How Should Climate Change Uncertainty Impact Social Valuation and Policy?*

Michael Barnett* William Brock** Lars Peter Hansen‡
Hong Zhang‡‡

January 31, 2024

Abstract

We study the uncertain transition to a carbon-neutral economy. The requisite technological innovation is made more probable through research and development (R&D). We explore multiple channels of economic and geoscientific uncertainties that impact this transition, and we show how to assess the relative importance of their varied contributions. We represent the social benefit of R&D and cost of global warming as expected discounted values of social payoffs using a probability measure adjusted for concerns about model misspecification and prior ambiguity. Our quantitative results show the value of R&D investment even when the timing of its technological success is highly uncertain.

*The authors thank Fernando Alvarez, Adlai Fisher, Peter Hansen, Joanna Harris, Diana Petrova, Eric Renault, and Judy Yue for helpful suggestions and Bin Cheng for valuable research assistance. In addition, participants at the 2023 SITE Conference on Climate Finance, Innovation, and Challenges for Policy provided valuable feedback. This research was supported in part by the University of Chicago Griffin Economics Incubator.

**University of Wisconsin and University of Missouri, Columbia
‡University of Chicago
*Arizona State University
‡‡Argonne National Laboratory

1 Introduction

In this paper we derive social valuations and uncertainty contributions that underlie robustly optimal policy actions in response to the economic consequences of climate change. Characterizing the potential long-run risks of climate change has become a focal point for economic and policy analysis. An influential literature in macro-asset pricing studies the potential impact of long-term macroeconomic uncertainty on current-period valuation. See, for example, Bansal and Yaron (2004) and the subsequent literature. While motivated by insights from this literature, our contribution differs in two important ways. First, we study the long-term uncertainty associated with climate change. This leads us to explore specific structures of this uncertainty from multiple sources as it impacts a production-based model. Second, we study valuation from a social perspective rather than a market perspective. These two perspectives differ because of production externalities. We adopt the social vantage point because climate change induced by human activity is at the forefront of many policy discussions. The counterpart to market-determined asset prices in our analysis are (i) the social cost of global warming and (ii) the social value of research and development of new viable clean technologies. The social perspective that we feature provides a valuable benchmark for effectiveness of ad hoc policies implemented in market settings.

To design prudent policies that limit climate change requires that we confront multiple sources of uncertainty, some of which are difficult to gauge with full confidence. As Rising et al. (2022) have recently argued,

> The economic consequences of many of the complex risks associated with climate change cannot, however, currently be quantified. ... these unquantified, poorly understood and often deeply uncertain risks can and should be included in economic evaluations and decision-making processes.

While we do not take up the full challenge posed by uncertainties featured by Rising et al., we do propose an approach for confronting "deep uncertainty" that is formally integrated into policy analysis. We demonstrate the value of such an approach in focusing on uncertainty in four inputs into prudent climate policy-making:

i) *carbon-climate dynamics*: mapping from carbon emissions into temperature changes

ii) *economic damage functions*: reductions in output because of changes in atmospheric temperature

iii) *technological innovation*: future abatement cost reductions for CO_2-based energy through investment in research and development (R&D).

iv) *investment productivity*: exogenous shifts in the capital evolution that have persistent consequences.

The fourth of these shocks captures the familiar "long-run risk" contributions to the macroeconomy in the absence of climate change and provides a benchmark for comparison. The other three are explicitly linked to climate change and contribute to our analysis in specific ways. We capture these with both Brownian and jump risk components.

We push beyond the usual risk-based analyses by taking a broad approach to uncertainty, consistent with perspectives from multiple fields, including decision and control theory. These literatures formalize sagacious responses to alternative types of uncertainty, including *risk* within a model (unknown outcomes with known probabilities), *ambiguity* across models (unknown priors weighting alternative models or parameter configurations), and potential *model misspecification* (unknown flaws of fully specified probability models). The last of these three constructs gives a way to formulate concepts such as "deep uncertainty," or "radical uncertainty," that are present in many policy discussions. While many references to these concepts exist, as evident from the call in our quote from Rising et al. (2022), little has been done to address these uncertainties in a formal way.

We also borrow and extend insights from asset pricing theory to study two social assets: (i) the *social cost of global warming* and (ii) the *social value of research and development* capital for developing new green technologies. We think of both of these as assets characterized by their implied intertemporal social cash flows. The asset price representations we use here depend on a "tilted probability" distribution that reflects aversion to model ambiguity or model misspecification. While much is made of the so-called social cost of carbon (SCC), quantifying the social value of R&D directed at the discovery of new and economically viable green technologies has been largely ignored. But even the social cost of global warming, a central input into the SCC, is altered by our inclusion of R&D investment in the toolkit of the social planner, modifying some of previous analyses in a substantive way.

For the most part, in this paper we consider a discounted dynamic decision problem under uncertainty for a fictitious social planner. We introduce the problem to construct a benchmark target for ad hoc policies. There is scope for such policies to be socially productive because of the failure of the "invisible hand." Market solutions do not account for the climate change externality.

Using our fictitious social planner setting, we provide a novel decomposition to assess which uncertainty components are most consequential for policy related to climate change. To achieve this, we propose and implement refinements to uncertainty quantification methods. Our approach is based explicitly on a decision theory perspective. We solve alternative decision problems with different configurations of uncertainty, each featuring a specific source of the uncertainty. By varying the hypothetical exposure to the uncertainty, we show which of the four sources of uncertainty has the biggest impact on the prudent decisions of the planner.

1.1 Related literature

Several papers engage in uncertainty quantification by imitating approaches taken in the scientific literature. As a result, these assessments are done from the perspective of a planner who may unjustifiably ignore the uncertainty. This is in contrast to an external analyst who acknowledge this uncertainty and explores its ramifications. An important component of our analysis is to bring at least some of the uncertainty inside the decision problem of the planner.

Some important papers do this by embracing a narrowly defined construct of risk (unknown outcomes with known probabilities) targeting the quantification of the social cost of carbon. See for instance, Cai et al. (2017), Bansal et al. (2019), and Hambel et al. (2021). While we embrace a more encompassing notion of uncertainty, some of our analysis could be reinterpreted from a pure risk perspective. For reasons that we will explain, we prefer the broader uncertainty vantage point. Kelly and Kolstad (1999) explore an explicit Bayesian learning approach in their stylized study of climate change uncertainty. While they show that endogenous learning evolves very slowly, they do not explore a policy maker with prior ambiguity and its impact. Weitzman (2009) uses a subjective expected utility approach to the study of damage uncertainty, also abstracting from prior ambiguity.

Broader notions of uncertainty have been considered in a few papers within the climate-economics literature, including Hennlock (2009), Li et al. (2016), Lemoine and Traeger (2016), Barnett, Brock and Hansen (2020), Rudik (2020), Berger and Marinacci (2020), Barnett, Brock and Hansen (2022), and Barnett (2023). With the exception of Barnett, Brock and Hansen (2020), this literature does not explore the implied social valuations from an asset pricing perspective, nor does it explore simultaneously the four sources of uncertainty that we do here.[1] Specifically, we study a novel framework that allows for multiple endogenous policy responses to climate change in the form of optimal choices of emission and R&D investment directed at the discovery of new and economically viable green technologies. We exploit this setting to provide an enriched policy analysis in the form of novel social valuations and uncertainty contributions for green R&D investment and technological innovation, as well as the spillovers arising from such actions. Notably, the inclusion of R&D investment in the planner's policy choice set gives rise to some salient modifications of the uncertainty contribution to the social price of global warming as compared with our previous analyses.[2]

1.2 Organization

The remainder of the paper proceeds as follows. Section 2 outlines the dynamic decision theory toolkit used for our uncertainty analysis. Sections 3–6 detail the climate-economic modeling framework. Specifically, Section 3 describes the underlying economic production opportunities, abstracting from climate change; Section 4 gives the model of climate dynamics; Section 5 shows how these opportunities are proportionally reduced by global warming; and Section 6 gives the intertemporal preferences of the hypothetical social planner. We construct the composite state vector in Section 7, and we introduce a convenient transformation of this vector that facilitates both characterization and computation. Section 8 shows how to use impulse responses combined with state-dependent, probabilistic discounting to represent the social value of the stock of R&D investment and the social cost of global warming. Section 9 provides a novel robustness decomposition that is featured in our quantitative in-

[1] Barnett (2023) explores some standard asset pricing implications with a decentralization that supports a robustly optimal policy.

[2] We fully acknowledge that R&D has been studied previously, including the recent paper Jaakkola and van der Ploeg (2019) along with earlier research such as Acemoglu et al. (2016).

vestigation. Quantitative results from an example economy are given in Section 10. Section 11 provides model sensitivity analysis across various alternative specifications. Section 12 summarizes our conclusions and briefly comment on the value of new R&D investment.

2 Confronting uncertainty

We pose our hypothetical social planner's decision problem in a continuous-time environment. In this section we study the contributions to the planner's Hamilton–Jacobi–Bellman (HJB) equation that emerge because of aversion to model misspecification or alternatively to ambiguity over models. These uncertainty adjustments lead us to replace a recursive maximization problem with a two-player formulation where one player maximizes social well-being and the other adversarial player looks for baseline model or prior misspecifications with the most adverse consequences by solving a minimization problem. In effect, this becomes a two-player zero sum game posed in continuous time with modifications to the HJB equations necessary to accommodate the uncertainty concerns. The analysis in this section focuses on the minimizing player and the implications for terms that enter the HJB equation involving the evolution of the value function. We temporarily omit the other terms, remembering that these are also important for deriving the maximizing control law.

To confront potential model misspecification, we follow Anderson et al. (2003) by entertaining misspecification linked both to the Brownian contribution and to the jump contribution. We also explore model ambiguity through the potential misspecification of a subjective distribution over unknown models by implementing an approach proposed by Hansen and Miao (2018). In all cases we use a well-studied construct called relative entropy or Kullback–Leibler divergence scaled by a penalty parameter. This divergence is a well-known construct measured as a non-negative expected log-likelihood ratio that we use to restrain the search over potential model and prior discrepancies in the uncertainty analysis.

2.1 Brownian motion misspecification

Under a baseline probability specification, $W \stackrel{\text{def}}{=} \{W_t : t \geqslant 0\}$ is a multivariate standard Brownian motion, and $\mathfrak{F} \stackrel{\text{def}}{=} \{\mathfrak{F}_t : t \geqslant 0\}$ is the corresponding information filtration with $\mathfrak{F}_t$ generated information that is realized between dates zero and t. This information includes

the Brownian increments that have been realized up to date t. Initially, we let $\mathfrak{F}$ be the Brownian filtration, but we subsequently will augment this filtration to include realizations of jumps that will influence technology and damages induced by climate change.

As is familiar from derivative claims pricing, positive martingales with expectations equal to one parameterize changes in probability measures. From Girsanov theory, such martingales can be characterized by their implied drift distortions. In particular, under the martingale change in the probability measure, process $W \stackrel{\text{def}}{=} \{W_t : t \geqslant 0\}$ instead has a drift $H \stackrel{\text{def}}{=} \{H_t : t \geqslant 0\}$.

Suppose that the state vector process X has a local mean increment $\mu_x(X_t, A_t)dt$ and stochastic increment $\sigma_x(X_t, A_t)dW_t$, where A_t is a decision or action taken at time t. Throughout the essay we let lower-case variables capture potential realizations of random vectors. The realizations of the state vector, X_t, reside in a state space $\mathcal{X}$. For a value function, $\widehat{V}$, the drift or local mean of $\widehat{V}(X)$ is given by[3]

$$\frac{\partial \widehat{V}}{\partial x'}(X_t)\mu_x(X_t, A_t) + \frac{1}{2}\text{trace}\left[\sigma_x(X_t, A_t)'\frac{\partial^2 \widehat{V}}{\partial x \partial x'}(X_t)\sigma_x(X_t, A_t)\right]. \tag{1}$$

In this equation, we omit time dependence and think of $\widehat{V}$ as the value function for an infinite horizon discounted problem. Formula (1) captures the time increment to risk confronted by the decision-maker.

The decision-maker entertains possible misspecification uncertainty by replacing formula (1) with the solution to

Problem 2.1.

$$\min_h \frac{\partial \widehat{V}}{\partial x'}(x)\left[\mu_x(x, a) + \sigma_x(x, a)h\right] + \frac{1}{2}\text{trace}\left[\sigma_x(x, a)'\frac{\partial^2 \widehat{V}}{\partial x \partial x'}(x)\sigma_x(x, a)\right] + \frac{\xi}{2}h'h$$

for a penalty parameter ξ.

The minimization captures a form of uncertainty aversion, analogous to risk aversion, since the minimizing objective will be less than or equal to (1). The penalty parameter ξ restrains

[3]We use the notation $\frac{\partial \widehat{V}}{\partial x}(x)$ to denote a column vector of derivatives with respect to the column vector x and $\frac{\partial \widehat{V}}{\partial x'}(x)$ to be the corresponding row vectors of derivatives with respect to the row vector x'.

the concern for robustness to model misspecification. The quadratic penalty in h is a local measure of "relative entropy" or Kulback–Leibler divergence.[4] A limiting choice of $\xi \approx \infty$ implies a minimizing choice of $h = 0$ with an implied contribution given by (1). Since the minimization problem is quadratic in h, the minimizer is

$$h^* = -\frac{1}{\xi}\sigma_x(x,a)'\frac{\partial \widehat{V}}{\partial x}(x) \tag{2}$$

with a minimized objective:

$$\frac{\partial \widehat{V}}{\partial x'}(x)\mu_x(x,a) + \frac{1}{2}\text{trace}\left[\sigma_x(x,a)'\frac{\partial^2 \widehat{V}}{\partial x \partial x'}(x)\sigma_x(x,a)\right] - \frac{1}{2\xi}\frac{\partial \widehat{V}}{\partial x'}(x)\sigma_x(x,a)\sigma_x(x,a)'\frac{\partial \widehat{V}}{\partial x}(x).$$

Notice that the minimizing drift, h^*, is potentially state dependent. When σ depends on the action a, the drift of interest for valuation and interpretation depends on the maximizing action a expressed as a function of the state. The drift vector, h^*, has larger entries where the value function is more adversely exposed to the Brownian increments. Smaller values of ξ result in drift adjustments with larger magnitudes.[5]

Remark 2.2. *While we motivated the adjustment to an HJB equation in terms of robustness, it may also be viewed as a risk adjustment in conjunction with recursive utility. Consider the local counterpart or small ϵ counterpart to the risk adjustment:*

$$\begin{aligned}&\frac{1}{\epsilon}\left(\frac{1}{1-\gamma}\log \mathbb{E}\left[\exp\left((1-\gamma)\widehat{V}(X_{t+\epsilon})\right) \mid \mathfrak{F}_t\right] - \widehat{V}(X_t)\right)\\ &= \frac{1}{\epsilon(1-\gamma)}\log \mathbb{E}\left[\exp\left((1-\gamma)\left(\widehat{V}(X_{t+\epsilon}) - \widehat{V}(X_t)\right)\right) \mid \mathfrak{F}_t\right].\end{aligned} \tag{3}$$

This exponential risk adjustment of the continuation value induces a local log-normal type of adjustment given by

$$(1-\gamma)\frac{\partial \widehat{V}(x)}{\partial x'}\sigma(x,a)\sigma(x,a)'\frac{\partial \widehat{V}(x)}{\partial x},$$

where the term multiplying $1-\gamma$ is the local variance of the continuation value process $\widehat{V}(X)$.

[4]For instance, see Anderson et al. (2003) for an elaboration.

[5]While this looks obvious from formula (2), it is a bit more subtle because the value function implicitly depends on ξ.

Setting $\gamma - 1 = \frac{1}{\xi}$ *gives a mathematical equivalence between robustness and risk considerations, although the rationale for the two adjustments is very different. We will eventually show how to embed this into a full recursive utility specification. The resulting equivalence is a direct extension of a well-known result from risk-sensitive control theory.*

Remark 2.3. *As Anderson et al. (2003) emphasize, the negative implied drift distortions from a planner's problem are also the local shadow prices for concerns about misspecification. While they featured the case in which these shadow prices are also pertinent for competitive financial markets, the same insight carries over to social valuation in the presence of externalities that induce a wedge between market prices and social counterparts.*

2.2 Jump misspecification

Jump components play prominently in our uncertainty analysis. This is due to the inclusion of model ingredients that capture uncertainty in the magnitude and timing of damages to economic opportunities induced by climate change, as well as uncertainty in the arrival of a new, economically viable, green technology.

