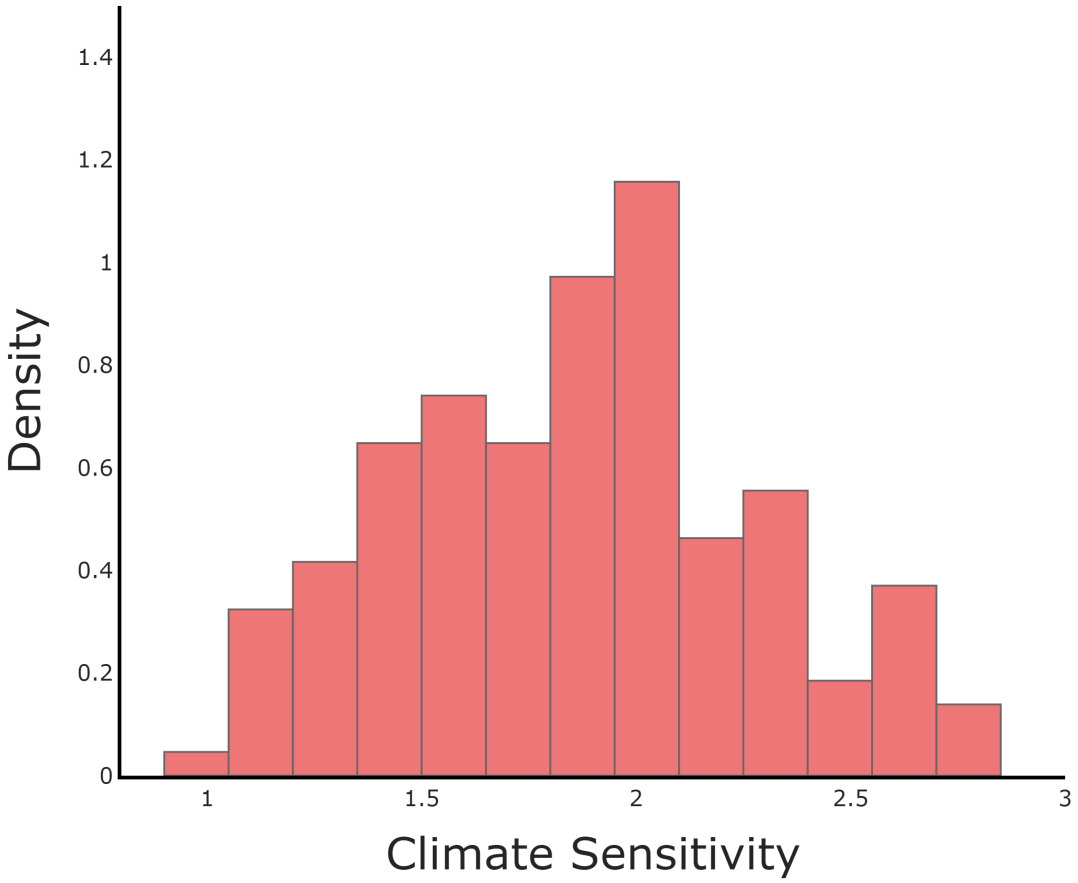

Figure 2: Histograms for the exponentially weighted average responses of temperature to an emissions impulse from 144 different models using a rate $\delta = .01$.

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.

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 processes that evolves as:

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

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

$$dW_t = \begin{bmatrix} dW_t^y \\ dW_t^k \end{bmatrix}$$

where the first component consists of the climate change shocks and the second component contains the technology shocks. Consider an emissions "pulse" of the form

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

where $\mathcal{E}_t$ is fossil fuel emissions and $\iota_y \cdot Z = \{\iota_y \cdot Z_t : t \geqslant 0\}$ is a positive process which we normalize to have mean one. The $\iota_y \cdot Z$-process captures "left out" components of the climate systems reaction to an emission of $\mathcal{E}_t$ gigatons into the atmosphere while the $\varsigma \cdot dW$ process captures short time scale fluctuations. We will use a positive Feller square root process for the $\iota_y \cdot Z$ process in our analysis.

Within this framework, we impose the "Matthews' approximation" by making the consequence of the pulse permanent:

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

where:

$$\mu_y(z, e) = e\, (\iota_y \cdot z)\, \theta$$
$$\sigma_y(z, e) = e\, (\iota_y \cdot z)\, \varsigma'$$

Throughout, we will use uppercase letters to denote random vector or stochastic processes and lower case letters to denote possible realizations. Armed with this "Matthew's 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. Barnett et al. (2021) found that when the prudent social planner has preferences with a low rate of discounting in the future, this initial buildup has very little impact on the policy prescription. 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.

4 A stochastic model of damages

We posit a damage process, $N_t = \{N_t : t \geqslant 0\}$ 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 implement these extensions using a piecewise log-quadratic function for mapping how temperature changes induced by emissions alter economic opportunities:

$$\log N_t = \begin{cases} \Gamma(Y_t) + \iota_n \cdot Z_t^w & Y_t < Z_t^p \\ \Gamma(Y_t - Z_t^p + \overline{y}) + \iota_n \cdot Z_t^w & Y_t \geqslant Z_t^p \end{cases} \tag{1}$$

for temperature anomaly Y_t where:

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

In this specification, 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. Let Z_t^p be the value of the temperature anomaly at the date of the Poisson event. After the jump, the threshold process, $Z^p \doteq \{Z_t^p : t \geqslant 0\}$, remains constant; but from the jump point on, the damages, as a function of the temperature anomalies, become more severe.

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. We model the Poisson intensity that governs the jump probability as an increasing function of the temperature anomaly. We specify it so that the Poisson event is triggered prior to the anomaly hitting the threshold $\overline{y}$. To model the Poisson event, we posit a jump process with m absorbing states. Each state corresponds to a value of γ_3 with possible values γ_3^m for $m = 1, 2, ..., M$. In what follows, we refer

to $\overline{y}$ as the upper threshold. The coefficients γ_3^m are specified so that the proportional damages are equally spaced after the threshold $\overline{y}$.

For a Poisson jump process, the intensity function, $\mathcal{I}$, determines, locally, the possibility of a jump over the next increment in time. For $Y_t = y$, $\epsilon\mathcal{I}(y)$ is the approximate jump probability over small time increment ϵ. In our computations, we use intensity function

$$\mathcal{I}(y) = \begin{cases} \mathsf{r}_1 \left(\exp\left[\frac{\mathsf{r}_2}{2} (y - \underline{y})^2 \right] - 1 \right) & y \geqslant \underline{y} \\ 0 & 0 \leqslant y < \underline{y} \end{cases}$$

as depicted in Figure 3. As $\mathcal{I}$ is increasing y, jumps become all the more likely as Y_t approaches the upper threshold $\overline{y}$. This intensity depends on $\underline{y}$, which we refer to as the lower threshold. We set the values of $(\mathsf{r}_1, \mathsf{r}_2)$ so that the probability that the Poisson event is realized prior to $Y_t = \overline{y}$ is essentially unity. Thus, the uncertainty is concentrated for state Y in the interval $[\underline{y}, \overline{y}]$. We use the intensity plotted in Figure 3 in computations that follow for $\underline{y} = 1.5$ and $\overline{y} = 2$.

