Uncertainty, Social Valuation, and Climate Change Policy*

Michael Barnett†, William Brock**, Lars Peter Hansen‡, and Hong Zhang‡‡

October 26, 2026

Abstract

We explore how uncertainty about climate change operates through multiple channels and shapes prudent policy dynamics. Drawing on and extending insights from asset pricing, we confront "deep uncertainties" and highlight their importance for constructing robustly optimal policies and social valuations. Using an example economy, we quantify the tradeoff between reducing fossil fuel emissions today versus waiting until more is known about future economic consequences of climate change. We also show that heightened uncertainty concerns about substantial but remote advances in clean energy technology lead to more proactive investment in research and development across a substantively interesting range of parameter settings.

Keywords— uncertainty, climate policy, research and development, social valuation

†Arizona State University, michael.d.barnett@asu.edu

**University of Wisconsin and University of Missouri, Columbia, wbrock@ssc.wisc.edu

‡University of Chicago, lhansen@uchicago.edu

‡‡Argonne National Laboratory. Email: hongzhang@anl.gov

*This paper is an outgrowth of a previously circulated document entitled, "How Should Climate Change Uncertainty Impact Social Valuation and Policy?" The authors thank Pengyu Chen, Adlai Fisher, Peter Hansen, Joanna Harris, Chun Hei Hung, Cosmin Ilut, Hagen Kim, Aleksei Oskolkov, Stavros Panageas, Diana Petrova, Eric Renault, Grace Tsiang, Noah Williams, and Judy Yue for helpful suggestions. Pengyu Chen, Bin Cheng, Zhaoyang Xu, Zhiyu Zheng, and especially Jiaying (Jessie) Liao contributed exceptional research assistance. The editor and three referees provided valuable feedback on an earlier version of this paper. We supply an online notebook with supplemental results and the code used to derive our model solutions at `https://lphansen.github.io/mitigation_book`. This research was made possible through the generous financial support from the University of Chicago Kenneth C. Griffin Economics Incubator, the Haddad Fund for Economics Research at the Becker Friedman Institute for Economics at the University of Chicago, and the Climate Systems Engineering initiative (CSEi) at the University of Chicago.

1 Introduction

While uncertainty plays a central role in asset pricing, it is often under-appreciated in policy analyses. Climate change is an example of this, although it is sometimes acknowledged as a challenge:

> 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 (Rising, Tedesco, Piontek, Stainforth, Nature, 2022).

We propose and apply a framework designed to confront such challenges formally. We show that important new dimensions of climate policy emerge as we explore the ramifications of uncertainty. Our analysis draws on insights from decision theory under (broadly conceived) uncertainty, and from asset pricing, but applied to social (rather than market) valuation. We highlight two uncertainty tradeoffs pertinent to policy design:

- do we act now or wait until we learn more?
- how much weight do we assign to the best guesses versus the potentially bad outcomes?

We confront both of these tradeoffs using example economies where a social planner has investment and production decisions to make each period. In addition to the familiar stock of capital from aggregate investment, the behavior and valuation of two other assets come into play: (i) the knowledge capital from R&D directed at fully green technologies, and (ii) the magnitude of global warming from production-based carbon emissions, which we measure using the temperature anomaly. The planner confronts the first uncertainty tradeoff in both anticipating and responding ex post to future big events related to the economic damages from climate change and to green technological innovation. The planner accounts for the second uncertainty tradeoff with preferences that are robust to model ambiguity and potential misspecification.

Our findings: We have two main substantive findings. First, while robustly optimal policies include substantial responses at the outset, these responses do become more nuanced and, in some ways, more pronounced after information is revealed about the magnitude of economic damages for more extreme amounts of global warming. As we show in Section 7, uncertainty aversion (preference for robustness) impacts not only modal responses, but also dispersion after the event is realized. This is particularly true for investments in R&D. Second, more uncertainty concerns about the future prospects of successfully

discovering a fully clean and economically viable energy source can result in a more proactive investment in R&D. This result does not, however, hold in a monotone way with the degree of uncertainty aversion. We characterize the two counteracting forces in play that lead to this non-monotonicity. One force, familiar from one-period investment theory, is that a higher aversion to uncertainty makes the investment less attractive. In contrast, from a valuation perspective, the consequences of a successful technology discovery that eliminates a significant source of uncertainty in the future are more pronounced with enhanced aversion. In other words, the relative importance of discovery to the planner depends on the presumed aversion to uncertainty. We suggest that this second force can dominate over an interesting range of robustness concerns, which provides a rather stark illustration of how increases in aversion to uncertainty can result in a more proactive approach to policy. This finding is developed in Sections 8.4 and 9. In contrast, we find that the first force dominates when the uncertainty concerns become more extreme.

Our approach: Our example economies include two potentially large outcomes. To address the "act now or wait" tradeoff in a stylized way, we include a Poisson informational or awareness event about climate damages. We initially have considerable uncertainty about the potential damage from global warming, but this Poisson event reveals information about the damage that could occur with large temperature changes. To support this interpretation, the Poisson intensity depends on the amount of actual global warming. This specification gives it an aspect of a *tipping point*, after which a hypothetical policy maker has a much clearer picture of potential damages.