The jumps depend on endogenously determined intensities that govern the probabilities of the jump realizations. Our specification of these intensities induces a corresponding endogeneity in the information structure.

We suppose there is a discrete set of jump states $\mathcal{Z}$. Let z denote a realized value in the set $\mathcal{Z}$. Let $\mathcal{J}$ denote a state-dependent jump intensity, and let $\pi(\tilde{z} \mid x, z)$, $z \in \mathcal{Z}$ give the jump distribution conditioned on a jump when $X_t = x$ in discrete state $Z_t = z$. Recall that the jump intensity, $\mathcal{J}$, implies an approximate jump probability, $\epsilon\mathcal{J}$, over a small time increment, ϵ. Following a jump, x changes as does the value function. For each choice of $\tilde{z}$, x jumps to the $\tilde{x}(\tilde{z})$, where $\tilde{z}$ is uncertain prior to the jump. The baseline probabilities are $\pi(\tilde{z} \mid x, z)$ for $z \in \mathcal{Z}$, where

$$\sum_{\mathcal{Z}} \pi(\tilde{z} \mid x, z) = 1.$$

With these jumps, a value function shifts from $\widehat{V}(x, z)$ to $\widetilde{V}[\tilde{x}(\tilde{z}), \tilde{z}]$. The jump process contributes the following term to the drift of $\widehat{V}(X, Z)$:

$$\mathcal{J}(x, z) \sum_{\tilde{z} \in \mathcal{Z}} \left[\widetilde{V}[\tilde{x}(\tilde{z}), \tilde{z}] - \widehat{V}(x, z) \right] \pi(\tilde{z} \mid x, z), \tag{4}$$

capturing the jump risk contribution to the decision problem.

To capture potential misspecification, we introduce a non-negative function f where the altered jump distribution is

$$\frac{f(\tilde{z} \mid x, z)\pi(\tilde{z} \mid x, z)}{\bar{f}(x, z)},$$

and intensity $\mathcal{J}(x, z)\bar{f}(x, z)$, where

$$\bar{f}(x, z) = \sum_{\tilde{z} \in \mathcal{Z}} f(\tilde{z} \mid x, z)\pi(\tilde{z} \mid x, z).$$

To restrain the exploration of potential misspecification, we introduce a convex cost:

$$\xi \mathcal{J}(x, z) \sum_{\tilde{z} \in \mathcal{Z}} \left[1 - f(\tilde{z} \mid x, z) + f(\tilde{z} \mid x, z) \log f(\tilde{z} \mid x, z)\right] \pi(\tilde{z} \mid x, z).$$

The term multiplying ξ is a local (in time) measure of relative entropy or Kullback–Leibler divergence applicable to jump processes.[6]

To confront misspecification, we solve

Problem 2.4.

$$\begin{aligned} \min_{f \geqslant 0} \quad & \mathcal{J}(x, z) \sum_{\tilde{z} \in \mathcal{Z}} \left[\widetilde{V}[\tilde{x}(\tilde{z}), \tilde{z}] - \widehat{V}(x, z)\right] f(\tilde{z} \mid x, z)\pi(\tilde{z} \mid x, z) \\ & + \xi \mathcal{J}(x, z) \sum_{\tilde{z} \in \mathcal{Z}} \left[1 - f(\tilde{z} \mid x, z) + f(\tilde{z} \mid x, z) \log f(\tilde{z} \mid x, z)\right] \pi(\tilde{z} \mid x, z). \end{aligned}$$

Minimization problem 2.4 has a quasi-analytical solution:

$$f^*(\tilde{z} \mid x, z) = \exp\left(-\frac{1}{\xi}\left(\widetilde{V}[\tilde{x}(\tilde{z}), \tilde{z}] - \widehat{V}(x, z)\right)\right),$$

with a minimized objective:

$$\xi \mathcal{J}(x, z) \sum_{\tilde{z} \in \mathcal{Z}} \left[1 - \exp\left(-\frac{1}{\xi}\left(\widetilde{V}[\tilde{x}(\tilde{z}), \tilde{z}] - \widehat{V}(x, z)\right)\right)\right] \pi(\tilde{z} \mid x, z), \tag{5}$$

which we use in place of (4).

[6] See, for instance, Anderson et al. (2003).

Remark 2.5. *We decompose the entropy contribution into two interpretable pieces. The first is the contribution from changing the intensity specification:*

$$\mathcal{J}(x,z)\left[1-\bar{f}(x,z)+\bar{f}(x,z)\log\bar{f}(x,z)\right].$$

The second is the contribution from changing the jump probability contribution:

$$\mathcal{J}(x,z)\bar{f}(x,z)\left[\sum_{\tilde{z}\in\mathcal{Z}}\frac{f(\tilde{z}\mid x,z)}{\bar{f}(x,z)}\left[\log f(\tilde{z}\mid x,z)-\log\bar{f}(x,z)\right]\pi(\tilde{z}\mid x,z)\right]. \tag{6}$$

The expression in the inside $[\,\cdot\,]$ is the usual static measure of relative entropy of a discrete distribution, where the discrete distribution conditions on x.[7] *The sum of these two contributions gives the relative entropy measure that we use.*

Remark 2.6. *For pedagogical reasons, using the decomposition in Remark 2.5, we solve subproblem 2.4 by first minimizing over $f(\tilde{z}\mid x)/\bar{f}(x)$ given $\bar{f}(x)$ and then minimizing over $\bar{f}(x)$. We represent the first minimization as follows:*

$$\begin{aligned}\mathcal{J}(x,z)\,\bar{f}(x,z)\,\min_{\frac{f(\tilde{z}|x,z)}{\bar{f}(x,z)}}\sum_{\tilde{z}\in\mathcal{Z}}\left[\widetilde{V}[\tilde{x}(\tilde{z}),\tilde{z}]-\widehat{V}(x,z)\right]\frac{f(\tilde{z}\mid x,z)}{\bar{f}(x,z)}\pi(\tilde{z}\mid x,z)\\ +\xi\,\mathcal{J}(x,z)\,\bar{f}(x,z)\left[\sum_{\tilde{z}\in\mathcal{Z}}\frac{f(\tilde{z}\mid x,z)}{\bar{f}(x,z)}\left[\log f(\tilde{z}\mid x,z)-\log\bar{f}(x,z)\right]\pi(\tilde{z}\mid x,z)\right].\end{aligned} \tag{7}$$

The term $\mathcal{J}(x,z)\bar{f}(x,z)$ is common to both contributions of the objective and hence effectively scales out of the problem. The result is a static robustness problem featuring the potential misspecification of the jump probabilities. It has the well-known minimized objective

$$-\mathcal{J}(x,z)\,\bar{f}(x,z)\,\xi\,\log\left[\sum_{\tilde{z}\in\mathcal{Z}}\exp\left(-\frac{1}{\xi}\left[\widetilde{V}[\tilde{x}(\tilde{z}),\tilde{z}]-\widehat{V}(x,z)\right]\right)\pi(\tilde{z}\mid x,z)\right].$$

[7] The jump contribution can be deduced by taking a discrete-time approximation with time interval ϵ and taking limits $\frac{1}{\epsilon}$ limits as ϵ declines to zero.

We then solve

$$\min_{\bar{f}} -\mathcal{J}(x,z)\,\bar{f}(x,z)\,\xi\,\log\left[\sum_{\tilde{z}\in\mathcal{Z}}\exp\left(-\frac{1}{\xi}\left[\widetilde{V}[\tilde{x}(\tilde{z}),\tilde{z}]-\widehat{V}(x,z)\right]\right)\pi(\tilde{z}\mid x,z)\right]$$
$$+\,\xi\,\mathcal{J}(x,z)\left[1-\bar{f}(x,z)+\bar{f}(x,z)\log\bar{f}(x,z)\right].$$

It is straightforward to verify that the minimized objective is given by (5).

Remark 2.7. *Following up remark 2.2, we may again interpret the outcome of the minimization problem as a recursive utility risk adjustment of the form* (3) *in the small* ϵ *limit with risk aversion parameter given by:* $\gamma - 1 = \frac{1}{\xi}$.

When solving a robust decision problem, we embed subproblems 2.1 and 2.4 inside coupled HJB equations.

2.3 Incorporating ambiguity aversion

Imagine there are alternative models of different components of the dynamics. We follow Hansen and Miao (2018) by supposing that the drift $\mu(x,z,a\mid\theta)$ depends on an unknown parameter θ residing in a set Θ. The parameter, θ, could index one of a discrete set of alternative models or depict a unknown parameter vector. The decision-maker has a baseline probability $dP_t(\theta)$ for each time instant, t, and makes an adjustment for ambiguity by solving

Problem 2.8.

$$\min_{q,\,\int_\Theta q(\theta)dP_t(\theta)=1}\ \frac{\partial\widehat{V}}{\partial x'}(x,z)\int_\Theta \mu_x(x,z,a\mid\theta)q(\theta)dP_t(\theta)$$
$$+\,\chi\int_\Theta q(\theta)\log q(\theta)dP_t(\theta),$$

where χ *is a penalty parameter.*

This problem is known to have a solution that entails exponential tilting as a function of the drift of the value function for alternative values of θ:

$$q_t^*(\tilde{\theta})=\frac{\exp\left(-\frac{1}{\chi}\frac{\partial\widehat{V}}{\partial x'}(x,z)\mu_x(x,z,a\mid\tilde{\theta})\right)}{\int_\Theta\exp\left(-\frac{1}{\chi}\frac{\partial\widehat{V}}{\partial x'}(x,z)\mu_x\left(x,z,a\mid\theta\right)\right)dP_t(\theta)}.$$

The minimized objective is

$$-\chi \log \int_{\Theta} \exp\left(-\frac{1}{\chi}\frac{\partial \widehat{V}}{\partial x'}(x,z)\mu_x(x,z,a \mid \theta)\right) dP_t(\theta).$$

Notice that this formulation implies an exponential adjustment for model ambiguity concerns.[8] We allow the baseline probability to be time dependent to allow for recursive learning, although we will abstract from this learning in our application.

Problem 2.8 and Problem 2.4 show a notable similarity. The smooth ambiguity model applies to Brownian uncertainty, and the objective of interest is the local evolution of the value function. In the case of jump uncertainty, this is replaced by the intensity times the difference between the post-jump and pre-jump value functions. The counterpart to χ for the smooth ambiguity adjustment is the intensity times ξ. The relative density q in Problem 2.8 plays a role analogous to $f/\bar{f}$ in Problem 2.4 when deducing the worst-case distribution. With this mapping, the two robustness adjustments are mathematically equivalent. As we noted, however, the required specification of the intensity introduces an additional source of potential misspecification for the case of jump uncertainty.

3 Economic framework

We start by considering a simple production-based model with an AK production technology subject to adjustment costs. By design, this formulation has close ties to much of the long-run risk literature with exogenously specified consumption endowments. Prior to introducing climate change, we include two modifications. First, we include an energy input in a way that is mathematically similar to what is used in "DICE" models as developed by Nordhaus (2017) and others. Second, we allow for R&D that could eventually remove the need for energy input. With regard to the second modification, some suggest that an economically viable version of nuclear fusion might achieve this aim.[9]

[8] As noted by Hansen and Miao (2018), this exponential adjustment can equivalently be viewed as a continuous-time version of a smooth ambiguity adjustment.

[9] See Chang (2022) for a recent discussion of progress in the development of an economically viable nuclear fusion technology.

3.1 Production and innovation

The economic component of our model has two endogenous state variables: the stock of productive capital, K_t, and the stock of R&D capital, R_t.

3.1.1 Output

Output is split between consumption, capital investment, I_t^k, and R&D investment, I_t^r (the superscripts denote the investment type):

$$C_t + I_t^k + I_t^r = \alpha K_t \left(1 - \phi_0(Z_t) \left(A_t^b\right)^{\phi_1}\right) \tag{8}$$

for $\phi_1 \geqslant 2$ and $0 < \phi_0(Z_t) \leqslant 1$, where

$$A_t^b \stackrel{\text{def}}{=} \left(1 - \frac{\mathcal{E}_t}{\beta \alpha K_t}\right) \mathbf{1}_{\{\mathcal{E}_t < \beta \alpha K_t\}}, \tag{9}$$

and $\mathbf{1}$ is an indicator function that assigns one to the event in the $\{\cdot\}$ brackets. The jump process Z has an evolution which will be described subsequently.

Emissions $\mathcal{E}_t$ are a proxy for a "dirty" energy input into production.[10] When emissions fall short of the threshold $\beta \alpha K_t$, there is a corresponding convex adjustment in the output given by the right-hand side of (8). This technology is, by design, homogeneous of degree one. For a fixed K_t, the implied production function is flat when emissions exceed the threshold of $\beta \alpha K_t$ and has a zero left derivative at this point. The function equals $1 - \phi_0(Z_t)$ when $\mathcal{E}_t = 0$ and increases up the threshold as a concave function with curvature dictated by the parameter ϕ_1. We feature the case in which $\phi_1 = 3$.

Remark 3.1. *This type of mathematical formulation has showed up in many climate-economics papers since the work of Nordhaus. See, for instance, Nordhaus (2017). The term*

$$\phi_0(Z_t) \left(A_t^b\right)^{\phi_1}$$

is commonly referred to as an abatement cost, perhaps even a cost that is external to the firm. This, however, is only a "cost" in our specification relative to a technology that does

[10] While we could include A_t^b as an entry in the control vector, A_t, this is unnecessary because A_t^b is fully determined emissions and capital.

not require emissions. The right-hand side of relation (8) *shows output increasing in* $\mathcal{E}_t$ *until it reaches a threshold* $\beta\alpha K_t$*, after which additional emissions are unnecessary as capital is the constraining factor. Embedding this* $\mathcal{E}_t \geqslant \beta\alpha K_t$ *implication in a fixed-proportion technology for emissions lower than* $\beta\alpha K_t$*, corresponds to setting* $\phi_0(Z_t) = 1$ *and* $\phi_1 = 1$*. By increasing* ϕ_1 *we increase the production possibilities, as does lowering* $\phi_0(Z_t)$*. In light of this, we view the right-hand side of* (8) *as a production relation without reference to an "abatement cost." This, of course, is a matter of interpretation, but it also impacts how we think of plausible calibrations of* $\phi_o(Z_t)$*.*[11] *We will have more to say about this subsequently.*

We suppose initially that $\phi_0(Z_t) > 0$ and that at some point in the future a fully green technology becomes economically viable, in which case $\phi_0(Z_t) = 0$ and dirty energy is no longer needed to produce output. For instance, think of a substantial advance such as nuclear fusion.[12] Investment in R&D makes this discovery more likely. Specifically, the jump to the clean technology is governed by a Poisson intensity $\mathcal{J}(R_t) = R_t$.

3.1.2 Productive capital evolution

The stock of productive capital, K_t, evolves as

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

where σ_k is a row vector with the same dimension as the underlying Brownian motion. New investment, I_t^k, augments the capital stock, K_t, subject to an adjustment cost captured by the curvature parameter κ. Capital is broadly conceived to include human capital and intangible capital.

3.1.3 R&D capital evolution

A process R captures the stock of R&D-induced knowledge capital and evolves as

$$dR_t = -\zeta R_t dt + \psi_0 \left(I_t^r\right)^{\psi_1} \left(R_t\right)^{1-\psi_1} dt + R_t \sigma_r dW_t, \tag{10}$$

[11] See appendix A.3 for a further discussion based on computing first and second derivatives of a proposed production function.

[12] See Chang (2022) for a recent discussion of the state of this technology and its promise.

where $0 < \psi_1 < 1$ and I_t^r is an investment in research and development. While we will solve a social planner's problem, this evolution equation potentially includes an externality associated with R&D. For pedagogical simplicity, we consider the case of a single technology jump to a fully productive green technology. The parameter ζ captures depreciation in the stock of knowledge pertinent for future technological progress. The term $\sigma_r dW_t$ reflects an exogenous stochastic inflow of information about the future likelihood of a technological advance.[13]

Our decision maker (social planner) has two investment opportunities given our two forms of capital: one is our broad-based notion of a aggregate capital stock, and the other is the stock of R&D devoted to the discovery of a new green technology. The costs of these investments by design are the same, but the payoffs are quite different.