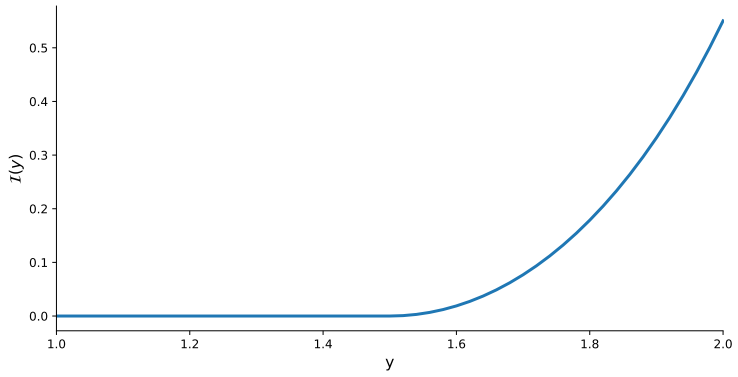

Figure 3: Intensity function, $\mathsf{r}_1 = 1.5$ and $\mathsf{r}_2 = 2.5$. With this intensity function, the probability of a jump at an anomaly of 1.6 is approximately .02 per annum, increasing to about .08 per annum at an anomaly of 1.7, increasing further to approximately .18 per annum at an anomaly of 1.8 and then to about one third per annum when the anomaly is 1.9.

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 function beyond $\overline{y}$. 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

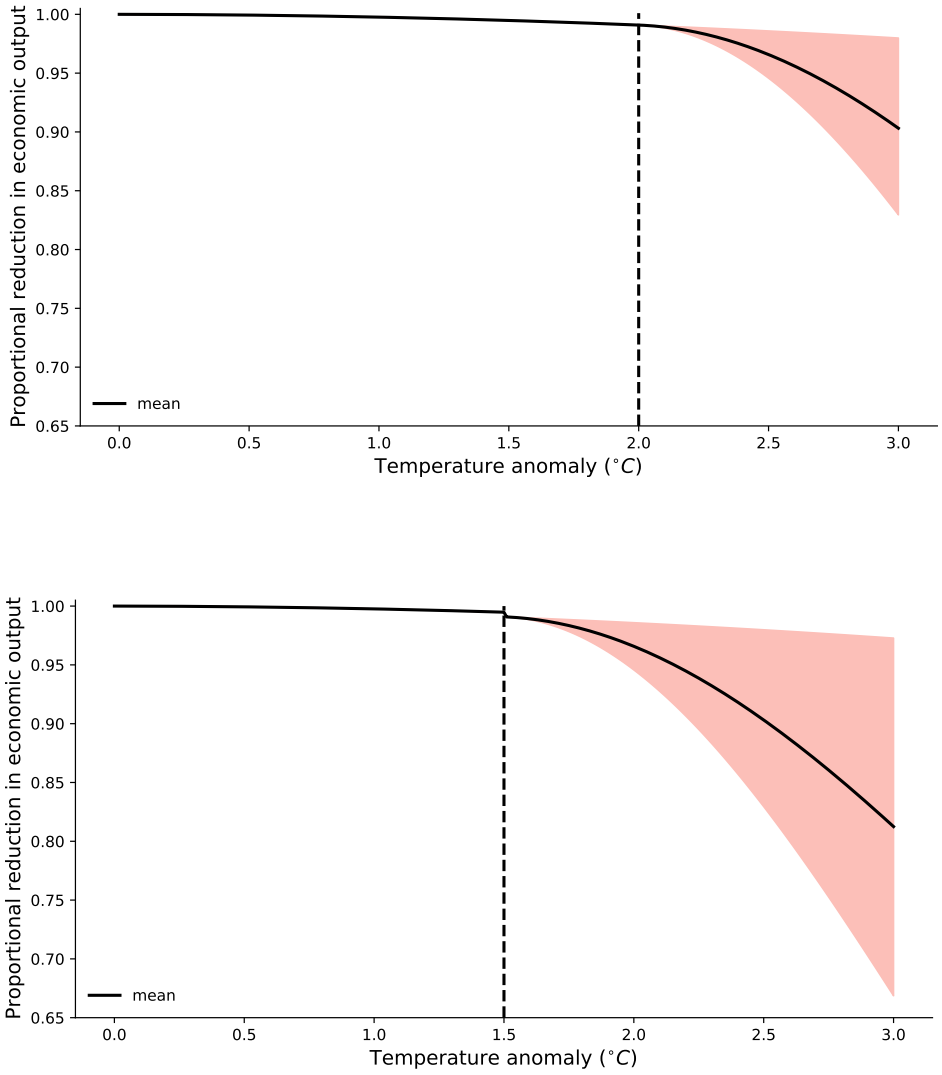

Figure 4: Range of possible damage functions for different jump thresholds. The shaded regions in these plots gives the range of possible values for exp(-n), which measures the proportional reduction of the productive capacity of the economy. The top panel shows the damage function curvature when the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

jumps are "good news."

Remark 4.1. *Our choice of 1.5 and 2 degree celsius thresholds for the temperature anomaly is motivated by discussions in the climate science literature. 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 that 1.5 degree temperature increase as a goal for limiting human damages of climate system. As this goal may be unachievable, they point to 2 degree target in line with 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.*

5 An illustrative economy

To illustrate how the perceived unraveling of uncertainty influences prudent decisions, we use a stylized economic setting that allows for tractability and clarity. (See Barnett et al. (2021) for more complete economic frameworks that allow us to explore additional climate policy components.) We start by specifying the economy in the absence of environmental damages. We pose 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. 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. Given the output constraint and capital evolution imposed by the AK technology, it suffices to let the planner choose the investment-capital ratio.

Formally, "undamaged" capital evolves as

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

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

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

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 temperature shifts proportion-

ately consumption and capital by a multiplicative factor N_t that captures damages to the productive capacity induced by climate change. For instance, the damage adjusted consumption is $\widetilde{C}_t = \frac{C_t}{N_t}$ and the damage adjusted capital is $\widetilde{K}_t = \frac{K_t}{N_t}$.

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:

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

where we will denote the subjective rate of discount used in preferences as δ.