In contrast to much of the applied research on climate policy, we also consider an early stage, big investment project aimed at the discovery of new fully green technologies as an important part of a menu of policy responses. For pedagogical simplicity, we treat the discovery as the outcome of a Poisson event with an endogenous intensity, one that depends on the prior history of investments. This endogeneity gives a real-options-type character to the discovery, albeit one that replaces a simple choice of when to exercise an option with an investment in making the Poisson outcome more likely. Although the Poisson specifications are very stylized, they are both tractable and revealing as a way to explore policy challenges and aims. Section 2 provides model specification details, and Section 5 illustrates stochastic model output.

Our use of quantitative models and methods is explicitly normative and prospective and differs from the familiar moment-matching approaches commonly adopted in structural analyses in corporate finance and asset pricing. With motivations similar to those articulated by Rising et al. with regard to what is sometimes called "deep uncertainty", we propose and implement methods for assessing the impact of

uncertainty when the nature of this uncertainty is broadly conceived. Our broad approach to uncertainty draws on perspectives from multiple fields, including decision and control theory with contributions from statistics, engineering, and economics. To capture formally robustness concerns about model ambiguity and potential misspecification, preferences are defined via a minimization over possible probabilities subject to penalization.[1] We refer to the outcomes of our policy-maker decision problem as being *robustly optimal*. Our exploration requires us to take a stand on uncertainty aversion in social preferences. Since it is not the role of researchers to impose a specific level of aversion on society, we instead provide ways to quantify the consequences of alternative degrees of this aversion. Section 3 describes our approach to uncertainty and how to operationalize it.

We find it valuable to use three different conceptual tools to characterize and analyze the implications for policy in the presence of uncertainty. They add to the usual toolkit for uncertainty quantification featured in many scientific disciplines. The application of the third one is original to this paper and builds on market valuation insights familiar to finance.

i) uncertainty decomposition by channels

   Our model has four channels by which uncertainty impacts prescriptions: (i) investment productivity, (ii) carbon-climate dynamics, (iii) economic damages, and (iv) technological innovation. The first channel is familiar to production-based asset pricing models, while the others are included to reflect additional challenges from climate-change concerns. To assess the importance of each channel, we activate uncertainty concerns one channel at a time and compare the model outcomes to when we activate all simultaneously. While this global decomposition is not additive, we nevertheless find it to be a revealing way to understand better our main quantitative findings. We report the outcome of this approach in Section 6.1, isolating the most salient source of uncertainty that underlies our findings.

ii) Uncertainty-adjusted probabilities that support robust decision making

   For the resulting policy problems, we construct a probability that delivers the same decision rule with just optimization. We call this constructed probability measure the *uncertainty-adjusted probability measure*. We follow the advice of robust Bayesians by computing and inspecting this probability distribution to enhance the understanding of our findings.[2] We report these distributions to elucidate

[1] For theoretical formulations and justifications, see Anderson et al. (2003), Maccheroni et al. (2006), Hansen and Miao (2018), and Cerreia-Vioglio et al. (2026).

[2] See Good (1952) for an initial articulation of this proposal, and Chamberlain (2000) for an econometrics example.

some of our main findings and to help assess the plausibility of uncertainty aversion parameters. We report on implied uncertainty-adjusted probabilities in Section 6.2.

iii) marginal social valuation

We apply novel representations of marginal valuations including for the social cost of global warming and social value of research and development to enhance our understanding of the robustly optimal policy responses. Such constructs are often referred to in policy discussions more generally (*e.g.* see Dasgupta and Mäler (2001)) and enter first-order conditions for robustly optimal policies (see Section 4). The representations we describe and apply are in a form familiar to intertemporal asset valuation but applied to marginal social cash flows. They open the door to additive decompositions of the alternative sources of the social cash flows, and provide insights into the impact of feedback mechanisms among state variables. Moreover, they embrace uncertainty through the use of the uncertainty-adjusted probabilities and support robust implementation – even in environments characterized by "deep uncertainties." This tool is developed and applied in Section 8 in explaining our substantive findings. While the use of change in probability measure is a familiar tool for derivative claims pricing, the uncertainty-adjusted probability is different from risk-neutral probability used in options theory.

Although our main conclusions are substantive, applications of the uncertainty decompositions based on marginal social valuations are original to this paper. They have more general applicability in uncertainty quantification and robust policy assessment. In particular, they highlight the value of our uncertainty decomposition and marginal social valuation toolkit.

The model economies we consider use deliberately stark simplifications along some dimensions, as is typical of many models in macroeconomics and finance, designed to reveal potentially nonlinear and durable transition mechanisms. We solve and analyze our economies with the requisite use of global solution methods, deliberately avoiding local approximations because the entire focus is on long-duration transitional dynamics not intended to be close to interesting steady states. Among other features, the uncertain transmission dynamics incorporate endogenous adjustments and allow for abrupt movements in the climate-economic dynamics, including the Poisson events described previously. Thus, global numerical solutions are essential to properly characterizing the nonlinear implications embedded in our model and in the implications analyzed using our conceptual toolkit.