4 Climate dynamics

Here we follow the simplified climate dynamics used in Brock and Xepapadeas (2017) and Barnett, Brock and Hansen (2022). Their approach is based on an approximation from the geoscience literature used to support model comparisons. Specifically, Matthews et al. (2009) and others have purposefully constructed an approximation for climate model output:

$$\text{temperature anomaly } (Y_t) \approx \text{ TCRE}(\theta) \times \text{ cumulative emissions },$$

where TCRE is an acronym for the **T**ransient **C**limate **R**esponse to cumulative **E**missions. This simplified formulation abstracts from transitory "weather" fluctuations in temperature. Instead, emissions today have a long-lasting impact on temperature in the future where TCRE is a measure of climate sensitivity.

Our specific form

$$dY_t = \mathcal{E}_t[\theta(\ell)dt + \varsigma dW_t]$$

for $\theta(\ell) \in \Theta$, where each $\theta(\ell)$ is a TCRE obtained from the set Θ of TCRE's implied by alternative climate models. The term ςdW_t captures short time scale fluctuations. Figure 1

[13]For a recent exploration of the policy implications of R&D for a green breakthrough technology, see Jaakkola and van der Ploeg (2019).

gives a histogram of the θ's that we use in our computations. Appendix A.1 describes how we constructed this histogram.[14] For baseline probabilities across the alternative climate models, we presume that each model has subjective probability $\frac{1}{L_y}$, where L_y is the number of climate models we use as inputs. In other words, we treat each such model as having the same subjective probability. We abstract from parameter learning since learning about such parameters has been slow.[15]

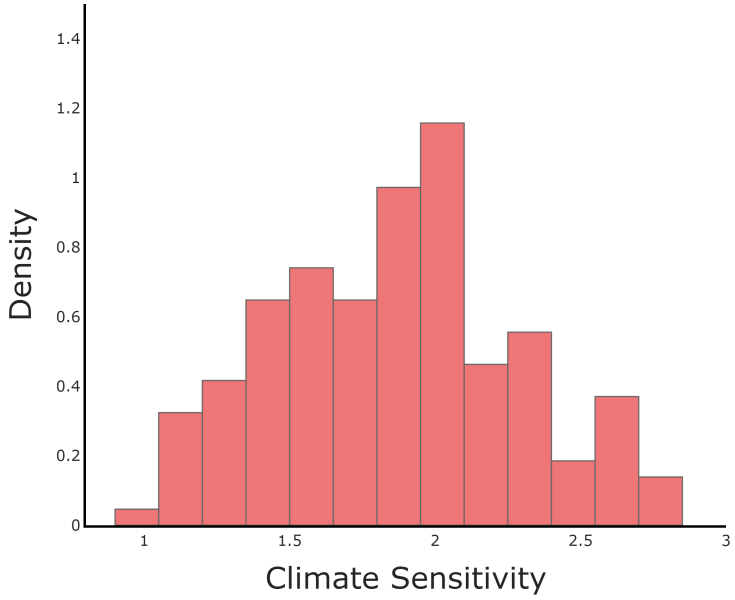


Figure 1: Histograms for the exponentially weighted average responses of temperature to an emissions pulse from 144 different models using a rate $\delta = .01$.

[14] As is well known from the climate science literature, the models actually imply an emissions response that builds from zero to a peak effect in about ten years followed by an approximate flattening at heterogeneous values. See Ricke and Caldeira (2014) for a discussion of these findings. Roughly speaking, the heterogeneous values at which the responses flatten out are equal to the model-specific TCRE's that we use. As we argue in Barnett et al. (2022), these transient dynamics have little impact on the model's implications for policy. Thus here we adopt this simpler specification to avoid having to include an additional state variable.

[15] One could imagine that in the future observations on more extreme temperatures could result in learning becoming more evident.

5 Damages to economic opportunities

We assume that capital, output, both investments, and consumption are all diminished proportionately by N_t. Our damage specification uses a piecewise log quadratic specification as a function of the temperature anomaly y. We suppose that the derivative of the logarithm of damages $\hat{n}$ with respect to the temperature anomaly is

$$\begin{aligned} \frac{d\hat{n}}{dy} &= \lambda_1 + \lambda_2 y & y \leqslant \tilde{y} \\ \frac{d\hat{n}}{dy} &= \lambda_1 + \lambda_2 \left(y - \tilde{y} + \bar{y}\right) + \lambda_3(z_n)(y - \tilde{y}) & y > \tilde{y} \end{aligned} \tag{11}$$

for $z_n \in \{1, 2, ..., L_n\}$. This equation has an initial condition $\hat{n}(0) = 0$. In the stochastic version of what follows, $\tilde{y}$ will be triggered by a Poisson jump prior to a temperature threshold $\bar{y}$. We specify the intensity so that this jump takes place in the interval $[\underline{y}, \bar{y}]$. We shift the derivative of damages with respect to temperature to the right as captured by the change from $\lambda_2 y$ to $\lambda_2 \bar{y}$. We also increase the slope by including a term $\lambda_3(z_n)(y - \tilde{y})$, where the coefficient $\lambda_3(z_n)$ is *ex ante* uncertain.

In Figure 2 we plot the implied damage functions for thresholds $\tilde{y} = 1.5$ and $\tilde{y} = 2$. These two thresholds are often referred to in discussions of the consequences of climate change. Each of the plots includes a range of $\lambda_3(z_n)$'s used in our quantitative policy assessment. The plotted solutions use the formula in Appendix A.2.[16]

As we mentioned previously, there has been considerable heterogeneity in the damage functions used in climate-economic models including a recognition that the curvature is best viewed as uncertain. One *ex ante* reasonable approach would be to activate learning over time of this curvature based on future information.[17] While conceptually appealing (computationally challenging), we find this stylized approach well suited to capture a situation

[16]There are plots for a very similar differential specification in Barnett et al. (2022). Our interpretation of common equations differs, which changes the solutions being plotted. Their construction results in discontinuous damage functions and potential discrete jumps in the damages. Under our construction, the level of damages does not jump at the time of the Poisson event.

[17]Rudik (2020) has a comprehensive treatment of learning and robustness to misspecification within the context of damage function uncertainty. To make this tractable, the unknown damage parameters are learned from a regression of aggregate output growth, absent damages, onto temperature. This type of exercise may open the door to a more comprehensive treatment of learning and a more targeted robustness analysis expressed in terms of prior/posterior ambiguity. Rudik (2020) featured one of the four uncertainty channels that we investigate in this paper.

in which learning is slow now but could be much accelerated once damages become more severe. Since the jump intensity depends on the temperature anomaly, like learning, this dependence introduces a mechanism by which emissions impact information available to the decision maker. Importantly, this allows us to explore the trade-off between acting now versus waiting until more information is available.

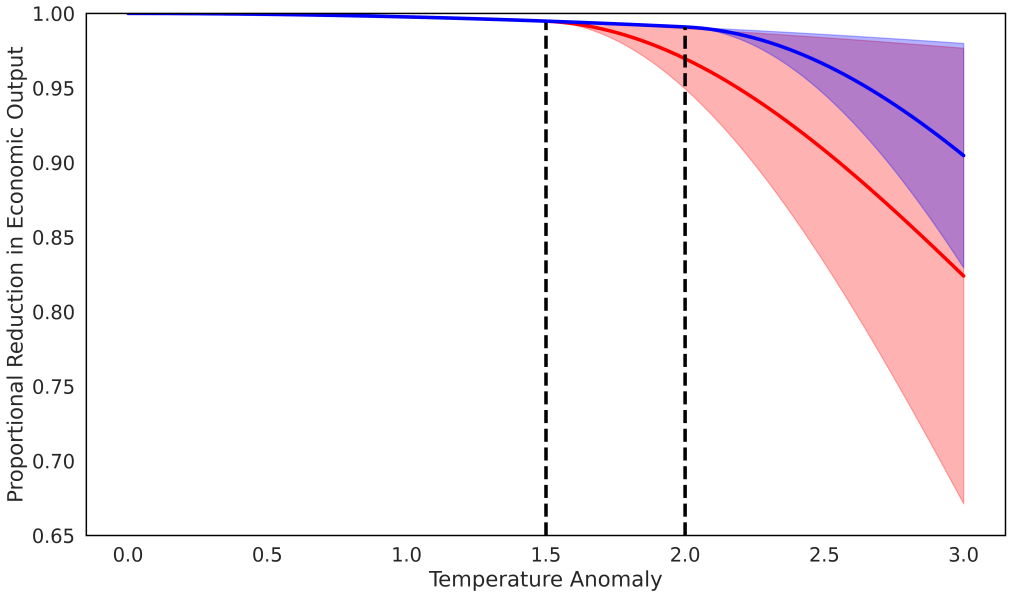


Figure 2: Range of possible damage functions for two cases with different jump thresholds. The solid lines show the average values, and the shaded regions give the range of possible values for $\exp(-n)$, which measures the proportional reduction of the productive capacity of the economy. The blue line and region show the damage function curvature when the jump occurs at $Y_t = \bar{y} = 2.0$. The red line and region show the damage function curvature when the jump occurs at $Y_t = \underline{y} = 1.5$. The black dashed lines indicate the values of Y_t for the upper and lower jump thresholds for the temperature anomalies.

We introduce the $\widehat{Y}$ process to help represent damages, conveniently. This process has a single jump at a random date τ. Prior to this jump, the process coincides with the temperature anomaly, $Y = \widehat{Y}$. We specify a Poisson intensity $\mathcal{J}_n(y)$ for the jump date, τ, so that this jump happens prior to $\hat{Y}_t$ reaching $\bar{y}$ with probability very near one. At the random date τ, the $\widehat{Y}$ process jumps up to $\bar{y}$, but the temperature anomaly does not jump. Thus the $\widehat{Y}$

process post jump ceases to be equal to the temperature anomaly. Instead, $\widehat{Y}_t = Y_t - Y_\tau + \bar{y}$ for $t > \tau$. When $\widehat{Y}$ reaches $\bar{y}$ because of a jump, the local dynamics for the process N shift as captured by $\widehat{Y}_t$. See Figure 3 for a plot of the baseline intensity function used in our analysis.

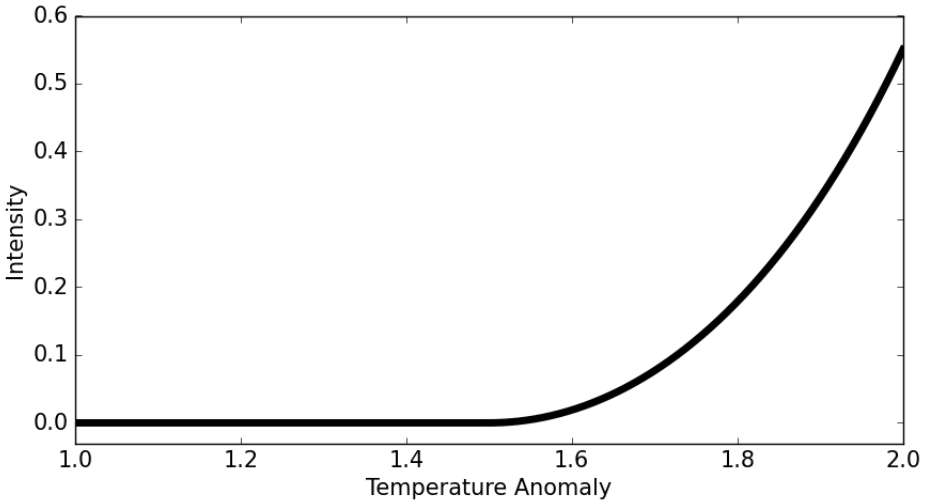


Figure 3: Plot of the intensity function given by $\mathcal{J}_n = r_1 \left(\exp\left(\frac{r_2}{2}(y - \underline{y})^2\right) - 1\right) \mathbf{1}_{\mathbf{y} \geqslant \underline{\mathbf{y}}}$ with $r_1 = 1.5$ and $r_2 = 2.5$. The implied probability of a jump at an anomaly of 1.6 is approximately .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.

As a second source of uncertainty, the jump triggers a move to a more concave region of the damage function rather than a simplistic tipping point with an instantaneous disaster. This incremental concavity is uncertain *ex ante* but known to the planner once the jump takes place. Prior to the jump, we represent each possible value $\lambda_3(z_n)$ for $z_n \in \{1, 2, ..., L_n\}$ with a baseline initial probability distribution $\pi_n(z_n)$.

The implied damage evolution has two branches for the baseline stochastic evolution, one prior to a jump and another after the jump:

$$d \log N_t = \begin{cases} \left(\lambda_1 + \lambda_2 \widehat{Y}_t\right) \left[\mathcal{E}_t \left(\frac{1}{L_y} \sum_{\ell=1}^{L_y} \theta(\ell) dt + \varsigma dW_t\right)\right] & \\ + \frac{\lambda_2 |\varsigma|^2 (\mathcal{E}_t)^2}{2} dt & t \leqslant \tau \\ \left(\lambda_1 + \lambda_2 \widehat{Y}_t + \lambda_3(z_n) \left(\widehat{Y}_t - \bar{y}\right)\right) \left[\mathcal{E}_t \left(\frac{1}{L_y} \sum_{\ell=1}^{L_y} \theta(\ell) dt + \varsigma dW_t\right)\right] & \\ + \frac{(\lambda_2 + \lambda_3(z_n)) |\varsigma|^2 (\mathcal{E}_t)^2}{2} dt & t > \tau, \end{cases}$$

for $z_n \in \{1, 2, \ldots, L_n\}$.

Remark 5.1. *As posed, our model is at the global scale. Note that our specification of a jump is not a "tipping point" discontinuity. Instead the jump triggers a move to a more concave region of the damage function. Existence of global scale tipping points is controversial within the climate literature. For example, see Brook et al. (2013) and Levitan (2013). We suspect that lower-order tipping points become a more salient concern for extended versions of the analysis with regional heterogeneity in the exposure to climate change. We intend our specification of damages only to be an initial platform for addressing the layers of important uncertainties in climate damages and prudent society responses. We find this to be more appealing than the so-called carbon-budgeting approach that imposes a Hotelling-type constraint on emissions to avoid crossing a temperature threshold.*[18]

6 Social planner preferences

We adopt a recursive representation of preferences in continuous time. Start with the special case of no uncertainty. Form the continuation value for each calendar date:

$$V_t = \left(\delta \int_0^\infty \exp(-\delta\tau) \left(C_{t+\tau}\right)^{1-\rho} d\tau \right)^{\frac{1}{1-\rho}}.$$

These preferences are dynamically consistent with a recursive representation, as is evident by raising V_t to the power $1-\rho$. In what follows, we use $\widehat{V} = \log V$ to represent preferences in an ordinally consistent manner.

The following differential equation gives the local representation:

$$\frac{d\widehat{V}_t}{dt} = -\frac{\delta}{1-\rho}\left(\left(\frac{C_t}{V_t}\right)^{1-\rho} - 1\right). \tag{12}$$

[18]Setting a carbon budget in terms of cumulative emissions typically abstracts from the inherent uncertainty in how emissions impact temperature. When it's taken to be a hard constraint, the implied damages when the constraint binds immediately become very substantial in contrast to damage function specifications like ours and others engaged in climate-economic research.

Integrating over a finite-time interval yields

$$\widehat{V}_0 = \frac{\delta}{1-\rho}\int_0^t \left(\left(\frac{C_\tau}{V_\tau}\right)^{1-\rho} - 1\right) d\tau + \widehat{V}_t, \tag{13}$$

which is a backward recursion linking future continuation values to the current one. As a baseline under risk, we use the counterpart to (12) replacing the time derivative of $\widehat{V}$ with its local mean or drift. The conditional expectation counterpart to (13) is

$$\widehat{V}_0 = \frac{\delta}{1-\rho}\mathbb{E}\left[\int_0^t \left(\left(\frac{C_\tau}{V_\tau}\right)^{1-\rho} - 1\right) d\tau \mid \mathfrak{F}_0\right] + \mathbb{E}\left[\widehat{V}_t \mid \mathfrak{F}_0\right]$$

with a local in time representation:

$$0 = \frac{\delta}{1-\rho}\left[\left(\frac{C_t}{V_t}\right)^{1-\rho} - 1\right] + \lim_{\epsilon \downarrow 0}\frac{1}{\epsilon}\left[\mathbb{E}\left(\widehat{V}_{t+\epsilon} \mid \mathfrak{F}_t\right) - \widehat{V}_t\right].$$

The $\rho = 1$ limiting version of this equation is

$$\delta\left(\log C_t - \widehat{V}_t\right) + \lim_{\epsilon \downarrow 0}\frac{1}{\epsilon}\left[\mathbb{E}\left(\widehat{V}_{t+\epsilon} \mid \mathfrak{F}_t\right) - \widehat{V}_t\right].$$

Aversions to ambiguity and model misspecification are captured by adjustments to the drift of $\widehat{V}$ as described in Section 2 for $\widehat{V}_t = \widehat{V}(X_t)$. The HJB equations to be solved include maximization along with the aversion adjustments captured by minimization.