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 5.1. *As Barnett et al. (2021) note, 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 5.2. *We obtain a further simplification by letting:*

$$\widetilde{\mathcal{E}}_t = \mathcal{E}_t\left(\iota_y \cdot Z_t^w\right).$$

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

6 HJB equations and robustness

We now describe our approach to solving the model and incorporation of concerns about robustness and ambiguity aversion. The uncertainty that we consider has a single jump point after which the damage function uncertainty is known. 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.

6.1 Post jump continuation value functions

We first compute the value functions pertinent after the Poisson event is realized. At this juncture, the damage specification uncertainty has been resolved but not the climate model uncertainty.

The state variables are the temperature anomaly and the exogenous Brownian uncertainty. While the argument of the function Γ depends on $Y_t - Z_t^p + \overline{y}$, post-jump Z_t^p is constant and only Y_t varies. We abuse notation a little bit by letting y denote a potential realization of $Y_t - Z_t^p + \overline{y}$.

We solve the optimization problems for the continuation value functions ϕ_m conditioned on each of the damage functions, $m = 1, 2, ..., M$. The ϕ_m are then computed with a threshold of $\overline{y}$, still conditioning on each of the damage functions, $m = 1, 2, ..., M$. The optimization problem is well-posed whenever $y \geqslant 0$, but we only care about the range in which $y \geqslant \overline{y}$.

Conditioned on damage specification m, and abstracting ambiguity aversion over the L climate models, the HJB equation for each conditional ϕ_m is given by:

$$\begin{aligned} 0 = \max_{\tilde{e}} & -\delta\phi_m(y) + \eta \log \tilde{e} \\ & + \frac{d\phi_m(y)}{dy} \sum_{\ell=1}^{L} \pi_\ell^a \theta_\ell \tilde{e} + \frac{1}{2} \frac{d^2\phi_m(y)}{(dy)^2} |\varsigma|^2 \tilde{e}^2 \\ & + \frac{(\eta - 1)}{\delta} \left([\gamma_1 + \gamma_2 y + \gamma_3^m (y - \overline{y})] \sum_{\ell=1}^{L} \pi_\ell^a \theta_\ell \tilde{e} + \frac{1}{2} (\gamma_2 + \gamma_3^m) |\varsigma|^2 \tilde{e}^2 \right). \end{aligned} \tag{2}$$

The threshold value we used to compute each continuation value function ϕ_m is $\bar{y} = 2$. 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. Robustness considerations can be included as before.

To incorporate ambiguity aversion, we follow Hansen and Miao (2018) by embedding a minimization problem in the HJB equation (2)

$$\begin{aligned} 0 = \max_{\tilde{e}} \min_{\omega_j, \sum_{\ell=1}^{L} \omega_\ell = 1} & -\delta\phi_m(y) + \eta \log \tilde{e} \\ & + \frac{d\phi_m(y)}{dy} \sum_{\ell=1}^{L} \omega_\ell \theta_\ell \tilde{e} + \frac{1}{2} \frac{d^2\phi_m(y)}{(dy)^2} |\varsigma|^2 \tilde{e}^2 \\ & + \frac{(\eta - 1)}{\delta} \left([\gamma_1 + \gamma_2 y + \gamma_3^m (y - \overline{y})] \sum_{\ell=1}^{L} \omega_\ell \theta_\ell \tilde{e} + \frac{1}{2} (\gamma_2 + \gamma_3^m) |\varsigma|^2 \tilde{e}^2 \right) \\ & + \xi_a \sum_{\ell=1}^{L} \omega_\ell \left(\log \omega_\ell - \log \pi_\ell^a \right) \end{aligned} \tag{3}$$

where the last term penalizes the search over alternative probabilities to assign to the L different climate models. The coefficient ξ_a is a penalty parameter. By setting this parameter to be sufficiently large, we approximate the solution without ambiguity aversion. For a given ω, the

minimization over $\tilde{e}$ has a quadratic objective and can be solved quasi-analytically. For a given $\tilde{e}$, the minimization problem can also be solved analytically, with:

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

We exploit the resulting formulas in our recursive computations. These give us M continuation value functions that we use in our analysis.

6.2 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. The first modification leads to:

$$\begin{aligned}
0 = \max_{\tilde{e}} \min_{\omega_j, \sum_{\ell=1}^L \omega_\ell = 1} & -\delta\phi(y) + \eta\log\tilde{e} \\
& + \frac{d\phi(y)}{dy}\sum_{\ell=1}^L \omega_\ell\theta_\ell\tilde{e} + \frac{1}{2}\frac{d^2\phi(y)}{(dy)^2}|\varsigma|^2\tilde{e}^2 \\
& + \frac{(\eta-1)}{\delta}\left(\left[\gamma_1+\gamma_2 y\right]\sum_{\ell=1}^L \omega_\ell\theta_\ell\tilde{e} + \frac{1}{2}\gamma_2|\varsigma|^2\tilde{e}^2\right) \\
& + \xi_a\sum_{\ell=1}^L \omega_\ell\left(\log\omega_\ell - \log\pi_\ell^a\right) \\
& + \mathcal{I}(y)\sum_{m=1}^M \pi_m^p\left[\phi_m(\overline{y}) - \phi(y)\right]. \qquad (4)
\end{aligned}$$

The continuation value functions ϕ_m are all evaluated at $y = \overline{y}$ 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 as is displayed in (1).

6.3 Robustness in the Poisson event specification

As we have seen, a jump process specification includes both a) a state-dependent intensity, $\mathcal{I}$, governing the probability of a jump and b) the distribution over the post-jump states, as depicted by π_3^m's. Both of these could be mistaken. We incorporate the two forms of potential misspecification in our setting as follows. Potential misspecification of the distribution over post jump states is captured by introducing positive random variables $G_t^m \geqslant 0$ for each alternative damage

model m with local evolution of the state given by

$$\mathcal{I}(Y_t) \sum_{m=1}^{M} G_t^m \pi_m^p \left[\phi_m(\overline{y}) - \phi(y)\right]$$

The altered probabilities resulting from the concern about misspecification are given by:

$$G_t^m \pi_m^p, \qquad \sum_{m=1}^{M} G_m \pi_m^p = 1, \qquad m = 1, 2, ..., M.$$

Potential misspecification of the probability of a jump is captured by introducing the positive random variable $H_t \geqslant 0$ for our given intensity function $\mathcal{I}(Y_t)$. The altered intensity function resulting from the concern about this type of misspecification is given by: $H_t \mathcal{I}(Y_t)$.