Related literature: Many contributions to what are referred to as integrated assessment models have

little if anything to say about endogenous R&D or uncertainty concerns fully integrated into the policy evaluation. A recent example in this literature is Barrage and Nordhaus (2024). Cai et al. (2017) and Daniel et al. (2019) explored the ramifications of risk aversion in different settings in which a policy maker has confidence in the probabilistic specifications and without endogenous R&D investments. Robustness and social pricing considerations were introduced in an integrated assessment model by Barnett et al. (2020), but again they did not consider endogenous R&D nor did they entertain fully dynamic implications of the resulting uncertainty aversion. While our motivation for exploring R&D investment is consistent with the arguments in Acemoglu et al. (2016), Aghion et al. (2022) and Alhamdan et al. (2023), our aim is to explore the potential benefits within a formal quantitative analysis with robustness concerns. This prior literature does not fully consider both uncertainty tradeoffs we highlighted initially, and it does not analyze the two substantive findings that we feature.

From a more methodological perspective, previous research such as Ricke and Caldeira (2014), Cappelli et al. (2021), and Barnett et al. (2022) has similar ambitions for the channel-based uncertainty decompositions, but with different implementations. What is more novel is our development and use of asset-pricing-like representations and decompositions for marginal social valuation. The resulting formulas incorporate model ambiguity and model misspecification concerns through the use of uncertainty-adjusted probabilities. Hansen and Souganidis (2025) provide a more formal mathematical justification for the representations and formulas, but our paper provides the initial impetus for that work and has an initial application of the methods.

2 Example economy

We consider a simple production-based model with an AK production technology along with adjustment costs. Our production framework includes two investment opportunities. While one of the investments is the commonly modeled investment in augmenting a broadly based capital stock, the second is an investment in R&D. Climate change is incorporated via damages induced by global warming. A successful R&D discovery eliminates carbon emissions from production, limiting any additional future climate damages.

Important components of our model specifications are two types of Poisson events, a *damage information event* that reveals the economic implications for further increases in global warming and the *technology discovery*. There are L possible initial jump realizations, with $\ell = 1, \ldots, L-1$ pertaining to the damage im-

plications for further increases in global warming and L denoting the final possible realization pertaining to the technology discovery. These Poisson events are essentially "metaphors" for more complicated phenomena, but ones that could unfold relatively rapidly.[3] The modeling simplicity allows the computations and characterizations to be particularly tractable. The resulting mechanisms and outcomes we elucidate will remain applicable in many models with more complicated dynamics. Importantly, the Poisson outcomes in our setup are triggered by intensity functions that depend on endogenous state variables, temperature increases and enhancements in the stock of knowledge supporting new green technology discoveries.

We use this setup to report some quantitative illustrations. Our model specification is highly stylized, which makes calibration challenging; but we do intend our computations to be meaningful and revealing. We will make reference to some of our calibration choices as we develop the model details, and we provide a more complete presentation and motivation in Appendix E. Given the stylized nature of our model specification, a more formal and complete calibration is not feasible.[4] As a consequence, the model outcomes we report should be thought of as illustrations of potentially important mechanisms. Since there are competing forces in play, we will find the quantitative outcomes to be a revealing and essential component to our economic analysis that follows. Perhaps this can help open the door to what Hansen and Heckman (1996) call a "symbiotic relationship" with a productive intellectual exchange between model builders and empiricists which will advance an overall climate-economics research program.

2.1 Production and innovation

Prior to introducing climate change, we include three modifications to a simple AK model of production: adjustment costs, two investment opportunities, and an energy input.

2.1.1 Output

Output is split between consumption and two different types of investment with distinct intertemporal contributions to production: a conventional capital investment, I_t^k, and an investment in R&D, I_t^r:

$$C_t + I_t^k + I_t^r = \alpha K_t \left[1 - \phi_0 \left(B_t\right)^{\phi_1}\right] \tag{1}$$

[3] For instance the damage information event could instead be an event that triggers collective awareness.

[4] Even if we had fully credible empirical inputs for all model components (which, in fact, we do not), repeated global solutions to our coupled partial differential equations (PDE) system over a nontrivial parameter space, as needed for formal estimation, would magnify dramatically the computational burden needed for formal estimation.

for $\phi_1 \geqslant 2$ and $0 \leqslant \phi_0 \leqslant 1$. Superscripts r and k denote investment type. The notation B_t is defined by

$$B_t \stackrel{\text{def}}{=} \left(1 - \frac{\mathcal{E}_t}{\beta \alpha K_t}\right) \mathbf{1}_{\{\mathcal{E}_t \leqslant \beta \alpha K_t\}}, \tag{2}$$

where $\mathbf{1}$ is an indicator function that assigns one to the event in the $\{\cdot\}$ brackets.

Emissions $\mathcal{E}_t$ are a proxy for a "dirty" energy input into production. We restrict the energy input to be nonnegative ($\mathcal{E}_t \geqslant 0$). 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 (1). For instance, one could imagine a fixed proportions technology for energy and capital, where dirty energy reduction less than the required proportion of capital must be replaced by a clean alternative subject to a convex "abatement cost." In this case, B_t is a measure of emissions abatement. Alternatively, one may think of this as a production function with curvature in the energy input.[5]

Including an energy input with this functional form is mathematically similar to what is used in "DICE" models as developed by Nordhaus (2017) and used by many others.