Remark 6.1. *As suggested previously, model misspecification aversion is mathematically equivalent to risk aversion in a recursive utility formulation. This insight holds for discrete-time specifications within the class studied by Kreps and Porteus (1978) and Epstein and Zin (1989). It also carries over to continuous-time recursive utility. This includes commonly used diffusion specifications within the class justified by Duffie and Epstein (1992) extended to incorporate jump risk.*[19]

[19]This mathematical equivalence extends an insight from control theory that establishes a connection between robust and risk-sensitive control. See Hansen and Sargent (2001), Anderson et al. (2003), Hansen and Sargent (2001), and Maenhout (2004) for further elaboration. Special cases of this equivalence first showed up in control theory in the important paper by Jacobson (1973) with many subsequent contributions. With the exception of Hansen and Sargent (1995), however, these papers did not make connections to the

Remark 6.2. *Recursive prior/posterior misspecification aversion in continuous time is mathematically equivalent to a recursive version of smooth ambiguity preferences as in Hansen and Miao (2018). The smooth ambiguity preferences characterized in Hansen and Miao are in turn a continuous-time counterpart to the recursive smooth ambiguity preferences of Klibanoff et al. (2009).*

Remark 6.3. *While we have denoted the continuation value dynamics presuming C is consumption, the actual consumption in our model is C/N since climate change induces proportional reduction in macroeconomic aggregates. We make this adjustment in what follows.*

7 Transformed state and control variables

In our computations we use the following state variables:

$$X_t \stackrel{\text{def}}{=} \begin{bmatrix} X_t^1 \\ \widehat{N}_t \end{bmatrix} \quad \text{where} \quad X_t^1 \stackrel{\text{def}}{=} \begin{bmatrix} \widehat{K}_t \\ \widehat{R}_t \\ \widehat{Y}_t \end{bmatrix},$$

where

$$\begin{aligned} \widehat{K}_t &\stackrel{\text{def}}{=} \log K_t \\ \widehat{R}_t &\stackrel{\text{def}}{=} \log R_t \\ \widehat{N}_t &\stackrel{\text{def}}{=} \log N_t. \end{aligned}$$

We treat the damage jump and technology jump realizations as implying continuation values for the post-jump outcomes. These become inputs into HJB equations prior to the jump. We compute a representation for the continuation values as

$$\widehat{V}(X_t, Z_t) = \widehat{V}^1(X_t^1, Z_t) - \widehat{N}_t\,,$$

where it is straightforward to verify the additive separability in the logarithm of damages. This separability simplifies our numerical solutions. We have three controls: $\frac{I_t^k}{K_t}$, $\frac{I_t^r}{K_t}$, and $\mathcal{E}_t$.

recursive utility literature.

Consumption is then determined by the output constraint (8). We solve the model using an iterative scheme whereby we iterate between maximization and minimization controls and a finite difference solution for the social value function.

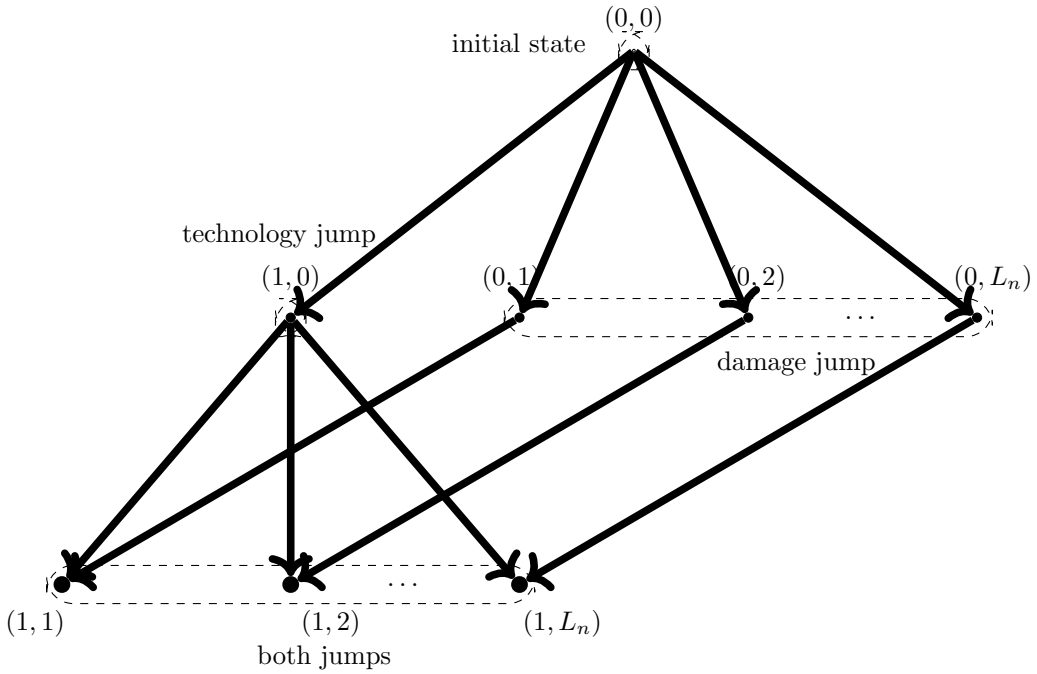


Figure 1: Structure of jump states

Figure 4: Evolution of the jump states. This figure illustrates the dynamic structure of jump states starting from an initial state of all zeros prior to any jump realizations and ending with two jumps being realized. There are two paths to any given endpoint.

Represent the jump states in Z as the ordered pair (z_g, z_n). The first component, z_g, a technology jump state, is equal to zero or one, where zero is the technology state prior to a new discovery and one is the state after the discovery. The second component, z_n, is in the set $\{0, 1, ..., L_n\}$, where the zero is prior to a damage jump and the other L_n realizations depict one of the L_n possible damage curves. Thus there are $2L_n + 2$ possible pairs. The structure and evolution of these states are illustrated in Figure 4.

We start with the transitions from the jump state $z = (0, 0)$. There are $L_n + 1$ possible

transitions. We represent jump probabilistic transitions as

$$\begin{aligned}
\mathcal{J}[x,(0,0)] &= \mathcal{J}_g(r) + \mathcal{J}_n(y) \\
\pi[(1,0) \mid x,(0,0)] &= \frac{\mathcal{J}_g(r)}{\mathcal{J}_g(r) + \mathcal{J}_n(y)} \\
\pi[(0,z_n) \mid x,(0,0)] &= \frac{\mathcal{J}_n(y)}{\mathcal{J}_g(r) + \mathcal{J}_n(y)}\, \pi_n(z_n), \quad z_n = 1,2,\ldots,L_n \\
\tilde{x}[(1,0)] &= x \\
\tilde{x}[(0,z_n)] &= \bar{x}, \quad z_n = 1,2,\ldots,L_n \\
&,
\end{aligned}$$

where $\bar{x}$ replaces y by $\bar{y}$ and leaves the remainder of the entries in x the same and where

$$\sum_{z_n=1}^{L_n} \pi_n(z_n) = 1.$$

Notice that $\mathcal{J}$ is a composite intensity of a first jump, either a technology jump or a jump to one of damage curvature states. Since the first jump could be one of two types, we include the respective probabilities given by the intensity fractions

$$\frac{\mathcal{J}_g(r)}{\mathcal{J}_g(r) + \mathcal{J}_n(y)} \quad \frac{\mathcal{J}_n(y)}{\mathcal{J}_g(r) + \mathcal{J}_n(y)}.$$

The $\pi_n(z_n)$'s give the probabilities for the different values of $\lambda_3(z_n)$ conditioned on the first jump being a damage jump.

The intensity starting from state $z = (1,0)$ is

$$\mathcal{J}[x,(1,0)] = \mathcal{J}_n(y)$$

with a jump distribution

$$\pi_n(z_n), \quad z_n = 1,\ldots,L_n.$$

conditioned on a damage jump happening.

The intensity starting from states $z = (0,z_n)$ and jumping with probability one to state

$z = (1, z_n)$ is

$$\mathcal{J}[x, (1, z_n)] = \mathcal{J}_g(r)$$

for $z_n = 1, 2, \cdots L_n$.

Post-jump value functions are inputs in the pre-jump evolution. Thus we solve for these first, consistent with a dynamic programming perspective. Specifically, we proceed as follows.

i) compute the valuation function, $\widehat{V}[\cdot, (1, z_n)]$ for $z_n = 1, \ldots, L_n$ conditioned on both jumps happening;

ii) compute the value function $\widehat{V}[\cdot, (1, 0)]$ for the $z_g = 1$ realized technology jump allowing for a damage state jump in the future and taking the step i) value functions as input;

iii) compute the value functions, $\widehat{V}[\cdot, (0, z_n)]$ for $z_n = 1, \ldots, L_n$ for each realized damage function allowing for a technology jump in the future and taking the step i) value function as an input;

iv) compute the value function $\widehat{V}[\cdot, (0, 0)]$ prior to any jumps occurring while taking the value functions from steps ii) and iii) as inputs.

8 Social valuation

In this section we characterize the robust social values for the R&D capital stock and for the social cost of global warming. We generalize the approach suggested by Barnett, Brock and Hansen (2020), which imports asset pricing methods into problems of social valuation extended to account for aversion to ambiguity or potential misspecification. Performing computations under a change in probability measure is well known to be mathematically and conceptually convenient for pricing derivative claims. In our application the minimization produces a new probability measure that adjusts for misspecification or ambiguity concerns. Since we apply this to a social planning problem, this change in measure is pertinent for social valuation potentially distinct from market valuation. The derivation of this measure follows directly from the characterizations we described.[20] While Barnett, Brock and Hansen

[20] As Hansen and Miao (2022) note, the argument for valuation adjustments induced by smooth ambiguity is more subtle. Nevertheless, it also has a convenient probabilistic interpretation that builds directly on the solution to the minimization problem we described.

(2020) focused on the uncertainty contribution to the social price of emissions, in this paper we are also interested in the implications of uncertainty for the social cost of global warming and the social value of the stock of R&D directed to new and economically viable green technologies.

8.1 Basic approach

The interpretation of a partial derivative of a value function (scaled by marginal utility) as a shadow value is familiar from the solution to a continuous-time stochastic control problem. Investment theory and "Q-theory" give prime examples of this. We add to this by formally representing the shadow value as a marginal-utility-adjusted expected discounted value of a local impulse response function. Moreover, we extend this insight to allow for the inclusion of a minimizing or adversarial agent introduced to address uncertainty concerns.

The solution to our HJB equation can also be viewed as a solution to the so-called Feynman–Kac (FK) equation expressed in terms of $\widehat{V}$ and its derivatives, with the solutions to the max and min problems substituted in. Solutions to FK equations are typically discounted expected values corresponding to the continuation value, represented as a function of the Markov state, $\widehat{V}$. Using this as a starting point, we construct an equation for a partial derivative of $\widehat{V}$ with respect to the state variable of interest, say, the i^{th} component with realization, x_i given by the i^{th} entry of x:

$$\widehat{F}_{x_i} \stackrel{\text{def}}{=} \frac{\partial \widehat{V}}{\partial x_i}.$$

In Appendix A.4 we provide a representation of the partial derivatives in terms of a forward-looking asset-pricing formula. Let Imp_t^i denote the local stochastic impulse response at date t to a marginal change in the i^{th} component of X_0. We build an asset pricing formula with the following ingredients:

i) the expectation looking forward uses the stochastic evolution implied by the (penalized) worst-case stochastic dynamics;

ii) the discounted stochastic social cash flow has the following two components:

(a) an impulse response process measuring marginal-utility scaled response of damaged consumption to a marginal change in the initial state variable of interest

$$MC_t^i \stackrel{\text{def}}{=} \delta \left[\frac{1}{C_t} \frac{\partial c}{\partial x'}(X_t) - (\mathsf{u}_n)' \right] Imp_t^i, \tag{14}$$

where c is the implied control law for consumption and u_n is a coordinate vector constructed so that $\mathsf{u}_n \cdot X = \widehat{N}$.

(b) an impulse response measuring the marginal impact of the jump intensity on the conditional relative entropy

$$\begin{aligned} ME_t^i \stackrel{\text{def}}{=} & \xi \sum_{z \in \mathcal{Z}} \left[1 - f^*(z \mid X_t, Z_t) + f^*(z \mid X_t, Z_t) \log f^*(z \mid X_t, Z_t) \right] \pi(z \mid X_t, Z_t) \\ & \times \frac{\partial \mathcal{J}}{\partial x'}(X_t, Z_t) Imp_t^i, \end{aligned} \tag{15}$$

where we specialized to the case in which $\rho = 1$.[21]

Contribution MC_t^i, given by formula (14), is the product of the marginal utility of (damaged) consumption times the marginal impulse response of this consumption at date t to a marginal change in the i^{th} coordinate of the state vector at date zero. Contribution ME_t^i, given by formula (15), is special to misspecification considerations with endogenous jump intensities. By reducing the intensity, the planner limits vulnerability to jump misspecification.

Taken together, these contributions give rise to the expected discounted response representation

$$\widehat{F}_{x_i}(x) = \widetilde{\mathbb{E}} \left[\int_0^\infty \exp(-\delta t) \left(MC_t^i + ME_t^i \right) dt \mid X_0 = x \right],$$

where $\tilde{\cdot}$ expectations are the ones implied by *minimizing solution* to the planner's problem.

[21] There is an analogous formula when $\rho \neq 1$, which is more tedious to express.

8.2 Investment alternatives

The social planner has two investment opportunities in our model: investment in new capital and investment in R&D. We investigate the first-order conditions prior to the realization of either technology or damage jump. In both cases we see the role of the shadow value of the corresponding asset stock.

The first-order conditions for investment in new capital are

$$\frac{\partial \widehat{V}}{\partial \hat{k}}(X_t, Z_t)\left(1 - \kappa \frac{I_t^k}{K_t}\right) - \delta\ (C_t)^{-\rho}\,(N_t)^{\rho-1}\,K_t\,\exp\left((\rho-1)\widehat{V}(X_t, Z_t)\right) = 0,$$

where we replace C_t with $\frac{C_t}{N_t}$ in the preference specification. Thus we obtain the formula for investment:

$$\frac{I_t^k}{K_t} = \frac{1}{\kappa}\left[1 - \frac{\delta\ (C_t)^{-\rho}\,(N_t)^{\rho}\,\exp\left((\rho-1)\widehat{V}(X_t, Z_t)\right)}{N_t\,(K_t)^{-1}\,\frac{\partial \widehat{V}}{\partial \hat{k}}(X_t, Z_t)}\right],$$

where the term

$$\frac{N_t\,(K_t)^{-1}\,\frac{\partial \widehat{V}}{\partial \hat{k}(X_t, Z_t)}}{\delta\ (C_t)^{-\rho}\,(N_t)^{\rho}\,\exp\left((\rho-1)\widehat{V}(X_t, Z_t)\right)}$$

is the "Q" from the theory of investment adjusted for damages. In our analysis this is a social valuation, which may be distinct from a marginal valuation.

The first-order conditions for the socially efficient R&D investment are

$$\frac{\partial \widehat{V}}{\partial \hat{r}}(X_t, Z_t)\psi_0\psi_1\left(\frac{I_t^r}{R_t}\right)^{\psi_1-1} - \delta\ (C_t)^{-\rho}\,(N_t)^{\rho-1}\exp\left((\rho-1)\widehat{V}(X_t, Z_t)\right) = 0.$$

Thus

$$\left(\frac{I_t^r}{R_t}\right)^{1-\psi_1} = \psi_0\psi_1 R_t\left[\frac{N_t\,(R_t)^{-1}\,\frac{\partial \widehat{V}}{\partial \hat{r}}(X_t, Z_t)}{\delta\ (C_t)^{-\rho}\,(N_t)^{\rho}\,\exp\left((\rho-1)\widehat{V}(X_t, Z_t)\right)}\right].$$

The term in square brackets is the social value of the knowledge stock of R&D expressed in units of (damaged) consumption.