We follow Anderson et al. (2003) by restraining the search over potential misspecifications with a relative entropy penalty:

$$\mathcal{I}(Y_t) \left[1 - H_t + H_t \log H_t + H_t \sum_{m=1}^{M} \pi_m^p G_t^m \log G_t^m\right] \geqslant 0$$

This measure is nonnegative, in part, because the convex function $h \log h$ exceeds its linear gradient approximation, $h-1$, where the gradient evaluated at $h = 1$ is unity. Moreover, Jensen's Inequality guarantees that:

$$\sum_{m=1}^{M} \pi_m^p G_t^m \log G_t^m \geqslant 0$$

To determine a local contribution to an HJB equation, we follow Anderson et al. (2003) by solving:

$$\min_{h \geqslant 0} \min_{g_m : \sum_{m=1}^{M} g_m \pi_m^p = 1} \mathcal{I} h \sum_{m=1}^{M} g_m \pi_m^p \left[\phi_m\left(\overline{y}\right) - \phi\right] + \xi_p \mathcal{I} \left[(1 - h + h \log h) + h \sum_{m=1}^{M} \pi_m^p g_m \log g_m\right]$$

where $\xi_p > 0$ is a penalty parameter limiting the search over the G_t^m's and H_t with realized values given by the g_m's and h.

We add this outcome to the HJB equation of the decision maker to adjust for robustness to

jumps misspecification, which gives the following HJB equation:

$$
\begin{aligned}
0 = \max_{\tilde{e}} \min_{h} & \min_{g_m, \sum_{m=1}^{M} g_m \pi_m = 1} \quad \min_{\omega_j, \sum_{\ell=1}^{L} \omega_\ell = 1} \quad -\delta\phi(y) + \eta \log \tilde{e} \\
& + \frac{d\phi(y)}{dy} \sum_{\ell=1}^{L} \omega_\ell \theta_\ell \tilde{e} + \frac{1}{2} \frac{d^2\phi(y)}{(dy)^2} |\varsigma|^2 \tilde{e}^2 \\
& + \frac{(\eta - 1)}{\delta} \left([\gamma_1 + \gamma_2 y] \sum_{\ell=1}^{L} \omega_\ell \theta_\ell \tilde{e} + \frac{1}{2} \gamma_2 |\varsigma|^2 \tilde{e}^2 \right) \\
& + \xi_a \sum_{\ell=1}^{L} \omega_\ell \left(\log \omega_\ell - \log \pi_\ell^a \right) \\
& + \mathcal{I}(y) h \sum_{m=1}^{M} g_m \pi_m^p \left[\phi_m \left(\overline{y} \right) - \phi(y) \right] \\
& + \xi_p \mathcal{I}(y) h \sum_{m=1}^{M} \pi_m^p \left(g_m \log g_m \right) \\
& + \xi_p \mathcal{I}(y) \left(1 - h + h \log h \right).
\end{aligned} \tag{5}
$$

Minimizing over the g_m's, we find that:

$$
g_m^* \;\propto\; \exp\left(\frac{1}{\xi_p} \left[\phi - \phi_m \left(\overline{y} \right) \right] \right),
$$

which do not depend directly on the intensity $\mathcal{I}$. The relation is a proportional one because the right-hand side needs to be scaled so that g_m^*'s multiplied by baseline probabilities sum to one.

Notice that the g_m^* tilt exponentially towards the states for which the post-jump value function $\phi_m \left(\overline{y} \right)$ is low relative to the pre-jump value function ϕ. Combining all the terms involving the minimizing g_m's for a given h:

$$
\mathcal{I}(y) h \left[-\xi_p \log \sum_{m=1}^{M} \pi_m^p \exp\left(\frac{1}{\xi_p} \left[\phi - \phi_m \left(\overline{y} \right) \right] \right) \right]
$$

In light of this formula, for a sufficiently large temperature anomaly, y, we expect:

$$
\phi(y) \approx -\xi_p \log \sum_{m=1}^{M} \pi_m^p \exp\left[-\left(\frac{1}{\xi_p} \right) \omega_m \phi_m(y) \right]
$$

Minimizing over h, we find that:

$$
h^* = \sum_{m=1}^{M} \pi_m^p \exp\left(\frac{1}{\xi_p} \left[\phi - \phi_m \left(\overline{y} \right) \right] \right).
$$

with the minimized value given by:

$$\xi_p \mathcal{I}(y) \left[1 - \sum_{m=1}^{M} \pi_m^p \exp\left(\frac{1}{\xi_p} \left[\phi - \phi_m \left(\overline{y} \right) \right] \right) \right]$$

7 A Quantitative Illustration

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 specification of aversions to climate model ambiguity and damage function misspecification were captured by two parameters, ξ_a and ξ_p. We use these parameters to penalize 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. Second, these altered probabilities provide an adjustment for social valuation implied by model ambiguity and misspecification uncertainty. This adjustment is analogous to the risk neutral probability measure used in derivative claims pricing that absorbs the risk adjustments pertinent to valuation. The probabilities that we feature follow from our broad perspective on uncertainty that extends beyond risk considerations.

7.1 Robust adjustments to climate model uncertainty

For the 144 carbon-climate dynamic models, we take as our baseline an equal weighting of all of the models. We explore the consequences of changing these probability, and in so doing we follow Barnett et al. (2021) by setting $\xi_a = .01$. 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 concerns 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. We note that, while the concerns about geo-scientific uncertainty are state-dependent, the distortion in the probability distribution of θ_ℓ values remains roughly constant over the course of our simulations.[3]

[3]Distorting the probabilities over the carbon-climate models alters the drift or local mean in log damage and temperature anomaly dynamics. As a complementary perspective, Barnett et al. (2021) compute and interpret the implied relative entropy of this drift distortion.