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 output function equals $1 - \phi_0$ when $\mathcal{E}_t = 0$ and increases up to the threshold as a concave function with curvature dictated by the parameter ϕ_1. We feature the case in which $\phi_1 = 3$. We suppose initially that $\phi_0 > 0$ and that at some future point a fully green technology becomes economically viable, in which case $\phi_0^L = 0$ and dirty energy is no longer needed to produce output. We use the superscript L to denote a realization of technology advance in the future. For instance, think of a substantial advance such as nuclear fusion, although we choose to view the source of the breakthrough more generally.[6]

2.1.2 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. The new investment

[5] We elaborate more on this mathematical formulation and interpretation in Appendix C and mention that we could allow for $\mathcal{E}$ to be negative as a way to capture carbon capture.

[6] See Chang (2022) and Ball (2023) for recent discussions of the state of nuclear fusion technologies and their promise.

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.

2.1.3 R&D capital evolution

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

$$dR_t = -\zeta R_t dt + \psi_0 \left(I_t^r\right)^{\psi_1} \left(R_t\right)^{1-\psi_1} dt + R_t \sigma_r dW_t, \tag{3}$$

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 potential depreciation in the stock of knowledge pertinent for future technological progress. The term $\sigma_r dW_t$ reflects an exogenous stochastic inflow of information on the future likelihood of a technological advance.[7] We impose that $\sigma_r \cdot \sigma_k = 0$ making the stochastic increments in the two types of capital uncorrelated. The jump intensity for a new discovery $\mathcal{J}^L(R_t)$ is proportional to the knowledge stock R. We denote the constant of proportionality by χ, so that $\mathcal{J}^L(R_t) = \chi R_t$. We refer to the resulting jump as a *technology jump*.

This endogenous R&D investment specification, while related to the seminal work of Romer (1990) and Grossman and Helpman (1993), as well as the recent climate change focused contribution of Acemoglu et al. (2016), emphasizes large-scale technological innovation and the uncertainty associated with such investments. The parameters for the evolution of the knowledge stock are set so that the simplified model without a damage information event produces outcomes that are roughly in line with empirical evidence on the returns to R&D investment from Lucking et al. (2019) and Bloom et al. (2019), and for major U.S. R&D investment programs reported in Stine (2008). We provide details in Appendix E.3.

2.2 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:

[7] See Jaakkola and van der Ploeg (2019) for recent work on the policy implications of green breakthrough technology R&D.

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

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 is given by

$$dY_t = \mathcal{E}_t[\theta dt + \varsigma dW_t] \tag{4}$$

where θ is obtained from the set Θ of TCREs implied by alternative climate models. The term ςdW_t captures short-term climate sensitivity fluctuations.[8] We impose that $\sigma_k \cdot \varsigma = \sigma_r \cdot \varsigma = 0$ making the stochastic increments in the temperature anomaly and the two types of capital uncorrelated. Figure 1 gives a smoothed histogram of the θ's used in our computations. The inputs are constructed using pulse experiments applied to models of carbon and climate dynamics from the climate science literature. Appendix E.5 describes the histogram construction.[9]

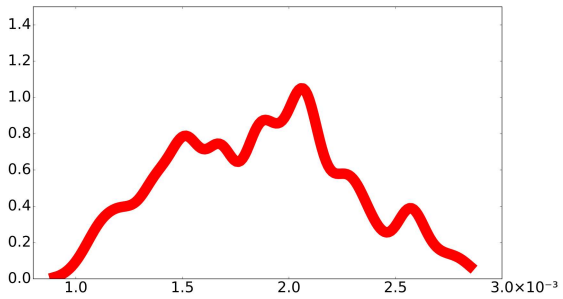


Figure 1: Smoothed density for the exponentially weighted (over horizon) responses of temperature to an emissions pulse based on input from 144 different models denoted by $\theta \in \Theta$ for $|\Theta| = 144$. The computations are based on an exponential decay rate matching the discount rate $\delta = .01$.

For baseline probabilities across the alternative climate models, we assign to each model the same subjective probability. For tractability, we abstract from parameter learning, as learning about such parameters is challenging given the underlying model complexities and has been historically slow.[10] Rather

[8] We are motivated by the TCRE literature in our specification of the mean temperature dynamics. The Brownian increment component of our temperature dynamics has permanent impacts. Such long-term contributions have not been detected in narrowly specified temperature dynamics. In more general depictions of climate change, however, durable uncertainty arising from enhanced emissions is more readily justified.

[9] 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 we use. As argued in Barnett et al. (2022), these transient dynamics have little impact on the model's implications for policy. Thus, we adopt this simpler specification to avoid including an additional state variable.

[10] One could imagine that future observations on more extreme temperatures result in learning becoming more evident. An aspect of this phenomenon is reflected in our damage information event.

than choose a particular θ, we use weighted averages of the θ's across the different climate models based on a form of robust model averaging as is described in Section 3.2.