8.3 First-order conditions for emissions

Prior to both jump realizations, the first-order conditions for emissions are[22]

$$\left[\frac{\partial \widehat{V}}{\partial y}(X_t, Z_t) - \lambda_1 - \lambda_2 Y_t\right]\left[\frac{1}{L_y}\sum_{\ell=1}^{L_y} q(\ell \mid X_t, Z_t)\theta(\ell) + \varsigma H_t\right] + \mathcal{E}_t\left[\frac{\partial^2 \widehat{V}(X_t, Z_t)}{\partial y^2} - \lambda_2\right]|\varsigma|^2$$
$$+ \delta\left(C_t\right)^{-\rho}\left(N_t\right)^{\rho-1}\exp\left((\rho-1)\widehat{V}(X_t, Z_t)\right)\frac{\phi_0\phi_1}{\beta}\left(\frac{A_t^b}{\beta\alpha K_t}\right)^{\phi_1-1}\mathbf{1}_{\{\mathcal{E}_t<\beta\alpha K_t\}} = 0.$$

The implied social cost of carbon is

$$\frac{\left[-\frac{\partial \widehat{V}}{\partial y}(X_t, Z_t) + \lambda_1 + \lambda_2 Y_t\right]\left[\frac{1}{L_y}\sum_{\ell=1}^{L_y} q(\ell \mid X_t, Z_t)\theta(\ell) + \varsigma H_t\right] - \mathcal{E}_t\left[\frac{\partial^2 \widehat{V}(X_t, Z_t)}{\partial y^2} - \lambda_2\right]|\varsigma|^2}{\delta\left(C_t\right)^{-\rho}\left(N_t\right)^{\rho}\exp\left((\rho-1)\widehat{V}(X_t, Z_t)\right)}, \tag{16}$$

and the social benefit is

$$\frac{1}{N_t}\frac{\phi_0\phi_1}{\beta}\left(\frac{A_t^b}{\beta\alpha K_t}\right)^{\phi_1-1}\mathbf{1}_{\{\mathcal{E}_t<\beta\alpha K_t\}},$$

where the formulas are evaluated at the socially efficient emissions and the minimizing $q(\cdot \mid X_t, Z_t)$ and H_t, inclusive of the misspecification adjustment. Again we see a central role for the social valuation of an endogenous state; in this case it is the social cost of global warming given by

$$\frac{-\frac{\partial \widehat{V}}{\partial y}(X_t, Z_t)}{\delta\left(C_t\right)^{-\rho}\left(N_t\right)^{\rho}\exp\left((\rho-1)\widehat{V}(X_t, Z_t)\right)}.$$

Notice also that social cost of carbon (16) includes an explicit volatility adjustment because emissions in our model alter the local exposure to Brownian motion risk.

9 A robustness decomposition

A focal point of our analysis will be the quantitative assessment of alternative sources of uncertainty. To do so, we start by computing solutions to two benchmark problems used as bases for comparison. One is a planner's problem that imposes what Cerreia-Vioglio

[22] See Appendix A.6 for the HJB equation used to deduce the first-order conditions.

et al. (2021) refer to as "misspecification neutrality." Formally this is achieved by imposing $\xi = \infty$. The second imposes a finite ξ entertaining misspecification of all four channels, i)–iv), described in the introduction, simultaneously. The uncertainty decompositions that we implement explore intermediate problems with alternative limits on the forms of misspecification.[23] For instance, we will activate each of the four uncertainty channels separately and compare the results with the two benchmarks.[24] [25]

More formally, recall the vector of possible Brownian misspecifications for dX_t:

$$\sigma(X_t, A_t)H_t dt.$$

To focus the misspecification on a component of the state vector, restrict h so that

$$\mathbb{U}\sigma(x, a)h = 0$$

in subproblem 2.1 , where $\mathbb{U}$ selects the components of the state vector that are immune from misspecification concerns.

For the jump process contribution, isolate the discrete jump state dynamics that are the targets of robustness. For instance, these could be the discovery of a new green technology represented as changes in the first coordinate of z, or they could be transitions to a damage curve represented by the second state. We capture both of these as restrictions on the potential distortion functions $f(\cdot \mid x, z)$ in a constrained version of subproblem 2.4 that includes only the targets of robustness when forming an HJB equation.

By considering alternative specifications on (h, f) in the minimization problem, we obtain

[23]Recently, Cappelli et al. (2021) have explored the source dependence of uncertainty from an abstract decision-theoretic perspective, but without reference to misspecification aversion.

[24]Ricke and Caldeira (2014) implemented a conceptually distinct but similarly motivated approach as displayed in Figure 3 of their paper. Ricke and Caldeira featured three alternative sources of uncertainty that contribute to climate change, where uncertainty is equated to model heterogeneity. Whereas they eliminate the uncertainty from each channel in their decomposition, we close down the *aversion* to uncertainty within a decision-making framework. We also focus on different channels, including damages and technological uncertainty. But we do not look at a more refined decomposition of temperature dynamics as Ricke and Caldeira (2014) do.

[25]Our approach in this paper differs from the one used in our previous work Barnett et al. (2020) and Barnett et al. (2022). There we featured decompositions of the social cost of carbon based on a Feynman–Kac approach. Moreover, in these earlier papers we did not explore the impact of endogenous R&D as we do here.

the decompositions of interest. We use the solutions of the resulting HJB equations to ascertain the relative importance of each of the component misspecifications, seeing which ones get us closest to the full robust solution relative to the misspecification neutral solution. Given the use of minimization and the possible lack of separability of the local entropy across the different states, this uncertainty decomposition is not additive. That is, the sum of the decrements from the solution under misspecification neutrality will not sum to the reduction implied under full robustness.

We could construct an analogous ambiguity-based uncertainty decomposition if, for instance, the different coordinates of a θ vector correspond to different state transitions. We may then solve alternative constrained minimization problems for q by focusing on alternative coordinates of the parameter vector. For the decompositions of interest, we could then explore the impact of misspecifying the distribution over each of the coordinates of interest.

10 Example economy

In this section we explore the implications of our analysis for an example economy. We focus exclusively on misspecification aversion for three penalty parameter values, $\xi = .15$, $\xi = .075$, and a baseline of misspecification neutrality $\xi = \infty$. We refer to the higher value of ξ as "less aversion" and the lower value as "more aversion."[26] To appreciate the impact of these settings, we display consequences for distorting probabilities, an approach that is commonly employed for robust Bayesian methods. After inspecting these distortions, we explore implications for social valuation and its consequences for R&D emissions, and climate change. Parameter values that we use for this economy are reported in Appendix A.5.

The calculations convey the following insights from our social planner solution:

- Including endogenous R&D is a particularly important tool for the planner when confronting the economic consequences of climate change.
- The impact of *technological uncertainty* dominates those of climate and damage uncertainty.

[26]For the corresponding risk aversion parameters in a recursive utility specification of preferences, less aversion is $\gamma \approx 7.7$, more aversion is $\gamma \approx 14.3$, and neutrality as $\gamma = 1$.

- R&D investment is substantially more sensitive to uncertainty than are reductions in emissions.

The subsections that follow give graphical characterizations of many of our salient findings.

10.1 Uncertainty-adjusted probability distributions

We consider in turn the implied worst-case probabilities that emerge from our analysis for the four sources of uncertainty: climate, damages, productivity, and technology. These are solutions to the minimization problem, evaluated at solutions to the maximization problem. These are not "best-guess" distributions. Rather they are the "worst-case" distributions subject to penalization that are a vehicle by which the planner constructs robustly optimal courses of action. We report results for two specifications of the penalty parameter: $\xi = .075$ (more aversion) and $\xi = .15$ (less aversion). It is straightforward to run our solution code for other choices of these parameters. As we noted previously, these could be reinterpreted as recursive utility risk aversion parameters, although we prefer the robustness motivation.[27] While it is hard to interpret directly the magnitudes of ξ, we find it valuable to adopt an approach from robust Bayesian methods by inspecting worst-case probability specifications isolating what probability specifications the robust decision rules are optimal with respect to.

We begin the Brownian motion risk distortion prior to any Poisson jumps. These are captured by changes in local mean or drift and impact only three of our four channels. We abstract from a Brownian contribution to damages. We report these as conditional mean changes in a multivariate standard normal distribution. The results are reported in Table 1 for each of the three endogenous state variables. The computed "worst-case" drift distortions are modest, since the Brownian increments are standardized. With a sign change for capital and the stock of knowledge, these translate into instantaneous uncertainty compensations for social valuation. The most notable one is for the shock to the capital stock, but with "more aversion" this is a mean compensation of .12 for a standardized normal shock, and for "less aversion" this compensation is only .06. The adjustments for the two other shocks distributions are much smaller and not noteworthy.

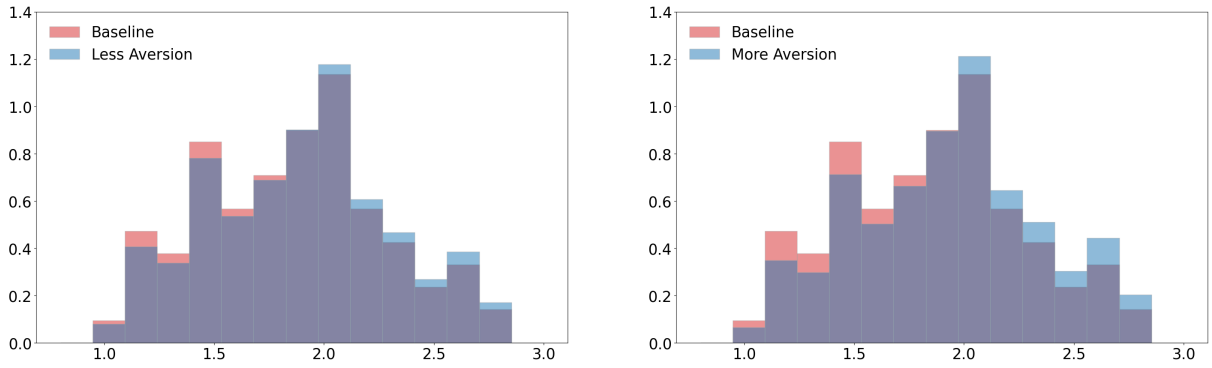
[27]The corresponding recursive utility risk aversion parameters are 12.3 and 5.67.

Recall that in addition to the Brownian shocks, our model has jump contributions to both damages and technology. As we will see, the jump uncertainties are much more prominent in our calculations.

Drift distortion for the endogenous state variables			
uncertainty	capital	temperature	knowledge stock
more aversion	-.122	.027	-.006
less aversion	-.063	.008	-.002

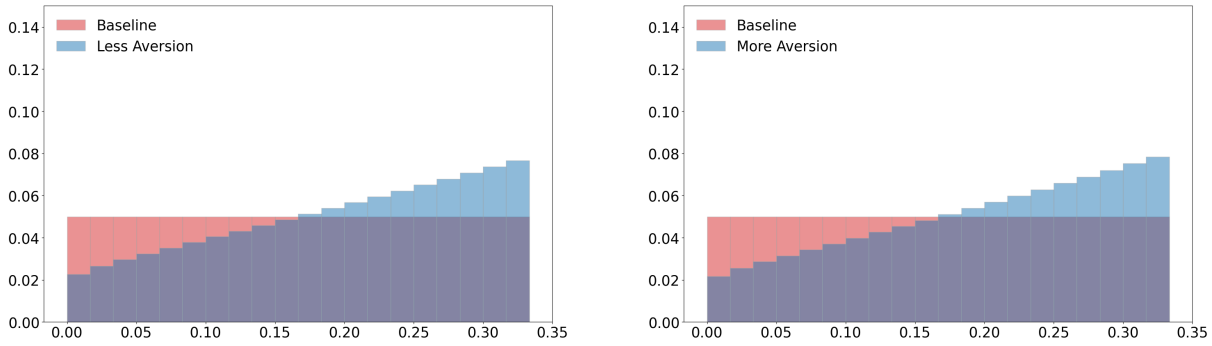
Table 1: Conditional mean distortions to Brownian increments for the productive capital, K_t, the temperature anomaly, Y_t, and knowledge capital, R_t. The table reports numbers both under more aversion ($\xi = .075$) and less aversion ($\xi = .15$) to potential model misspecification. The conditional mean shifts are calculated in the initial period.

From Figures 5 and 6 we see that uncertainty aversion distorts the probability distributions used by the planner for the climate and damage models away from the low values and toward the higher values of θ and λ_3. The effect is much more modest for the climate models than the damage models and, as expected, is more pronounced as the aversion is increased (as the penalty parameter, ξ is decreased.) Figure 7 shows the probability of a technology or damage jump occurring during the simulation. The clear effect here is that uncertainty has little impact on the expected probability of a damage jump occurring but substantially reduces the probability of a technology jump occurring used by the uncertainty averse planner in our model. These results highlight technology uncertainty as a key force influencing robustly optimal policy choices of the planner.



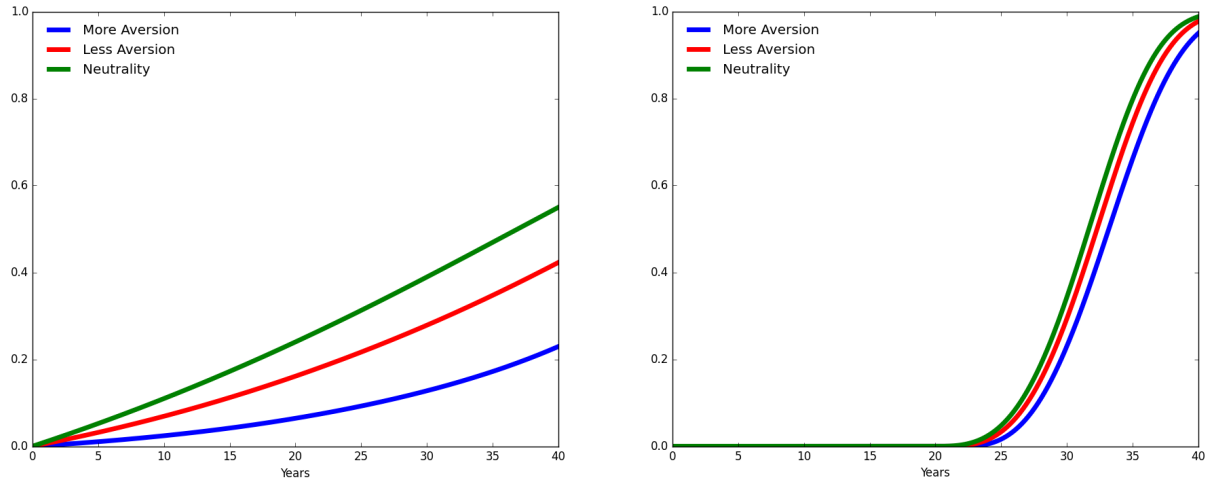
(a) Less Aversion (b) More Aversion

Figure 5: Distorted climate model distribution. The left plot shows the undistorted and distorted distributions with less misspecification aversion. The right plot shows the undistorted and distorted distributions with more misspecification aversion. The histograms are calculated at year 0, with the trajectories leading up to year 0 simulated under the baseline probabilities and abstracting from the intrinsic randomness. Recall that Brownian motion misspecification aversion induces a shift in the local mean. The distorted histograms were constructed by imposing instead ambiguity aversion where the ambiguity penalization was set to capture the same unfavorable conditional mean distortion as the misspecification aversion. The "no aversion" mean of this distribution is 1.86; with less aversion it is 1.91, and with more aversion it is 1.94.



(a) Less Aversion (b) More Aversion

Figure 6: Distorted damage model distribution. The plot shows the undistorted and distorted distributions with less misspecification aversion in the left panel and more misspecification aversion in the right panel. The histograms are calculated at year 0, with the trajectories leading up to year 0 simulated under the baseline dynamics abstracting from the intrinsic randomness.



(a) Uncertainty Aversion Comparison

(b) Uncertainty Aversion Comparison

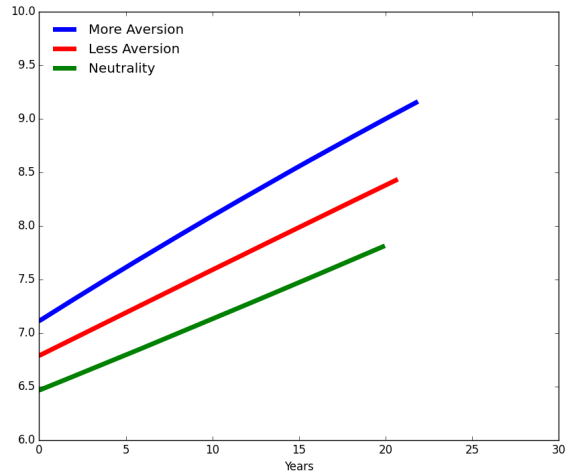
Figure 7: Simulated pathways of the distorted probability of a technology change jump (left plot) and a damage function jump (right plot) occurring. The plot compares outcomes for different values of the misspecification aversion parameter ξ. The state variable trajectories are simulated under the baseline dynamics abstracting from the intrinsic randomness.