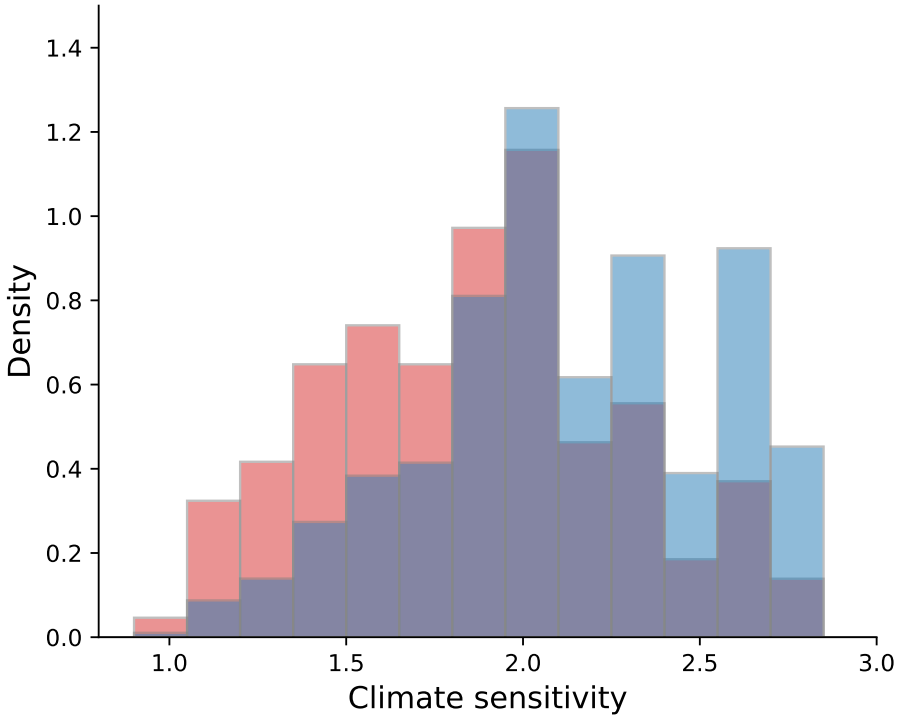

Figure 5: Histograms of climate sensitivity parameters. The red histogram is the outcome of equally weighting all 144 climate models. The blue histogram is the outcome of the minimization in the recursive formulations of our social planners problem.

7.2 Robust adjustments to damage function uncertainty

We next consider the penalty parameter ξ_p that governs concerns about misspecifying 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 in the steepness in the damage function and timing for 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 $\xi_p = 5, \xi_p = 1, \xi_p = .3$ in Figure 6. The distorted histogram for the lowest value, $\xi_p = .3$, is arguably extreme, although the other two choices seem considerably harder to dismiss.

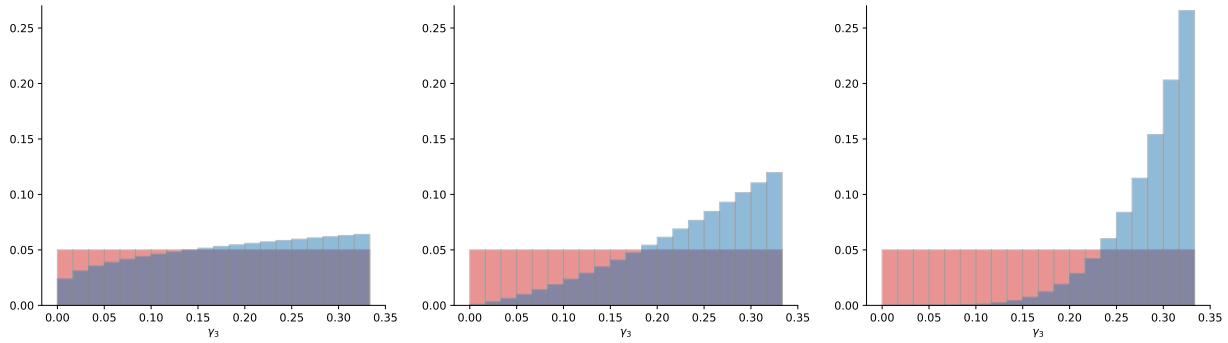

Figure 6: Distorted probabilities of damage functions. Baseline probabilities for damage functions are 1/40 (blue bars), and the red bars are robust adjustments to the probabilities induced by model misspecification concerns (left panel: $\xi_p = 5$, center panel $\xi_p = 1$, right panel: $\xi_p = 0.3$). 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: $\xi_p = 1$. 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:

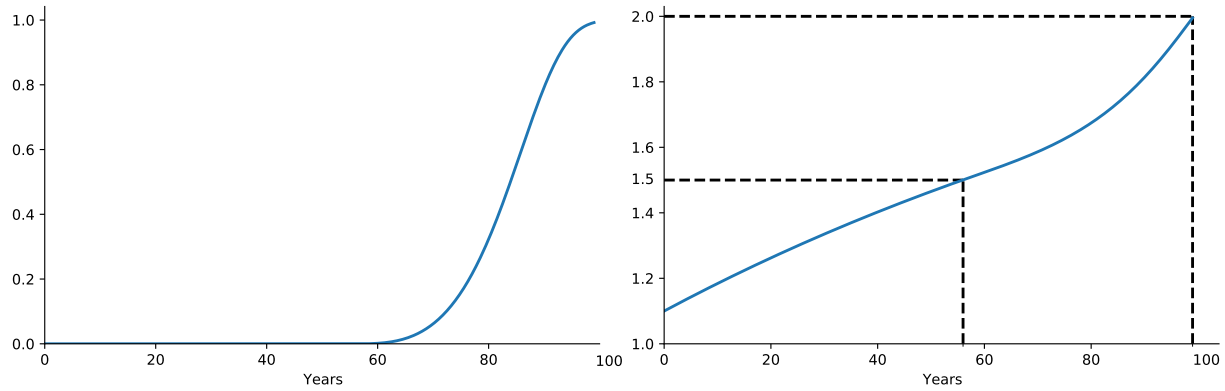

Figure 7: The left panel shows the probabilities that a jump will occur prior to the date given on the horizontal axis. The right panel shows the simulated pathway for the temperature anomaly and the points where the anomaly reaches $\underline{y} = 1.5$ and $\overline{y} = 2.0$ (black dashed lines).

7.3 Emission and anomaly trajectories

To show the effects of concerns about damage function uncertainty on policy decisions of the planner, we now turn to the emissions trajectories across the different cases. We first show reporting the pre-jump emissions pathways for three values of ξ_p in Figure 8. For comparison,

this figure also includes the emissions pathway when the planner has full commitment to the baseline probabilities. We abstract from the role of the Brownian shocks. Their inclusion would add potentially interesting fluctuations needed to account for actual time series behavior. Each line in the figure ends at the point where $y = \underline{y} = 1.5$ in the simulation. Beyond this point the probability of a jump occurring becomes non-zero. This substantial level of precaution employed by the planner is due to the concerns about the unknown future damage state as the emission trajectory for say, $\xi_p = 1$, is about twenty percent lower than the baseline trajectory. We also see that as the value of ξ_p is decreased the initial caution is amplified and the choice of emissions is lowered even further. The lower trajectories necessarily reach the $\underline{y} = 1.5$ threshold later.