2.3 Damage functions

We assume that consumption, C_t, is diminished proportionately by climate change damages, N_t. Damages follow a piecewise log quadratic function of the temperature anomaly y. We assume a *damage information event* not only reveals the damage curvature for more extreme temperature changes, but it also induces a discrete increase in the logarithmic derivative, $d\hat{n}/dy$, in accordance with one of $L-1$ possible specifications. Formally, the derivative of logarithmic damages $\hat{n}$ with respect to the temperature anomaly is:

$$\frac{d\hat{n}}{dy} = \lambda_1 + \lambda_2 y \qquad \text{pre damage information event} \tag{5}$$

$$\frac{d\hat{n}}{dy} = \lambda_1 + \lambda_2 (y - \hat{y} + \overline{y}) + \lambda_3(\ell)(y - \hat{y}) \qquad \text{post damage information event} \tag{6}$$

for $\ell = 1, ..., L-1$, where $\hat{y}$ denotes the temperature anomaly when the damage information event takes place. The threshold $\overline{y}$ denotes the upper limit on $\hat{y}$. The value of λ_3 is initially latent, but is subsequently revealed through the realization of a damage information event as captured by $\lambda_3(\ell)$ and $\hat{y}$.[11]

The intensity function governing this jump is given by

$$\mathcal{J}^n(y) = \begin{cases} 0 & 0 \leqslant y \leqslant \underline{y} \\ \mathsf{d}_0 \exp\left[\frac{\mathsf{d}_1}{2}(y-\underline{y})^2 + \frac{\mathsf{d}_2}{2}(y-\overline{y})^{-2}\right] - \mathsf{d}_0 \exp\left[\frac{\mathsf{d}_2}{2}(\underline{y}-\overline{y})^{-2}\right] & \underline{y} < y < \overline{y} \end{cases},$$

where $\underline{y}$ denotes the lower bound for when the damage information event can occur. The numerical values for the d's are provided in Appendix E along with a plot of the implied intensity function. The d_2 term guarantees that the damage information event occurs prior to the temperature anomaly reaching $\overline{y}$. Later we will show graphically probabilities induced by our parameter settings.

The intensity $\mathcal{J}^n$ determines probabilistically the time of a damage information event as a function of the temperature anomaly. Given that such an event takes place there are $L-1$ possible damage curves that could be realized, each with equal probability. Curve ℓ has intensity:

$$\mathcal{J}^\ell(y) = \left(\frac{1}{L-1}\right)\mathcal{J}^n(y), \qquad \ell = 1, 2, ..., L-1.$$

[11] Notice that the post jump derivatives can be expressed as functions of $y - \hat{y}$, a property that leads to a transformation of the state variable, post jump, and contributes some computational simplicity as is discussed in Appendix A.

We refer to each of these jumps as a damage curvature jump; the damage information event is the collection of these $L-1$ potential jumps. We take their intensities, the $\mathcal{J}^{\ell}$'s for $\ell = 1, 2, ...L-1$, to be baseline specifications, but we allow for robust adjustments to them.

Figure 2 displays the damage functions included in our analysis. A damage information event at a lower temperature ($\hat{y} = 1.75$) generates a range of damage curves that essentially includes those for a damage information event at a higher value ($\hat{y} = 2.25$) over the range plotted. The lower value for $\hat{y}$ also includes substantially more severe damage curves than those for the higher value. This functional form, including the values of λ_1, λ_2, and $\lambda_3(\ell)$ that we use, is roughly consistent with the range of climate damage specifications from the literature, including Nordhaus (2019), Weitzman (2012), and the recent study by Waidelich et al. (2024).

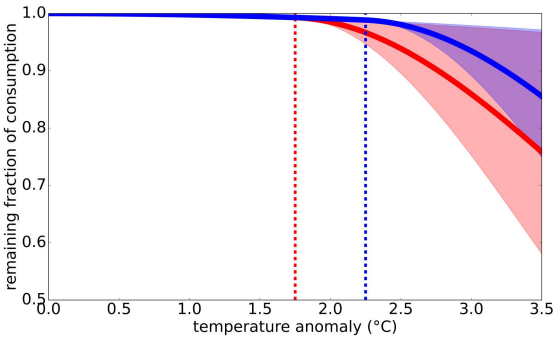


Figure 2: Range of possible damage functions for two cases with different temperature thresholds. Formally, the vertical axis measures $\exp(-\hat{n})$. The solid curves show the average values, and the shaded regions give the range of possible values for $\exp(-\hat{n})$, representing the remaining fraction of the economy's capacity to produce consumption. We constructed this figure for $\underline{y} = 1.5$ and $\overline{y} = 2.5$. The red curve and region show the potential damage functions when the damage information event occurs at $Y_t = \hat{y} = 1.75$. The blue curve and region show the potential damage functions when the event occurs at $Y_t = \hat{y} = 2.25$.

Notice that an early occurrence of the damage information event is in part "bad news" in our specification due to the more extreme concavity of the damage function, in contrast with a late occurrence. Our specification of damage uncertainty thus focuses on recognition of the severity of climate damage function curvature in the future rather than stark "falling off the cliff" type damage specifications used in some other settings with tipping points. [12] We intend this as an initial illustrative platform for addressing the layers of important uncertainties in climate damages and prudent society responses.

[12] For instance, we find this more appealing than the so-called carbon-budgeting approach often referred to in discussions of climate change policy. A carbon budget constraint imposes a Hotelling-type constraint on emissions to avoid a temperature threshold. When the constraint binds, damages immediately become very substantial in contrast to damage function specifications like ours and others engaged in climate-economic research. Moreover, setting a carbon budget in terms of cumulative emissions typically abstracts from the inherent uncertainty in how emissions impact temperature.