10.2 Social valuations

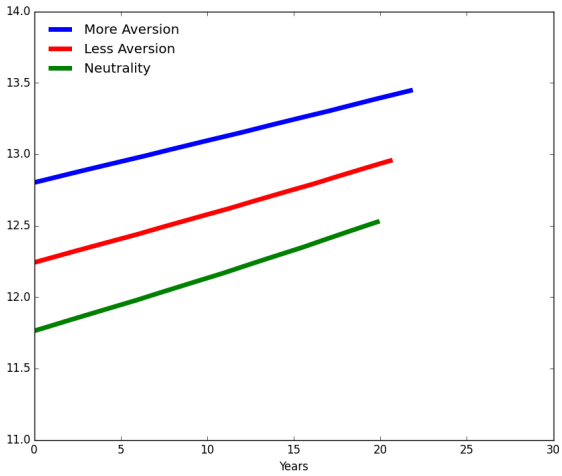
When reporting social valuations, we divide the quantities by the marginal utility of (damaged) consumption) prior to taking logarithms. We perform this scaling so that the units are expressed in terms of this consumption numeraire.

Figure 8 shows the implied trajectories for the social value of R&D and for the social cost of global warming. From the left panel of Figure 8 we see about a 25% increase in the social value of R&D for less aversion and an increase in excess of 50% for more aversion relative to the more usual specification that abstracts from misspecification concerns. The findings tell a similar story for the social costs of global warming, except that the uncertainty adjustments are twice as large. The trajectories in both of these figures condition on no jumps in the technology and the trajectories end prior to any possible damage jump realization.

Table 2 gives the uncertainty decompositions for both the social value of R&D and the social cost of global warming. For both of these computations, the technological uncertainty accounts for the bulk of the uncertainty enhancement. For instance, there is only a 6% reduction in the social value of R & D and a 9% reduction in the social cost of global



(a) Uncertainty Aversion Comparison



(b) Uncertainty Aversion Comparison

Figure 8: Simulated pathways of the logarithm of the social valuations from the model. The left plot shows the trajectory for logarithms of the social values of R&D, and the right plot shows the trajectory of the logarithms of the social costs of global warming. Both plots compare outcomes for different values of the misspecification aversion parameter ξ. The trajectories are constructed from the baseline dynamics abstracting from the intrinsic randomness. The pathways stop when the temperature anomaly reaches $1.5^{\circ}C$.

warming when only the technology uncertainty channel is activated.[28]

Uncertainty channel	$\Delta \log SVRD$	$\Delta \log SCGW$
climate uncertainty	.31	.45
damage uncertainty	.30	.43
productivity uncertainty	.33	0.48
technology uncertainty	.06	.09

Table 2: Uncertainty decompositions for social valuations for the four different channels under the"less aversion" model specification. The table entries are the difference between the logarithms of the values or costs when all four channels are activated minus the logarithms for each of the different channels. The numbers are for the initial time period.

[28]The contributions are not constructed to be additive. That is, the fractions need not add up to one, which is particularly apparent for the social cost of global warming.

10.3 Robust actions

This subsection explores implications for robustly optimal actions. Actions and values are closely tied via first-order conditions, but it remains revealing to show the outcomes for both R&D and emissions. We explore these in Figures 9 and 10, respectively. The figures include both zero shock trajectories and five stochastic trajectories illustrating some of the different realizations of the Poisson events:

1. no jumps for forty years;
2. damage jump at year 28 is accompanying by a low damage curve realization;
3. damage jump at year 36 is accompanied by a high damage curve realization;
4. technology jump to a clean technology at year 26;
5. technology jump to a clean technology at year 20.

The reported paths were computed under the preference specification with less aversion.

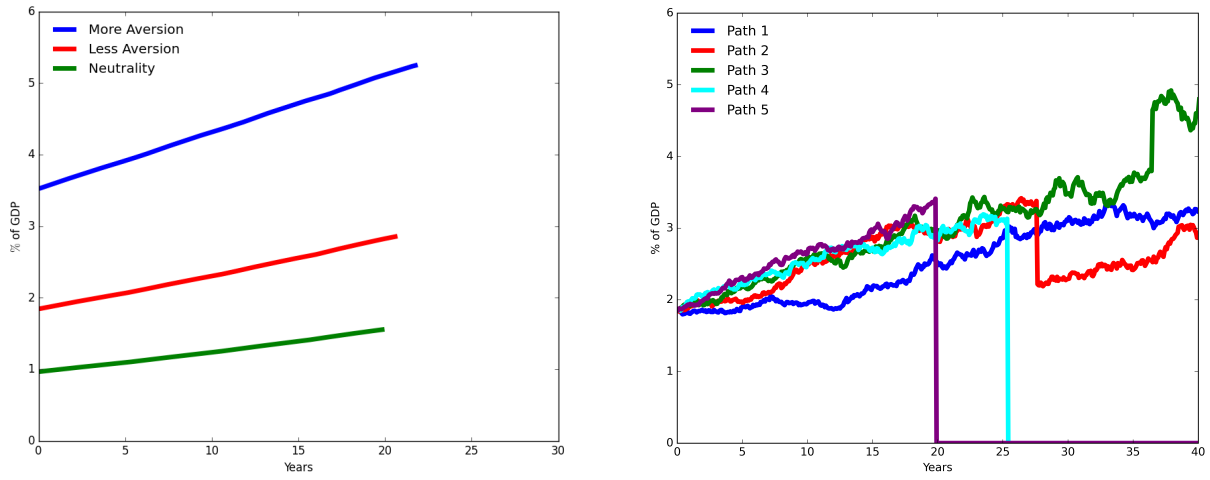
The left panel of Figure 9 shows a substantial increase in R&D investment in the presence of misspecification aversion. Even with the more modest specification of this aversion, the investment is almost doubled. Perhaps not surprisingly, technology uncertainty dominates the other channels, accounting for about 90% of uncertainty increase in the R&D investment.[29] We report the stochastic simulations in right panel of Figure 9, which illustrate some of the impacts of the Poisson events on the R&D investment trajectories.

The increase in R&D investment as potential misspecification aversion increases is striking. The planner reduces the investment-output ratio to offset the increase in R&D. Since R&D investment is much smaller than investment in the productive stock, the proportional reduction in the latter investment is much smaller. Since R&D investment would seem more exposed to uncertainty, as reflected by our Poisson jump specification of discovery, than is the productive capital stock, this R&D investment increase might seem surprising. Notice that this is "big project" R&D analogous perhaps to the Manhattan project or the Apollo program. Moreover, the uncertainty in the R&D investment is about the timing of a successful outcome or social payoff. In this setting, more R&D leads to an increased likelihood that

[29] In contrast, the other sources by themselves account for somewhere between 33 and 40%.

the new, clean technology will be discovered sooner, in contrast to say a standard portfolio allocation problem.

To push this analysis further, we verified numerically that a sufficiently large increase in the aversion to misspecification uncertainty (a sufficiently small reduction in ξ) reduces the R&D investment-output to essentially zero with implied worst-case probabilities that make a discovery very unlikely for the relevant decision-making horizons of our planner. Thus the impact of aversion impact is not monotonic. It increases, however, over a range that we find to be interesting and plausible.



(a) Uncertainty Aversion Comparison

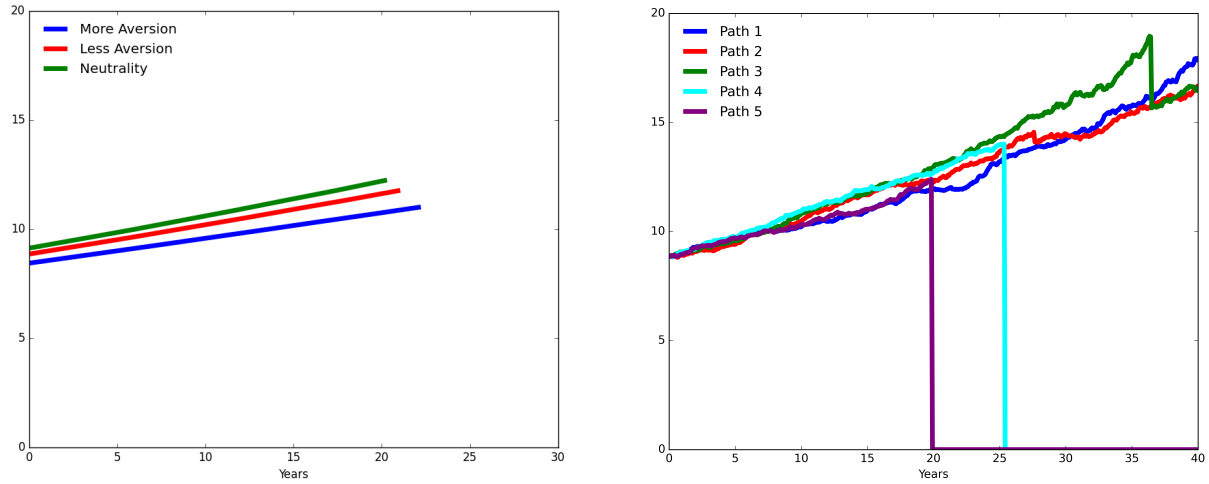
(b) Stochastic Simulation Pathways

Figure 9: Simulated pathways of R&D investment as a fraction of output. The left plot compares outcomes for different values of the misspecification aversion parameter ξ. The trajectories are simulated under the baseline transition dynamics abstracting from the randomness. The pathways stop when the temperature anomaly reaches $1.5^{\circ}C$. The right plot depicts the R&D investment as a fraction of output for the five illustrative stochastic paths under less aversion.

We next consider emissions. Panel A of Figure 10 shows that while emissions are reduced when the fictitious planner is more concerned about potential model misspecification, the impact is quite muted. The secular increase in the emission trajectories occurs for two reasons. First, they condition on the absence of a discovery of a new, economically viable, green technology. Second, our production technology requires an energy input. Panel B of Figure 10 shows some stochastic simulations, with the dramatic impact of the technology

discovery allowing for emissions to drop to zero along some trajectories, including one that happens in year twenty. In addition, when a damage jump occurs before the technology jump, the path of emissions can shift down for a bad realization of the damage curvature γ_3, as in path 3, or possibly stay the same or even increase for a relatively good γ_3 realization, as in path 2.

We conclude this subsection by looking at the consequences for temperature along the illustrative stochastic trajectories as reported in Figure 11. Notice that in the next forty years only one of the trajectories approaches a two degree anomaly for a prudent social planner with modest misspecification aversion. This highlights that relatively cautious, though increasing, emissions pathways are chosen to try to avoid the most severe climate change consequences.



(a) Uncertainty Aversion Comparison

(b) Stochastic Simulation Pathways

Figure 10: Simulated pathways for emissions. The left plot compares outcomes for different values of the misspecification aversion parameter ξ. The trajectories are simulated under the baseline transition dynamics abstracting from the randomness. The pathways stop when the temperature anomaly reaches $1.5^{\circ}C$. The right plot displays five illustrative paths of stochastic simulations for emissions under less aversion.

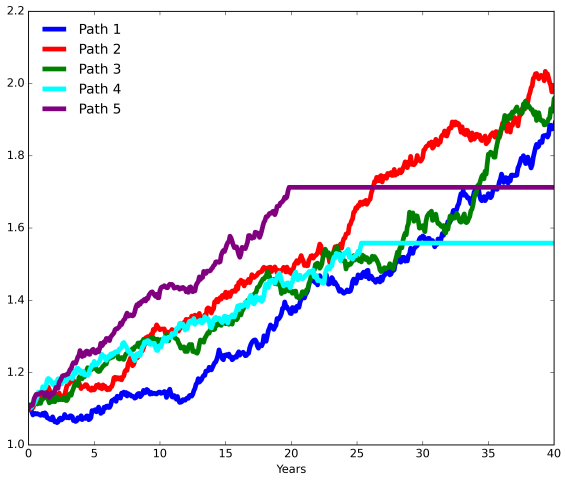


Figure 11: Stochastic simulation for temperature anomaly under lower aversion. This figure displays temperature trajectories for five illustrative paths simulations under less aversion.

	R&D-output ratio		emissions	
Uncertainty aversion	$\phi_0 = .5$	$\phi_0 = .1$	$\phi_0 = .5$	$\phi_0 = .1$
More Aversion	3.52%	1.20%	8.44	6.75
Less Aversion	1.85%	0.94%	8.86	7.43
Neutrality	0.97%	0.44%	9.14	8.18

Table 3: Initial robust actions for two alternative initial specifications of ϕ_o.

Remark 10.1. *As we noted in Remark 3.1, the presentation of environmental economics models of climate change often make reference to an "abatement technology." We instead start with a production function that includes a specific role for fossil fuels as an input. There is a loss in output when this input is reduced, which could be labeled as abatement. In the extreme case in which* $\mathcal{E}_t = 0$*, the fraction of output remaining is* $1 - \phi_0(Z_t)$*. Previous work by Nordhaus and others assumes* $\phi_0 \leqslant .1$ *net of technological progress. In our example,* ϕ_0 *is initially* .5*, which implies a substantially larger measure of the production cost of decarbonization than other literature, albeit one that we find to be more substantively relevant. The results are indeed sensitive to this aspect of our example specification.*

Since the output loss induced by a reduction in emissions would be substantially lower when ϕ_0 *is initialized at .1 instead of .5, this change leads to i) substantially less R&D as a fraction of output and ii) a more modest but notable proportional reduction in emissions.*

We illustrate this in Table 3.

11 Preference sensitivity

In preceding sections we have featured the sensitivity of social valuation and policy outcomes to changes in the aversion to misspecification, aversion that is reflected in implied uncertainty-adjusted probability distributions. In this section we study sensitivity to two other aspects of preferences. We initially explore the impact of increasing the subjective rate of discount on preferences. We then explore changes to the intertemporal elasticity of substitution (IES) in preferences.

11.1 Increasing the subjective rate of discount

The prior environmental economics literature has explored sensitivity of the SCC to changes in the "discount rate." Often these analyses feature the implied discount rate used in computing present values abstracting from stochastic discounting. In our setting with misspecification or ambiguity aversion, stochastic discounting is a central ingredient in valuation as is expected from our understanding of asset pricing. In particular, we have featured how uncertainty aversion preferences impact an endogenously determined change in the probability measure that is pertinent for valuation. So far, we have kept fixed the subjective rate of discount used by the social planner by setting $\delta = .01$. We show how changes in this subjective discount rate alters the social value of the R&D stock and the social cost of global warming. In Table 4 we report computation for both the social value of R&D and the social cost of global warming. Since the reported numbers are in a logarithmic scale, we see drops in the valuations by about 30% by increasing δ from .01 to .015, confirming a sensitivity that is often noted in the environmental economics literature.

Subjective discount rate	$\log SVRD$	$\log SCGW$
$\delta = .010$	6.79	11.61
$\delta = .015$	6.46	11.33
$\delta = .020$	6.07	11.01

Table 4: Social values at the initial time period for less aversion to misspecification uncertainty.

11.2 Intertemporal elasticity of substitution (IES)

Our numerical results have so far featured the case where the IES, $\frac{1}{\rho}$, is unitary. We now consider two other specifications of the IES: $\rho = 2/3$ and $\rho = 3/2$. Much of the asset pricing literature has studied the consequences of changing the IES on asset valuation within the setting of an endowment economy. Our economy is a production economy, however, and changing the IES has a big impact on production outcomes, as we show in Table 5. As expected from growth models, the investment in both types of capital relative to output are higher when the elasticity is greater (ρ is smaller). For instance, in the initial time period R&D investment is more than doubled when we increase the IES to 3/2, with just the opposite effect when we decrease the IES to 2/3. Table 6 shows that the valuations of the corresponding capital stocks move in the opposite way, as expected.