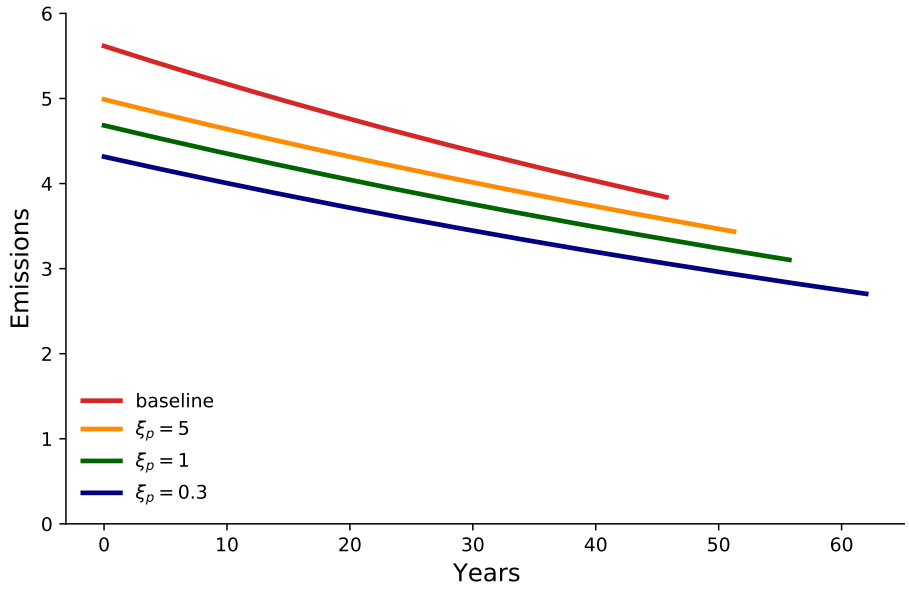

Figure 8: Emissions trajectories before a Poisson event under different penalty configurations. The trajectories are simulated under the baseline probabilities abstracting from the intrinsic randomness. The pathways stop when the temperature anomaly reaches 1.5 and the jump intensity becomes non-zero.

When it reaches temperature anomaly values $y \in [\underline{y}, \overline{y}] = [1.5, 2]$, the Poison event revealing the full damage function will at some point be realized. Abstracting from Brownian shocks, once the event is realized, the emissions trajectory will follow one of forty possible paths. Figure 9 gives summary information about these paths, including the minimum, the maximum and two quantiles for the case in which $\xi_p = 1$. The overall range of possible outcomes is quite large, with emissions values ranging from 12 GtC initially down to 5 GtC. Each trajectory is downward reflecting the downward slops of the individual damage function specifications. Interestingly, there is a big reduction in emissions going from the damage function with least curvature to the .25

quantile as reflected in the blue and orange curves. The emissions trajectories below this quantile are much closer. The realization of a very low damage function specification at the time of the Poisson event is good news to the planner, resulting in an increase in emissions in contrast to the many of the other damage function specifications that could be realized.

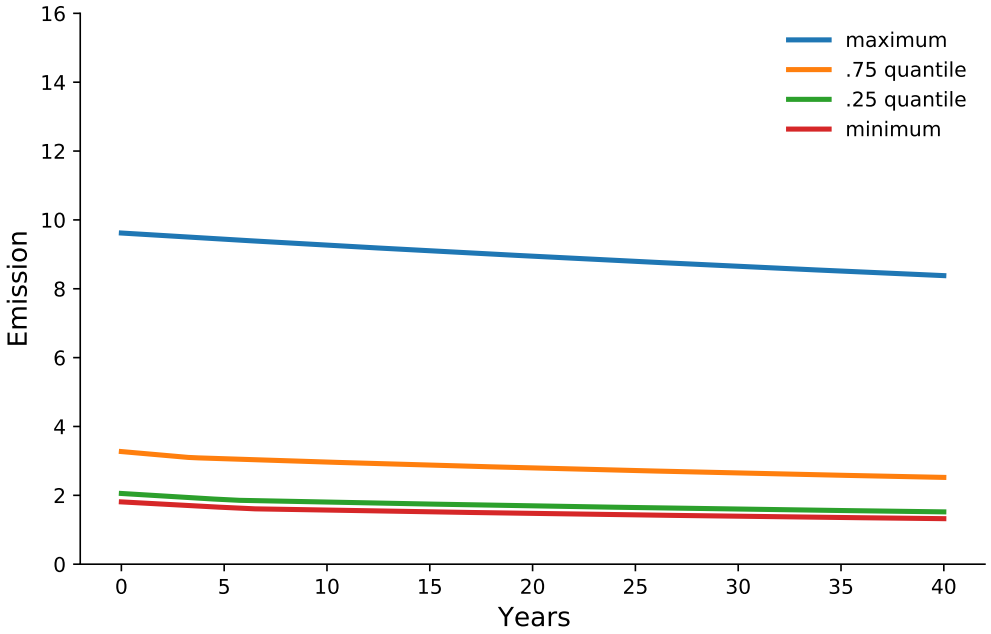

Figure 9: Emissions pathways, conditioned on a jump having occurred, for different realized damage function parameters: γ_3^m. Shown are the extreme minimum (red line) and maximum (blue line) emissions pathways, as well as the 25th percentile (green line) and 75th percentile (orange line) for 0 to 40 years after the jump has occurred.

Given the probabilistic nature of the Poisson event, the emissions and resulting temperature anomalies behave probabilistically. While the planner has uncertainty about these probabilities, we find it revealing to report the distributions under the baseline probabilities. We show these distributions conditioned on the Poisson event being realized in in Figures 10 and 11.[4]

[4]The number of outcomes in the histograms is determined by the number of values of γ_3^m, and the time discretization used in the simulation.

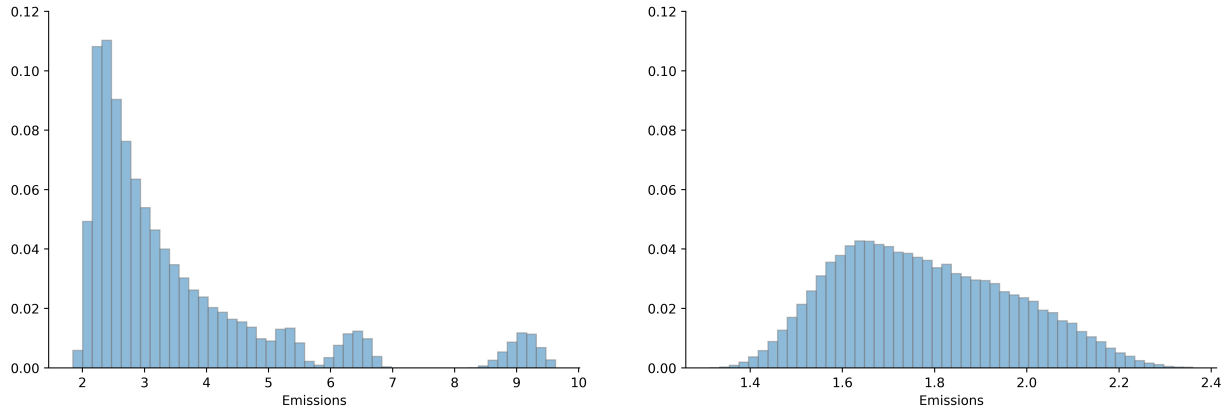

Figure 10: Approximate distribution emissions values for the scenario when $\xi_p = 1$. the emissions values are for year 99, which is when the simulation conditional on no jump occurring reaches a temperature anomaly of 2.0 degrees Celsius, for cases where a jump has occurred prior to then. The simulation is done under the baseline probabilities and abstract from Brownian motion uncertainty. The left histogram shows the outcomes for the 20 least severe values of γ_3^m and the right histogram shows the outcomes for the 20 most severe values of γ_3^m.