Remark 2.1. *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, such as those highlighted by Drijfhout et al. (2015) and Armstrong McKay et al. (2022).*

2.4 Social planner preferences

We adopt a recursive representation of preferences in continuous time for the planner. We start by forming the forward-looking continuation value, V_t, under risk for each calendar date, t:

$$\begin{aligned} V_t &= \delta \int_0^\infty \exp(-\delta u)\mathbb{E}\left[(\log C_{t+u} - \log N_{t+u}) \mid \mathfrak{F}_t\right] du \\ &= \delta \int_0^\tau \exp(-\delta u)\mathbb{E}\left[(\log C_{t+u} - \log N_{t+u}) \mid \mathfrak{F}_t\right] du + \exp(-\delta\tau)\mathbb{E}\left(V_{t+\tau} \mid \mathfrak{F}_t\right). \end{aligned}$$

where $\mathfrak{F}_t$ reflects the date t information available to the planner. Thus under risk, the preferences assume logarithmic utility in "damaged consumption", C/N. The second equality expresses a backward recursion linking future continuation values to the current one. When solving a planner's decision problem, we impose that the value function V expressed in terms of a state vector, X_t, satisfies: $V_t = V(X_t)$. The following differential equation gives the local representation:

$$0 = \delta\left(\log C_t - \log N_t - V_t\right) + \lim_{\tau \downarrow 0} \frac{1}{\tau}\left[\mathbb{E}\left(V_{t+\tau} - V_t \mid \mathfrak{F}_t\right)\right]. \tag{7}$$

This equation becomes a Hamilton-Jacobi-Bellman (HJB) equation once we substitute a candidate value function expressed as a function of the state vector. The HJB equations that interest us include maximization along with the aversion adjustments captured by minimization. As we will see, the robustness adjustments replace the local mean of the continuation value in (7) with a robust counterpart to:

$$\lim_{\tau \downarrow 0} \frac{1}{\tau}\mathbb{E}\left(V_{t+\tau} - V_t \mid \mathfrak{F}_t\right).$$

Remark 2.2. *The representation deduced in this section assumes a unitary elasticity of substitution which we feature in most of our computations, although we do explore sensitivity to this modeling choice in Appendix F and in more detail in the Online Appendix. In Appendix B we describe the extension allowing for other values of this substitution parameter.*

2.5 State and control variables

In our computations that follow, we use the state variables:

$$X_t \stackrel{\text{def}}{=} \left[\widehat{K}_t, \widehat{R}_t, Y_t\right]', \quad \text{where} \quad \widehat{K}_t \stackrel{\text{def}}{=} \log K_t \quad \widehat{R}_t \stackrel{\text{def}}{=} \log R_t$$

For the two capital stocks, our computations of state variable derivatives are expressed in terms of logarithms. When we quantify social values of the capital stocks, we make adjustments based on:

$$\frac{\partial V}{\partial \log k} = \left(\frac{\partial V}{\partial k}\right) k \quad \text{and} \quad \frac{\partial V}{\partial \log r} = \left(\frac{\partial V}{\partial r}\right) r.$$

We treat the damage curvature jumps and technology jump realizations as implying continuation values for the post-jump outcomes. These become inputs into HJB equations prior to the jump. Throughout the remainder of this 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 $\mathcal{X}$.

We have three controls:

$$\frac{I_t^k}{K_t}, \quad \frac{I_t^r}{K_t}, \quad \mathcal{E}_t.$$

The first two controls are expressed relative to capital to simplify the numerical analysis. This construction is consistent with the homogeneity properties of the technology. When we deduce first-order conditions for robust optimality, we find it more valuable pedagogically to use I_t^k and I_t^r as the choice variables without the division by capital. In the social valuation formulas, we use damaged consumption as the numeraire. Consumption prior to damages is determined by the output constraint (1), and damages are given by:

$$N_t = \exp\left[\hat{n}(Y_t)\right].$$

Consequently, consumption inclusive of damages (damaged consumption) satisfies:

$$C_t^d \stackrel{\text{def}}{=} \frac{C_t}{N_t} = \frac{K_t}{N_t}\left(\alpha\left[1 - \phi_0 \left(B_t\right)^{\phi_1}\right] - \frac{I_t^k}{K_t} - \frac{I_t^r}{K_t}\right).$$

3 Confronting uncertainty

We now analyze the contributions to the planner's Hamilton–Jacobi–Bellman (HJB) equation that emerge because of aversion to model misspecification. We pose our hypothetical social planner's decision problem

in a continuous-time environment. The uncertainty adjustments for model misspecification concerns 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 a baseline model or prior misspecifications with the most unfavorable consequences by solving a minimization problem. The uncertainty aversion of the social planner is reflected in penalties that restrain the adversarial choice. The investigation 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 omit the other terms, remembering that these are also important for deriving the maximizing control law and solving the HJB equation. The full HJB equations are provided in Appendix A and the Online Appendix.

We follow Anderson et al. (2003) by entertaining misspecification linked both to the Brownian contribution and to the jump contribution. We use a well-studied construct called relative entropy or Kullback–Leibler divergence scaled by a penalty parameter to quantify misspecification. This divergence is measured as a non-negative expected log-likelihood ratio that we use to limit the search over potential model disparities in the uncertainty analysis.[13]

3.1 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 by information that is realized between dates zero and t, including the Brownian increments.

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 $H \stackrel{\text{def}}{=} \{H_t : t \geqslant 0\}$, and hence the increment is

$$dW_t = H_t dt + dW_t^H$$

where W^H is a standard Brownian motion under the change in probability measure. Given the local normality implied by Brownian motions, the corresponding local measure of relative entropy is simply:

$$\frac{1}{2} H_t \cdot H_t,$$

[13] See Cerreia-Vioglio et al. (2026) for an axiomatic formulation of misspecification aversion.

which we use to penalize search over alternative possible misspecifications.