	R&D-output ratio			Investment-output ratio		
Uncertainty aversion	$\rho = 2/3$	$\rho = 1$	$\rho = 3/2$	$\rho = 2/3$	$\rho = 1$	$\rho = 3/2$
more aversion	7.6%	3.5%	1.5%	88.7%	74.3%	66.5%
less aversion	4.6%	1.8%	0.8%	92.3%	76.0%	67.3%
neutrality	2.8%	1.0%	0.4%	94.6%	77.0%	67.6%

Table 5: Initial investment to output ratios for different specifications of the IES.

IES	$\log SVRD$	$\log SCGW$
$\rho = 2/3$	4.89	9.71
$\rho = 1$	6.79	11.61
$\rho = 3/2$	6.86	11.75

Table 6: Social values at the initial time period for less aversion to misspecification uncertainty.

Figure 12 gives the implied emissions trajectories for the two alternative specification of the IES. We can see some modest differences at the outset and marked differences in the slope of the trajectories. For emissions, the initial values across ξ values for the case of $\rho = 2/3$ are between 8 and 9 GtC, lower than the initial emissions across ξ values for the case of $\rho = 3/2$, which are between 8.75 and 9.25 GtC. Over time, the emissions increase when $\rho = 2/3$ to above 12 and 14 GtC, as compared with the range of 10.5 to 12 GtC when $\rho = 3/2$. The growth sensitivities in emissions prior to any jumps are consistent with the

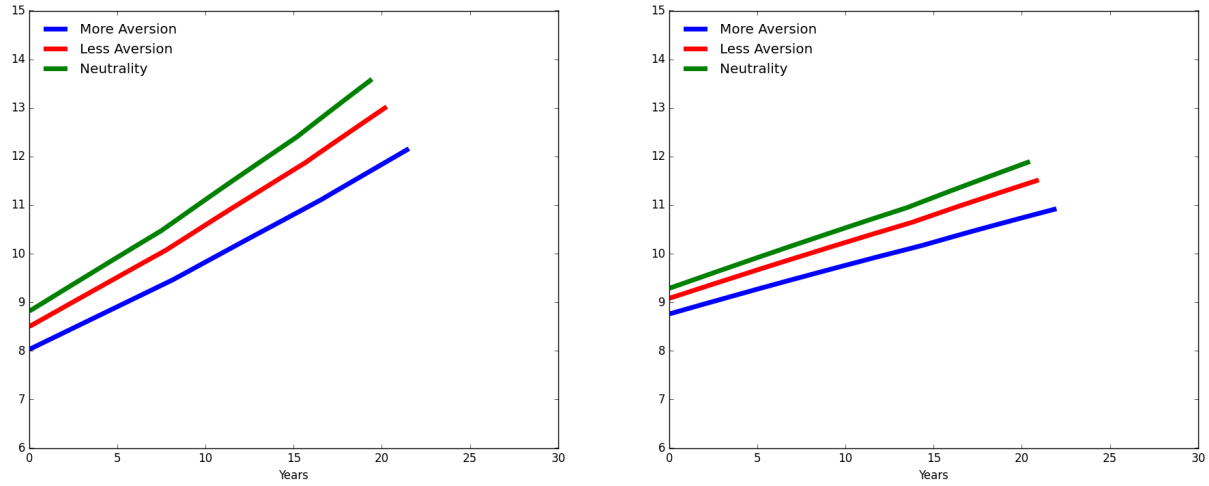


Figure 12: Simulated pathways of emissions for alternative values of ρ. The left panel shows outcomes for the case of $\rho = 2/3$ and the right panel for $\rho = 3/2$. The green, red, and blue lines show outcomes for different degrees of misspecification aversion. The trajectories are simulated under the baseline probabilities abstracting from the intrinsic randomness. The pathways stop when the temperature anomaly reaches $1.5^{\circ}C$.

investment differences reported in Table 5, since higher investment helps support more initial growth in output.[30]

12 Conclusions

Our findings emphasize the potential importance of including endogenous R&D investment into quantitative assessments of socially prudent courses of action along with current and near-term future carbon reductions. Our calculations expose the limitations of commonly proposed policy solutions that entail a gradual decrease of emissions to a net zero target. Further research is needed to explore robustness of such gradualist approaches to confronting the threat of climate change. In addition, R&D provides our fictitious social planner with a

[30]When discussing our paper, Eric Renault reminded us that there can be seemingly counterintuitive interactions between the IES and the risk aversion in recursive utility models, noting that the latter is not a "pure" risk aversion parameter. Indeed in dynamic stochastic settings, the intertemporal composition of risk comes into play when exploring the preference implications. See Cai and Lontzek (2019) and Hambel et al. (2021) for related discussions when exploring the SCC. The changing implications for consumption and investment induced by changes in the IES make the risk aversion comparisons all the more tricky.

second investment opportunity. More uncertainty in the payoff to this investment can result in an increase in the R&D investment relative to output needed to accelerate the possible discovery of a new economically viable green technology. This provides a substantively important example when *more aversion* to potential model misspecification leads to *bolder actions* on the part of the decision maker.

We end by giving our opinionated and speculative discussions of more pragmatic policy challenges. A common objection to taking strong immediate action on climate change by skeptics who accept the climate science behind global warming is that the costs of using inefficient government to confront the problem undermine attractiveness of public as opposed to market solutions. A blunt way of putting this objection is, "This problem won't be solved by government throwing money at it." While there are good reasons to be skeptical about political processes undermining the attractiveness of collective action, our robustness calculations suggest the uncertainties broadly conceived may be large enough to push for efforts to overcome political distortions and embark on investment directed to the discovery of economically viable clean alternatives. What is missing in our analysis is a more serious probe into the political economy of large-scale public investment projects. Nevertheless, our formulation and robustness analysis features the potential for stimulating R&D investment in truly new technologies rather than subsidies that create inefficiencies and special interest rent seeking.

A Appendix

Below we provide the details and derivations for various results in the main text. As in the main text, throughout the appendix we let lower-case variables capture potential realizations of random vectors. Additional plots and figures can be found in our online notebook:

`https://climatesocialpolicy.readthedocs.io/en/latest/index.html`.

A.1 Climate model uncertainty

We construct 144 different TCRE's by using 100 GtC pulse experiment results of Joos et al. (2013) tracing out the resulting carbon in the atmosphere for 9 different models. We then use these as inputs into 16 model approximations for temperature responses using the approximation in Geoffroy et al. (2013) to build the collection of θ_ℓ's used in our analysis.

A.2 Damage function

The solution to the differential equation (11) is

$$\hat{n}(y) = \begin{cases} \lambda_1 y + \frac{\lambda_2}{2} y^2 & y \leqslant \tilde{y} \\ \lambda_1 y + \frac{\lambda_2}{2}(\bar{y} + y - \tilde{y})^2 + \frac{\lambda_3(z_n)}{2}(y - \tilde{y})^2 - \frac{\lambda_2}{2}(\bar{y})^2 + \frac{\lambda_2}{2}\tilde{y}^2 & \tilde{y} < y, \end{cases}$$

or equivalently

$$\hat{n}(y) = \begin{cases} \lambda_1 y + \frac{\lambda_2}{2} y^2 & y \leqslant \tilde{y} \\ \lambda_1 y + \lambda_2 \bar{y}(y - \tilde{y}) + \frac{\lambda_2 + \lambda_3(z_n)}{2}(y - \tilde{y})^2 + \frac{\lambda_2}{2}\tilde{y}^2 & \tilde{y} < y. \end{cases}$$

A.3 Production function verification

We verify that the first derivatives are positive and that the second-derivative matrix is negative semi-definite when we interpret

$$\alpha k \left(1 - \phi_0(z) \left(\iota\right)^{\phi_1}\right)$$

for

$$\iota = \left(1 - \frac{e}{\beta \alpha k}\right) \mathbf{1}_{\{0 \leqslant e \leqslant \beta \alpha k\}}$$

as a production function.[31] Notice that the candidate production function is homogeneous of degree one in (e, k).

First, we consider the partial derivatives with respect to e:

$$\begin{aligned}
\frac{\partial}{\partial e}\text{output} &= \alpha k\, \phi_0(z)\, \phi_1\, (\iota)^{\phi_1 - 1}\, \frac{1}{\beta \alpha k} \\
&= \frac{\phi_0(z)\, \phi_1}{\beta}\, (\iota)^{\phi_1 - 1} \\
&> 0 \\
\frac{\partial^2}{\partial e^2}\text{output} &= -\frac{\phi_0(z)\phi_1(\phi_1 - 1)}{\beta^2 \alpha k}\, (\iota)^{\phi_1 - 2} \\
&< 0.
\end{aligned}$$

If $\phi_1 > 2$, both derivatives are zero at $e = \beta \alpha k$. This remains true for $e > \beta \alpha k$.

Next we consider derivatives with respect to k:

$$\begin{aligned}
\frac{\partial}{\partial k}\text{output} &= \alpha \left(1 - \phi_0(z)\, (\iota)^{\phi_1}\right) - \phi_1 \phi_0(z) \alpha k\, (\iota)^{\phi_1 - 1} \left(\frac{e}{\beta \alpha k^2}\right) \\
&= \alpha \left(1 - \phi_0(z)\, (\iota)^{\phi_1}\right) - \phi_1 \phi_0(z)\, (\iota)^{\phi_1 - 1} \left(\frac{e}{\beta k}\right) \\
\frac{\partial^2}{\partial k^2}\text{output} &= -\frac{\phi_0(z)\phi_1(\phi_1 - 1)}{\beta^2 \alpha k}\, (\iota)^{\phi_1 - 2} \left(\frac{e}{k}\right)^2 \\
&< 0.
\end{aligned}$$

The first derivative is $\alpha(1 - \phi_0) \geqslant 0$ when $k \to \infty$ and $\alpha > 0$ when $\beta \alpha k \leqslant e$. Given the negative second derivative, the first derivative remains positive for $k > 0$.

The simple relationship between the second derivatives with respect to e and k is to be anticipated, since the first derivatives are homogeneous of degree zero. Consistent with this relationship, the cross partial is

$$\frac{\partial^2}{\partial e \partial k}\text{output} = \frac{\phi_0(z)\, \phi_1(\phi_1 - 1)}{\beta^2 \alpha k}\, (\iota)^{\phi_1 - 2} \left(\frac{e}{k}\right).$$

[31] For notational simplicity, we drop the dependence of ϕ_0 on z.

The negative semi-definite Hessian matrix follows since

$$\begin{aligned}\begin{bmatrix}\mathsf{r}_1 & \mathsf{r}_2\end{bmatrix}\begin{bmatrix}\frac{\partial^2}{\partial e^2}\text{output} & \frac{\partial^2}{\partial e\partial k}\text{output}\\ \frac{\partial^2}{\partial e\partial k}\text{output} & \frac{\partial^2}{\partial k^2}\text{output}\end{bmatrix}\begin{bmatrix}\mathsf{r}_1\\ \mathsf{r}_2\end{bmatrix} &= \frac{\partial^2}{\partial e^2}\text{output}\begin{bmatrix}\mathsf{r}_1 & \mathsf{r}_2\end{bmatrix}\begin{bmatrix}1 & -\frac{e}{k}\\ -\frac{e}{k} & \left(\frac{e}{k}\right)^2\end{bmatrix}\begin{bmatrix}\mathsf{r}_1\\ \mathsf{r}_2\end{bmatrix}\\ &= \frac{\partial^2}{\partial e^2}\text{output}\left(\mathsf{r}_1 - \mathsf{r}_2\frac{e}{k}\right)^2\\ &\leqslant 0.\end{aligned}$$

A.4 Representing $\widehat{F}_{x_i}$

This appendix constructs representations for partial derivative $\widehat{F}_{x_i}$ with respect to the entry, i, of the state, x, of the value function in terms of expected discounted value of marginal impulse functions. See Borovička et al. (2014) for a discussion of the constructions for Markov diffusion building in part on the work of Fournie et al. (1999). This derivation is admittedly heuristic. A more formal and fully justified treatment requires a separate contribution.

We find it convenient to work with infinitesimal generators to deduce the formulas of interest. Let $\mathbb{A}$ be the generator of the composite process, (X, Z), which we partition as

$$\mathbb{A} = \begin{bmatrix}\mathbb{A}^x\\ \mathbb{A}^z\end{bmatrix}.$$

Following Fournie et al. (1999), we form the first variational process, M^x, that gives the marginal impact of future X of a marginal change in one of the initial states. By initializing the process at one of the alternative coordinate vectors we identity the initial state of interest.[32] The generator for this process is constructed from from the generator of (X, Z)

$$\frac{\partial \mathbb{A}^x}{\partial x'}m^x,$$

where m^x is a potential realization of M^x. The operator $\mathbb{A}^x$ depends on the state vector x. The notation $\frac{\partial \mathbb{A}^x}{\partial x'}$ differentiates with respect to this state dependence and becomes a matrix

[32]Our initial condition for M_0 differs from Fournie et al. (1999) in a superficial way. They treat M as a matrix with an identity as the initialization. In this way they consider all of the states of interest simultaneously. We take M to be a vector and characterize the marginal initial responses one at a time by letting the initial condition be any one of the coordinate vectors.

operator since x' is a vector. The post-multiplication by m^x captures the dependence on the vector m^x. As might be expected, the generator is linear in m^x.

Build the generator $\mathbb{A}^{\hat{v}}$, for the constructed process $\widehat{V}(X,Z)$ using the value function computation, and let $M^{\hat{v}}$ denote the generator for the implied first variation process for $\widehat{V}$ constructed in terms of the marginal changes in the alternative initial states in x. This generator can be deduced from the formula

$$m^z = \frac{\partial \hat{V}}{\partial x'} m^x.$$

The implied generator is

$$\frac{\partial \mathbb{A}^{\hat{v}}}{\partial x'} m^x.$$

Suppose that $\rho = 1$. Differentiate the FK equation with respect to row vector x', and post multiply by m^x to obtain

$$\begin{aligned}
&\delta \left[\frac{1}{c} \frac{\partial c}{\partial x'}(x,z) - (\mathsf{u}_n)' \right] m^x \\
&+ \xi \sum_{\tilde{z} \in \mathcal{Z}} \left[1 - f^*(z \mid x,z) + f^*(z \mid x,z) \log f^*(z \mid x,z) \right] \pi(z \mid x,z) \frac{\partial \mathcal{J}}{\partial x'}(x,z) m^x \\
&- \delta m^z + \frac{\partial \mathbb{A}^{\hat{v}}}{\partial x'} m^x = 0,
\end{aligned} \tag{17}$$

where $L = 2L_n + 1$.

The solution to the equation is in the form of a resolvent operator applied to the function given in the first two lines of (17). The contribution in the first line,

$$\delta \left[\frac{1}{C_t} \frac{\partial c}{\partial x'}(X_t, Z_t) - (\mathsf{u}_n)' \right] M_t^x,$$

is the date t impulse response of marginal utility of damaged consumption to a marginal change in one of the initial state variables. The second term is special to misspecification considerations applied to processes with endogenous jump intensities. It is the marginal impulse response at date t of conditional relative entropy to a marginal change in one of the initial states X_0.

By initializing M_0^x to alternative coordinate vectors, we obtain discounted expected rep-

resentations for the partial derivatives of the value function for each element of x. We denote the marginal impulse response process to a local perturbation when M_0^x is a coordinate vector with a one in position i, as Imp^i. Expectations are computed under the worst-case distributions.

In performing this calculation, we use the property that the generators map functions that are conditionally linear in m^x into other functions with this same linearity property.

We calculate the term: $\frac{\partial c}{\partial x_i}(x)$, by first writing:

$$\frac{C_t}{K_t} = \alpha \left(1 - \phi_0(Z_t) \left(A_t^b\right)^{\phi_1}\right) - \frac{I_t^k}{K_t} - \frac{I_t^r}{K_t}$$

for

$$A_t^b = \left(1 - \frac{\mathcal{E}_t}{\beta\alpha K_t}\right) \mathbf{1}_{\{\mathcal{E}_t < \beta\alpha K_t\}}.$$

This gives a formula for $\left(\frac{c}{k}\right)^*$ in terms of the maximizing control laws and states. Its derivative with respect to states that can be deduced by direct differentiation and from the derivative of the control vector with respect to the states where $\left(\frac{c}{k}\right)^*$ is the implied maximizing control law for the consumption capital ratio. Since c^* is the product of $\left(\frac{c}{k}\right)^*$ and k, we infer the derivative of interest by a simple application of the product rule.