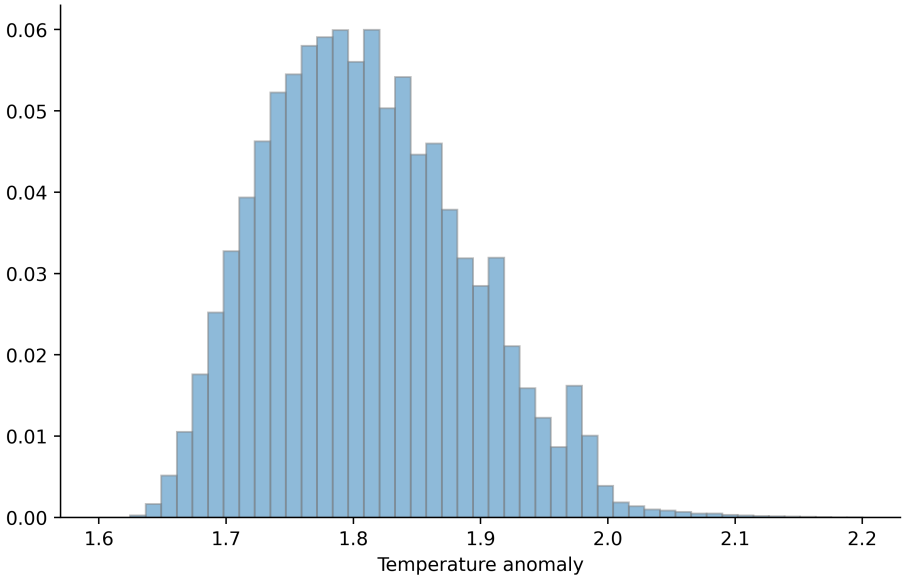

Figure 11: Histogram of possible temperature anomaly values for the scenario where $\xi_a = 0.01$, $\xi_p = 1$. The temperature anomaly values are for year 99, which is when the simulation conditional on no jump occurring reaches a temperature anomaly of 2.0 degrees Celsius, for cases where a jump has occurred prior to then. The simulation is done under the baseline probabilities and abstract from intrinsic randomness. The number of outcomes in the histograms is determine by the number of values of γ_3^m and the discrete number of possible points at which the Poisson event revealing the additional damage curvature could occur in the simulation.

Figure 10 shows the range of possible emissions choices by the social planner conditional on a jump having occurred. The right plot shows the emissions values for the 20 most severe damage function states having been realized. Within this set of damage functions, it is clear that the planner will continue to exercise significant caution, even more than before the jump being realized. This is due to the severity of the future costs of climate change associated with these damage function states. The left plot shows the emissions values for the 20 least severe damage function states having been realized. While we note that even within this set of damage functions, there are still significant realizations for which the planner continues to exercise significant caution with regards to climate change because of the severity of potential future climate damages, there are a number of cases for which there is significantly less restraint. Starting from the right tail of the distribution, we see three distinct minor "peaks." These emissions values correspond to the three least severe damage function states having been realized at different threshold values. Thus, we see that there are in fact "good news" states possible that the planner can experience that lead to far less restraint on the emissions following the Poisson event.

Figure 11 shows the range of possible temperature anomaly values corresponding to the emissions choices in Figure 10. The vast majority of temperature anomaly values fall below the 2.0 degree Celsius value, though there is a small right tail going beyond that and small "peaks" leading up to it. This relatively constrained distribution of temperature anomalies is driven largely by the initial caution exercised by the planner. Realizations of any of the small number of "good news" states eventually result in higher temperature anomalies because the planner's increase in emissions after the Poisson event.

8 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. Our use in this manner is consistent with macroeconomists quantifying efficiency wedges in labor markets and elsewhere. The planner equates marginal social costs and benefits of emissions. We represent the marginal benefits in units of damaged consumption so that:

$$SCC_t = \frac{\eta\left(\widetilde{C}_t\right)}{(1-\eta)\left(\widetilde{E}_t\right)} = \frac{\eta\left(C_t\right)}{(1-\eta)\left(N_t\right)\widetilde{E}_t}$$

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

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

As noted by Barnett et al. (2021), "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:

$$\log \eta - \log(1-\eta) - (\log N_t) - \log \widetilde{E}_t. \tag{6}$$

To deduce the emissions contribution, we start by differentiating the HJB equation (4) with respect to $\tilde{e}$ and solve for $\frac{\eta}{\tilde{e}}$:

$$\frac{\eta}{\tilde{e}} = -\frac{d\phi(y)}{dy}\sum_{\ell=1}^{L}\omega_\ell\theta_\ell - \frac{d^2\phi(y)}{(dy)^2}|\varsigma|^2\tilde{e} + \frac{(1-\eta)}{\delta}\left[(\gamma_1+\gamma_2 y)\sum_{\ell=1}^{L}\omega_\ell\theta_\ell + \gamma_2|\varsigma|^2\tilde{e}\right] \tag{7}$$

for $0 < \overline{y}$. 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.

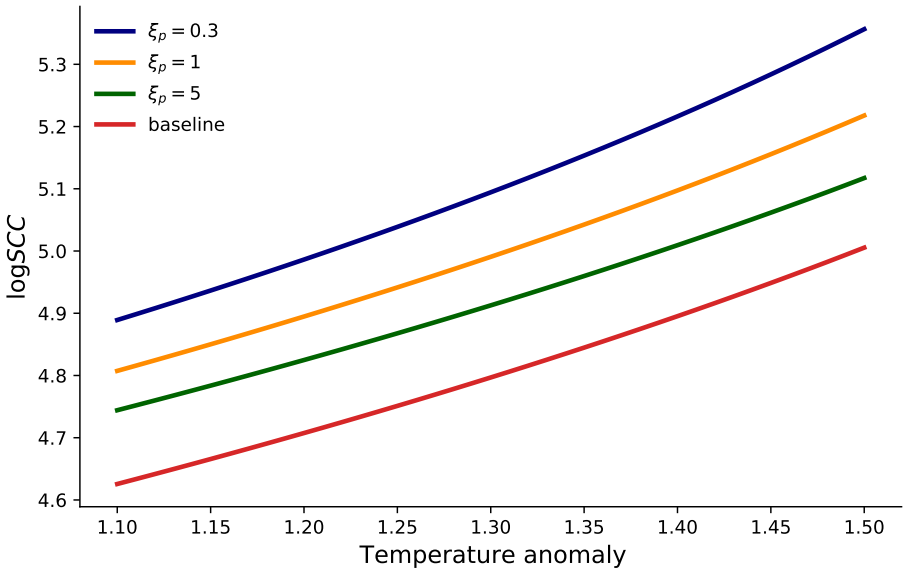

Figure 12: $\log(SCC)$ as functions of temperature anomaly under different penalty configurations. The SCC is state dependent, and we focus on the domain for which the anomaly is less than 1.5 degrees. The values logarithms are translated by the initial period logarithm of consumption.