Write the stochastic increment of the state vector process X that is contributed by the Brownian motion as $\sigma(X_t, A_t)dW_t$, where A_t is a decision or action taken at time t. Then for a hypothetical value function V and the potential misspecification, we write the value function exposure to the Brownian increment as:[14]

$$\frac{\partial V}{\partial x'}(X_t)\sigma(X_t, A_t)dW_t = \frac{\partial V}{\partial x'}(X_t)\left[\sigma(X_t, A_t)H_t dt + \sigma(X_t, A_t)dW_t^H\right].$$

Thus, the potential misspecification contributes a drift distortion $\sigma(X_t, A_t)H_t dt$ to the dynamic evolution of the state vector process X.

We now introduce a parameter ξ_m that penalizes the search over possible misspecifications. The penalty parameter ξ_m restrains the concern for robustness to model misspecification. This leads us to solve the following minimization problem

Problem 3.1.

$$\min_{H_t} \frac{\partial V}{\partial x'}(X_t)\sigma(X_t, A_t)H_t + \frac{\xi_m}{2}H_t \cdot H_t = -\frac{1}{2\xi_m}\frac{\partial V}{\partial x'}(X_t)\sigma(X_t, A_t)\sigma(X_t, A_t)'\frac{\partial V}{\partial x}(X_t),$$

with the minimizing H given by:

$$H_t^* = -\frac{1}{\xi_m}\sigma(X_t, A_t)'\frac{\partial V}{\partial x}(X_t).$$

This minimization problem captures a form of uncertainty aversion or preference for robustness to potential misspecification, analogous to risk aversion. A limiting choice of $\xi_m \approx \infty$ implies a minimizing choice of $H_t = 0$. More generally, small values of ξ_m imply a high degree of misspecification concerns. The implied drift vector, H_t^*, has a relatively larger contribution when the value function is more adversely exposed to the Brownian increments. The parameter ξ_m governs the magnitude of H_t^*. A smaller value of ξ_m results in drift adjustments with a larger magnitude.[15] We include the minimizing outcome of Problem 3.1 in the HJB equation.

3.2 Robust model averaging

Recall that we are using as inputs 144 different climate models capturing the link between emissions and temperature, the TCRE's as given in (4). The model output histograms that we reported in Figure 1

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

[15] While this looks obvious from the Problem 3.1 solution, it is subtle since the value function implicitly depends on ξ_m.

implicitly assigned equal weights across the models. Let θ index a possible TCRE in a set Θ, and let P_t be a baseline probability that assigns weight 1/144 to each member of the set. We allow for time dependence in the baseline probability in our notation to allow for recursive learning, although we abstract from this learning in our application. Even when P_t is time invariant, we will entertain alternative weights that are time varying in our robustness analysis.

We incorporate concerns about the potential misspecification of these baseline model weights by introducing alternative relative densities $Q_t(\theta) \geqslant 0$ that satisfy:

$$\int_{\Theta} Q_t(\theta) dP_t(\theta) = 1$$

where $Q_t dP_t$ gives an alternative weighting across models. With this alternative weighting, we form the model average TCRE as follows:

$$\tilde{\theta}_t \stackrel{\text{def}}{=} \int_{\Theta} \theta Q_t(\theta) dP_t(\theta)$$

with a baseline outcome $\bar{\theta}$ obtained by setting $Q_t = 1$. The discrepancy between the two temperature evolution drifts is then given by

$$\mathcal{E}_t \left(\tilde{\theta}_t - \bar{\theta}_t \right).$$

We now use a relative entropy measure that is expressed in terms of model weights:

$$\int_{\Theta} Q_t(\theta) \log Q_t(\theta) dP_t(\theta)$$

in conjunction with a penalty parameter ξ_a. We follow Hansen and Miao (2018) by introducing[16]

Problem 3.2.

$$\begin{aligned} \min_{Q_t, \int_{\Theta} Q_t(\theta) dP_t(\theta) = 1} & \frac{\partial V}{\partial y}(X_t) \mathcal{E}_t \int_{\Theta} \theta Q_t(\theta) dP_t(\theta) + \xi_a \int_{\Theta} Q_t(\theta) \log Q_t(\theta) dP_t(\theta) \\ = & -\xi_a \log \int_{\Theta} \exp\left[-\frac{1}{\xi_a} \frac{\partial V}{\partial y}(X_t) \mathcal{E}_t \theta \right] dP_t(\theta), \end{aligned} \tag{8}$$

with a minimizing solution

$$Q_t^*(\theta) = \frac{\exp\left[-\frac{1}{\xi_a} \frac{\partial V}{\partial y}(X_t) \mathcal{E}_t \theta \right]}{\int_{\Theta} \exp\left[-\frac{1}{\xi_a} \frac{\partial V}{\partial y}(X_t) \mathcal{E}_t \theta \right] dP_t(\theta)}.$$

[16] As noted by Hansen and Miao (2018), the outcome of Minimization Problem 3.2 more generally can be viewed as a continuous-time version of a smooth ambiguity adjustment of the type advocated by Klibanoff et al. (2005).