In these computations we use the Envelope theorem to net out the impacts of minimization. We do so to capture the perspective that minimization is external to the maximizing agent. In contrast, we purposely preserve the terms involving the derivatives with respect to the maximizing control vector to support out interpretation of the value function derivatives as forward-looking asset prices under the minimizing probability model.

A.5 Parameter values for the example economy

Parameter	Value
μ_k	0.045
κ	7
σ_k	[0, 0.01]

Table 7: Capital dynamics

Parameter	Value
ζ	0
ψ_0	0.1
ψ_1	0.5
σ_r	[0, 0.0078]

Table 8: Knowledge dynamics

Parameter	Value
α	0.12
$\bar{\beta}$	0.12
$\phi_0(0)$	0.5
ϕ_1	3

Table 9: Productivity

Parameter	Value
δ	.01
ρ	1
ξ	{0.075, 0.15}

Table 10: Preferences

The subjective discount rate is set to $\delta = 0.01$. This value is consistent with the value used by others in the literature, including Barnett et al. (2020, 2022) and Barrage and Nordhaus (2023), and it leads to emissions values in our model that are comparable to estimates of current and future annual global emissions from Figueres et al. (2018). The baseline choice of the IES is set to $\rho = 1$, which is the standard log utility case. For sensitivity analysis of our model, we examine outcomes for $\rho = 2/3$ and $\rho = 3/2$, which are similar to the values considered for sensitivity analysis by Cai and Lontzek (2019) and others in the literature.

The choices of parameters for the productivity and evolution of productive capital follow from Barnett et al. (2022), who use an undamaged version of the consumption capital model to calibrate the economic growth rate to a value of 2%, consistent with empirical values from the BEA and World Bank databases. The resulting values are set to $\alpha = 0.115, \kappa = 6.667$ and $\mu_k = 0.045$. The capital volatility is set to $\sigma_k = [0, 0.01]$, matching annual percent changes in the time series of GDP from the World Bank database.

Our choices for the emissions component of the production technology (i.e., the abatement cost parameters following the interpretation of Nordhaus and others) are as follows. We set $\phi_1 = 3$, similar to the estimated parameter values from Cai and Lontzek (2019) and Barrage and Nordhaus (2023). While the value of ϕ_0 is highly uncertain, we choose $\phi_0 = 0.5$ as a reasonable benchmark for the fraction of lost output in order to achieve zero emissions. We also consider $\phi_0 = 0.1$, consistent with Barrage and Nordhaus (2023), for a sensitivity analysis comparison. The value for the emissions intensity of output β comes from the implied emissions intensity value for 2020 from Cai and Lontzek (2019).

For the R&D investment parameters, we choose $\psi_1 = 0.5$ for computational tractability and set $\psi_0 = 0.1$ so that our model generates R&D investment values that are in line with major U.S. R&D investment programs as estimated by Stine (2008) and estimates for R&D investment from Bloom et al. (2019). For simplicity, we assume the depreciation of R&D stock is given by $\zeta = 0$. The knowledge stock volatility is set to $\sigma_r = [0, 0.0078]$, matching annual percent changes in the time series of U.S. R&D capital stock from the BLS database.

In our baseline analysis we choose $\xi \in \{0.075, 0.15\}$. As noted previously, the corresponding risk aversion parameters in a recursive utility specification of preferences implied by these values would be $\gamma \approx 7.7$ and $\gamma \approx 14.3$, similar to the values considered by Cai and Lontzek (2019) and others in the literature, and they provide probability distortions that we view as reasonable based on the outcomes reported in Figures 5–7.

The initial value of capital is set so that our initial GDP matches the 2020 World GDP value of \$85 trillion estimated by the World Bank National Accounts data. With our choice of $\alpha = 0.115$, we end up with $K_0 = 739.13$. The initial value of knowledge capital is set to $R_0 = 11.2/1120$, which converts and scales the value for current US R&D capital stock in the BLS database to a global value such that the expected arrival time of a breakthrough green technological change without additional R&D investment is the year 2100. The initial value of atmospheric temperature anomaly is set to $Y_0 = 1.1$ degrees Celsius to match recent estimates from the IPCC AR6.

A.6 HJB Equation for $\rho \neq 1$

Below we write out explicitly the HJB equation and quasi-analytical simplification for the pre-technological change, pre-damage function jump state when $\rho \neq 1$. The special case for

$\rho = 1$ is entirely similar. Starting from Equation (12),

$$0 = \frac{\delta}{1-\rho}\left(\left(\frac{C_t}{V_t}\right)^{1-\rho} - 1\right) + \frac{d\widehat{V}_t}{dt} \tag{18}$$

Equation (8), and noting that $\hat{V}(X_t, Z_t) = \log V(X_t, Z_t)$, gives rise to the HJB equation:

$$\begin{aligned}
0 \;=\; & \max_{i^k, i^r, e} \min_{h,g,f} \frac{\delta}{1-\rho}\left(\left(\frac{\alpha k - i^k - i^r - \alpha k \phi_0(z)\left(1 - \frac{e}{\beta \alpha k}\right)^{\phi_1}}{\exp(\hat{V}) n}\right)^{1-\rho} - 1\right) \\
& + \frac{\partial \hat{V}}{\partial \hat{k}}\left(-\mu_k + \frac{i^k}{k} - \frac{\kappa}{2}\left(\frac{i^k}{k}\right)^2 - \frac{|\sigma_k|^2}{2} + \sigma_k h\right) + \frac{\partial^2 \hat{V}}{\partial \hat{k}\, \partial \hat{k}'}\frac{|\sigma_k|^2}{2} \\
& + \frac{\partial \hat{V}}{\partial y}\left(\frac{1}{L_y}\sum_{\ell=1}^{L_y} q(\ell|x,z)\theta(\ell) + \varsigma h\right) e + \frac{\partial^2 \hat{V}}{\partial y\, \partial y'}\frac{|\varsigma|^2}{2} e^2 \\
& + \frac{\partial \hat{V}}{\partial \hat{n}}\left((\lambda_1 + \lambda_2 y)\left(\frac{1}{L_y}\sum_{\ell=1}^{L_y} q(\ell|x,z)\theta(\ell) + \varsigma h\right) e + \lambda_2 \frac{|\varsigma|^2}{2} e^2\right) \\
& + \frac{\partial^2 \hat{V}}{\partial \hat{n}\, \partial \hat{n}'}\frac{(\lambda_1 + \lambda_2 y)^2 |\varsigma|^2}{2} e^2 \\
& + \frac{\partial \hat{V}}{\partial \hat{r}}\left(-\zeta + \psi_0(z)(i^r)^{\psi_1}\exp(-\psi_1 \hat{r}) - \frac{|\sigma_r|^2}{2} + \sigma_r h\right) + \frac{\partial^2 \hat{V}}{\partial \hat{r}\, \partial \hat{r}'}\frac{|\sigma_r|^2}{2} \\
& + \xi \mathcal{J}_g(r)(1 - g(\tilde{z} \mid x, z) + g(\tilde{z} \mid x, z)\log g(\tilde{z} \mid x, z)) + \mathcal{J}_g(r) g(\tilde{z} \mid x, z)\left(\tilde{V}(\tilde{x}(\tilde{z}), \tilde{z}) - \hat{V}\right) \\
& + \xi \mathcal{J}_n(y) \sum_{\tilde{z} \in \mathcal{Z}} \pi(\tilde{z}|x,z)(1 - f(\tilde{z} \mid x, z) + f(\tilde{z} \mid x, z)\log f(\tilde{z} \mid x, z)) \\
& + \mathcal{J}_n(y) \sum_{\tilde{z} \in \mathcal{Z}} \pi(\tilde{z}|x,z) f(\tilde{z} \mid x, z)(\tilde{V}(\tilde{x}(\tilde{z}), \tilde{z}) - \hat{V}) + \xi \frac{h'h}{2} + \chi \frac{1}{L_y}\sum_{\ell=1}^{L_y} q(\ell|x,z)\log q(\ell|x,z),
\end{aligned}$$

where we have allowed for the various types of uncertainty outlined previously.

The solution has a quasi-analytical simplification of the form

$$\hat{V}(X_t, Z_t) = \hat{v}(X_t^1, Z_t) - \hat{N}.$$

After plugging this simplification into our HJB equation and removing common terms, we are left with the following simplified HJB to solve:

$$
\begin{aligned}
0 \;=\; & \max_{i^k, i^r, e} \min_{h,g,f} \frac{\delta}{1-\rho} \left(\left(\frac{\alpha k - i^k - i^r - \alpha k \phi_0(z) \left(1 - \frac{e}{\beta \alpha k}\right)^{\phi_1}}{\exp(\hat{v})} \right)^{1-\rho} - 1 \right) \\
& + \frac{\partial \hat{v}}{\partial \hat{k}} \left(-\mu_k + \frac{i^k}{k} - \frac{\kappa}{2} \left(\frac{i^k}{k} \right)^2 - \frac{|\sigma_k|^2}{2} + \sigma_k h \right) + \frac{\partial^2 \hat{v}}{\partial \hat{k} \, \partial \hat{k}'} \frac{|\sigma_k|^2}{2} \\
& + \frac{\partial \hat{v}}{\partial y} \left(\frac{1}{L_y} \sum_{\ell=1}^{L_y} q(\ell | x, z) \theta(\ell) + \varsigma h \right) e + \frac{\partial^2 \hat{v}}{\partial y \, \partial y'} \frac{|\varsigma|^2}{2} e^2 \\
& - \left((\lambda_1 + \lambda_2 y) \left(\frac{1}{L_y} \sum_{\ell=1}^{L_y} q(\ell | x, z) \theta(\ell) + \varsigma h \right) e + \lambda_2 \frac{|\varsigma|^2}{2} e^2 \right) \\
& + \frac{\partial \hat{v}}{\partial \hat{r}} \left(-\zeta + \psi_0(z) (i^r)^{\psi_1} \exp(-\psi_1 \hat{r}) - \frac{|\sigma_r|^2}{2} + \sigma_r h \right) + \frac{\partial^2 \hat{v}}{\partial \hat{r} \, \partial \hat{r}'} \frac{|\sigma_r|^2}{2} \\
& + \xi \mathcal{J}_g(r) (1 - g(\tilde{z} \mid x, z) + g(\tilde{z} \mid x, z) \log g(\tilde{z} \mid x, z)) + \mathcal{J}_g(r) g(\tilde{z} \mid x, z) \left(\tilde{v}(\tilde{x}(\tilde{z}), \tilde{z}) - \hat{v} \right) \\
& + \xi \mathcal{J}_n(y) \sum_{\tilde{z} \in \mathcal{Z}} \pi(\tilde{z} | x, z) (1 - f(\tilde{z} \mid x, z) + f(\tilde{z} \mid x, z) \log f(\tilde{z} \mid x, z)) \\
& + \mathcal{J}_n(y) \sum_{\tilde{z} \in \mathcal{Z}} \pi(\tilde{z} | x, z) f(\tilde{z} \mid x, z) (\tilde{v}(\tilde{x}(\tilde{z}), \tilde{z}) - \hat{v}) + \xi \frac{h' h}{2} + \chi \frac{1}{L_y} \sum_{\ell=1}^{L_y} q(\ell | x, z) \log q(\ell | x, z).
\end{aligned}
$$

This HJB equation characterizes only the pre-technological change, pre-damage function jump state, but the simplifications used carry through for each of the corresponding post-jump states needed to solve the full problem. The HJB equations characterizing those jump states are similar, with adjustments made for the realization of each of the jump processes.

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.

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.

Bansal, Ravi, Dana Kiku, and Marcelo Ochoa. 2019. Climate change risk. *Federal Reserve Bank of San Francisco Working Paper* .

Barnett, Michael. 2023. Climate change and uncertainty: An asset pricing perspective. *Management Science* .

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, and Lars Peter Hansen. 2022. Climate change uncertainty spillover in the macroeconomy. *NBER Macroeconomics Annual* 36 (1):253–320.

Barrage, Lint and William D Nordhaus. 2023. Policies, Projections, and the Social Cost of Carbon: Results from the DICE-2023 Model. Tech. Rep. w31112, NBER.

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.

Bloom, Nicholas, John Van Reenen, and Heidi Williams. 2019. A toolkit of policies to promote innovation. *Journal of Economic Perspectives* 33 (3):163–84.

Borovička, Jaroslav, Lars Peter Hansen, and Jose A. Scheinkman. 2014. Shock Elasticities and Impulse Responses. *Mathematics and Financial Economics* 8 (4).

Brock, W. and A. Xepapadeas. 2017. Climate change policy under polar amplification. *European Economic Review* 99:263–282.

Brook, Barry W., Erle C. Ellis, Michael P. Perring, Anson W. Mackay, and Linus Blomqvist. 2013. Does the terrestrial biosphere have planetary tipping points? *Trends in Ecology and Evolution* 28:396–401.

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, Kenneth L. Judd, and Thomas S. Lontzek. 2017. The Social Cost of Carbon with Climate Risk. Tech. rep., Hoover Institution, Stanford, CA.

Cappelli, Veronica, Simone Cerreia-Vioglio, Fabio MacCheroni, Massimo Marinacci, and Stefania Minardi. 2021. Sources of Uncertainty and Subjective Prices. *Journal of the European Economic Association* 19:872–912.

Cerreia-Vioglio, Simone, Lars Peter Hansen, Fabio Maccheroni, and Massimo Marinacci. 2021. Making Decisions under Model Misspecification.

Chang, Kenneth. 2022. Scientists Achieve Nuclear Fusion Breakthrough With Blast of 192 Lasers. *New York Times, December 13* .

Duffie, D and L G Epstein. 1992. Stochastic Differential Utility. *Econometrica* 60:353–394.

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.

Fournie, E, J M Lasry, J Lebuchoux, P L Lions, and N Touzi. 1999. Applications of Malliavin Calculus to Monte Carlo Methods in Finance. *Finance and Stochastics* 3:391–413.

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.

Hambel, Christoph, Holger Kraft, and Eduardo Schwartz. 2021. Optimal carbon abatement in a stochastic equilibrium model with climate change. *European Economic Review* 132:103642.

Hansen, Lars Peter and Jiamjun Miao. 2022. Asset pricing under smooth ambiguity in continuous time. *Economic Theory* 74:335–371.

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. 1995. Discounted Linear Exponential Gaussian Control. *IEEE Transactions on Automatic Control* 40:968–971.

———. 2001. Robust Control and Model Uncertainty. *The American Economic Review* 91 (2):60–66.

Hennlock, Magnus. 2009. Robust Control in Global Warming Management: An Analytical Dynamic Integrated Assessment. Tech. Rep. 9-19, RFF Discussion Paper.

Jaakkola, Niko and Frederick van der Ploeg. 2019. Non-cooperative and cooperative climate policies with anticipated breakthrough technology. *Journal of Environmental Economics and Management* 97:42–66.

Jacobson, David H. 1973. Optimal Stochastic Linear Systems with Exponential Performance Criteria and Their Relation to Deterministic Differential Games. *IEEE Transactions for Automatic Control* AC-18:1124–1131.

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.

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

Klibanoff, P, M Marinacci, and S Mukerji. 2009. Recursive Smooth Ambiguity Preferences. *Journal of Economic Theory* 144:930–976.

Kreps, David M. and Evan L. Porteus. 1978. Temporal Resolution of Uncertainty and Dynamic Choice. *Econometrica* 46 (1):185–200.

Lemoine, Derek and Christian P Traeger. 2016. Ambiguous tipping points. *Journal of Economic Behavior & Organization* 132:5–18.

Levitan, David. 2013. Quick-Change Planet: Do Global Climate Tipping Points Exist? *Scientific American* .

Li, Xin, Borghan Narajabad, and Ted Temzelides. 2016. Robust dynamic energy use and climate change. *Quantitative Economics* 7 (3):821–857.

Maenhout, P J. 2004. Robust Portfolio Rules and Asset Pricing. *Review of Financial Studies* 17:951–983.

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

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.

Rising, James, Marco Tedesco, Franziska Piontek, and David A Stainforth. 2022. The missing risks of climate change. *Nature* 610:643–651.

Rudik, Ivan. 2020. Optimal Climate Policy When Damages Are Unknown. *American Economic Journal: Economic Policy* 12 (2):340–73.

Stine, Deborah D. 2008. The Manhattan Project, the Apollo program, and federal energy technology R & D programs: A comparative analysis. Tech. rep., Congressional Research Service, the Library of Congress.

Weitzman, Martin L. 2009. On modeling and interpreting the economics of catastrophic climate change. *The Review of Economics and Statistics* 91 (1):1–19.