Figure 12 shows the log SCC for the baseline case of $\xi_p = \infty$ (red line), and three cases of increasing concerns about damage function uncertainty: $\xi_p = 5$ (green line), $\xi_p = 1$ (yellow line), and $\xi_p = 0.3$ (blue line). The results use the socially efficient emissions trajectories for the different values of ξ_p. The log SCC values are calculated using formula (7), where we infer $\log N_t$ from formula (1). 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_p).

ξ_p	1.1	1.2	1.3	1.4	1.5
5	11.5%	11.4%	11.4%	11.2 %	11.1%
1	17.3%	17.9%	18.7%	19.7%	20.9%
0.3	24.4%	26.0%	28.2%	30.9%	34.5%

Table 1: The fraction of the SCC due to uncertainty for three uncertainty configurations ($\xi_p = 5$, $\xi_p = 1$, $\xi_p = 0.3$) at temperature anomaly values in the simulations of $1.1, 1.2, 1.3, 1.4$, and 1.5. The post-jump uncertainty contribution to the SCC is approximately 10%.

To compute the uncertainty component to the SCC, we use the emissions trajectory from the planner's problem, but we recompute the value function ϕ using the baseline probabilities. We then use the right side of (7) as the counterpart to the SCC absent uncertainty.[5] This leads to a smaller SCC. The uncertainty contribution is the difference in logarithms between the original SCC and the one that abstracts from uncertainty.

In Table 1, we provide values for the contribution of uncertainty to our SCC numbers for three values of ξ_p: 5, 1, and 0.3. As the value of ξ_p decreases, the uncertainty contribution to the SCC is amplified. While the uncertainty contribution to the SCC stays at around 11% when $\xi_p = 5$, the contribution increases to almost 35% at the lower threshold $\underline{y} = 1.5$ when $\xi_p = 0.3$. The concerns about uncertainty create not only level adjustments, but dynamic differences in climate policy as measured by our SCC values.

9 Conclusions

As concerns about the human impact on the climate have come to the forefront of recent policy discussions, many have argued for significant and immediate policy to prevent future climate change damages. And yet, geo-scientists and economists have identified significant uncertainties about the magnitude and timing of climate change and its consequences. This raises questions about the magnitude and timing of the policy responses given that we could learn more about environmental and economic damages due to climate change in the future. We addressed these questions in the setting of a stylized social planning problem using dynamic decision theory under uncertainty.

In our 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

[5]The right and left sides of (7) are no longer equated because we purposefully do not re-optimize by choice of emissions.

future. The damage function uncertainty is resolved by a single Poisson event that becomes more 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 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. 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.

References

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.

Barnett, Michael, William A. Brock, and Lars Peter Hansen. 2020. Pricing Uncertainty Induced by Climate Change. *Review of Financial Studies* 33 (3):1024–1066.

Barnett, Michael, William Brock, Lars Peter Hansen, et al. 2021. Climate Change Uncertainty Spillover in the Macroeconomy. *Prepared for the 2021 Macoreconomics Annual* .

Cai, Yongyang, Timothy M Lenton, and Thomas S Lontzek. 2016. Risk of multiple interacting tipping points should encourage rapid CO 2 emission reduction. *Nature Climate Change* 6 (5):520–525.

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.

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.

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. 2020. Macroeconomic Uncertainty Prices when Beliefs are Tenuous. *Journal of Econometrics* published online.

Hausfather, Z and GP Peters. 2020. Emissions – the 'business as usual' story is misleading. *Nature* 577 (7792):618–620.

Joos, F., R. Roth, J. S. Fuglestvedt, 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.

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.

MacDougall, Andrew H., Neil C. Swart, and Reto Knutti. 2017. The Uncertainty in the Transient Climate Response to Cumulative CO_2 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.

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.

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 Strefler, 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.

Appendix A Value function components and parameter values

As in Barnett et al. (2021), the HJB equations in Section 6, are simplified HJB equation that depend only on the the state variable y. The full climate-economics based HJB equation following from the illustrative economy laid out in Section 5 are based on a larger system with additional state variables. The larger system can be analytically simplified and solved in subsystems to produce the full solution.

The parameters used for the climate-economics sub-system that we focus on in the paper are given in Table **??**.

Parameter	Value
ς'	$[2.23 \quad 0]$
γ_1	.000177
γ_2	.0044
γ_3^{min}	.000
γ_3^{max}	.333
$\underline{y}$	1.5
$\overline{y}$	2.0
r_1	1.5
r_2	2.5
η	.032
δ	.01

Table 2: Climate economics system parameters

Note that $\underline{y}$ and $\overline{y}$ define the upper and lower bounds of the interval over which the Poisson event occurs. Also note that γ_3^{min} and γ_3^{max} denote the upper and lower bounds of the interval over which values of γ_3^m are assumed to be within in our analysis. In practice, we use $M = 40$ values of γ_3^m, where the values of γ_3^m are chosen such that values of $exp(-n)$ at $y = 2.5$ for a threshold value of $\overline{y} = 2.0$ are equidistant between the value of $exp(-n)$ for $\gamma_3 = .000$ and the value of $exp(-n)$ for $\gamma_3 = .333$.

To put in context the choice of ς that we use in our analysis, note that for a constant emissions path, the implied standard deviation associated with the Matthew's approximation coefficient is .446 at twenty five years, .315 for fifty years and .223 for one hundred years.

We omit details for the derivation of consumption-capital dynamics subsystem, and direct the interested reader to the appendix of Barnett et al. (2021), which provides the derivation we use here. We set up the model so that the undamaged version of consumption capital model has a straightforward solution. The "guess and verify" method allows us to derive a solution for this

subsystem, where we guess a value function of the form $\upsilon_k \log k + \zeta(z)$.

For the calibration of this component of the model, we use values such that the model is consistent with previous work including Hansen and Sargent (2020) and Barnett et al. (2021). Therefore, we the values for the AK model productivity and capital adjustment costs to be $\alpha = .115$ and $\kappa = 6.667$, respectively. We also 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. As a result, the implied consumption dynamics are consistent with an annualized counterpart of those given in Hansen and Sargent (2020).