Notice that the adjustment in the weights in the minimizing solution has an exponential tilt based on the negative of the model-specific value function evolutions:

$$\frac{\partial V}{\partial y}(X_t)\mathcal{E}_t\theta$$

scaled by $1/\xi_a$. With this adjustment, smaller continuation-value increments induce larger model weights.

The uncertainty aversion adjustment for the model averaging incorporates term (8) into the HJB equation for the policy maker in place of $\frac{\partial V}{\partial y}(x)e\theta$, or its simple average over alternative θ's for the different models. In our implementation, we will close down less structured drift distortion for the temperature equation and replace it by this more structured approach.

3.3 Jump misspecification

As we will see, jump components play prominently in our uncertainty analysis. The jump probabilities depend on endogenously determined intensities. Our specification of these intensities thus induces a corresponding endogeneity in the information structure, as we describe below.

We assume a discrete set of jump states $\ell = 1, 2, ..., L$. Let $\mathcal{J}^\ell$ denote a state-dependent intensity for a jump of type ℓ. Recall that the jump intensity, $\mathcal{J}^\ell$, implies an approximate jump probability, $\epsilon\mathcal{J}^\ell$, over a small time increment, ϵ. Following a jump of type ℓ, the value function jumps to V^ℓ. In the absence of misspecification concerns, the jump process contributes the following term to the HJB equation for $V(X)$:

$$\sum_{\ell=1}^{L} \mathcal{J}^\ell(X_t)\left[V^\ell(X_t) - V(X_t)\right], \tag{9}$$

capturing the jump-risk contribution to the stochastic evolution of the value function.

To allow for potential misspecification, we introduce non-negative functions G_t^ℓ where the altered jump contributions have intensity $\mathcal{J}^\ell(X_t)G_t^\ell(X_t)$. To restrain the exploration of potential misspecification, we introduce a convex cost:

$$\xi_m \sum_{\ell=1}^{L} \mathcal{J}^\ell(X_t)\left[1 - G_t^\ell + G_t^\ell \log G_t^\ell\right].$$

The term multiplying ξ_m is a local (in time) measure of relative entropy or Kullback–Leibler divergence applicable to jump processes.[17] To confront misspecification, we solve:

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

Problem 3.3.

$$\min_{G_t^\ell \geqslant 0} \sum_{\ell=1}^{L} \mathcal{J}^\ell(X_t) G_t^\ell \left[V^\ell(X_t) - V(X_t) \right] + \xi_m \sum_{\ell=1}^{L} \mathcal{J}^\ell(X_t) \left[1 - G_t^\ell + G_t^\ell \log G_t^\ell \right]$$
$$= \xi_m \sum_{\ell=1}^{L} \mathcal{J}^\ell(X_t) \left[1 - \exp\left(-\frac{1}{\xi_m} \left[V^\ell(X_t) - V(X_t) \right] \right) \right] \tag{10}$$

with a minimizing solution:

$$G_t^{\ell *} = \exp\left(-\frac{1}{\xi_m} \left[V^\ell(X_t) - V(X_t) \right] \right).$$

Notice that the jump intensities increase or decrease depending on whether $V^\ell(X_t) - V(X_t)$ is negative or positive. The misspecification aversion for the prospective jumps contributes the RHS term from 10 to the policy-maker HJB equation in place of the term given in 9.

Remark 3.4. *In our example economy, we compute V^ℓ's by solving value functions conditioned on the damage curvature jumps ($\ell = 1, 2, ..., L-1$) and the technology jump ($\ell = L$) using the same baseline intensities for the remaining jump possibilities.*

To implement these specifications of uncertainty aversion or robustness concerns requires that we specify the parameters ξ_m and ξ_a. Since the ξ_m and ξ_a parameter settings are a feature of planner preferences, as external researchers we see little reason to commit to specific calibrated values. Instead we show implications of multiple specifications and provide some diagnostics that help in assessing their plausibility.

4 First-order conditions for robustly optimal policies

In addition to CO_2 emissions, the social planner has two investment opportunities to consider in our setup: investment in new capital and investment in R&D. We now investigate the first-order conditions prior to the realization of either the technology discovery or the damage information event. The first-order conditions for all three controls have a central role for the partial derivative of the value function with respect to an endogenous state, the logarithm of capital, the logarithm of the knowledge stock, or temperature. In all three cases, we will display the role of the shadow value of the corresponding asset stock.

The first-order conditions for investment in new capital are

$$\frac{1}{K_t} \left[\frac{\partial V}{\partial \hat{k}}(X_t) \right] \left(1 - \kappa \frac{I_t^k}{K_t} \right) - \delta \left(\frac{1}{C_t} \right) = 0,$$

where the initial division by K_t occurs because of our choice of $\log K_t$ as a state variable. Thus, we obtain the formula for the investment-capital ratio:

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

where the term in brackets on the far right of this formula, when scaled by N_t, is the reciprocal of "Q" from the theory of investment expressed in units of damage consumption.

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

$$\frac{\partial V}{\partial \hat{r}}(X_t)\psi_0\psi_1 \left(I_t^r\right)^{\psi_1-1}\left(R_t\right)^{-\psi_1} - \left(\frac{\delta}{C_t}\right) = 0. \tag{11}$$

Thus

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

where the term on the far right in brackets, when scaled by $1/N_t$, is what we will refer to as the social value of the knowledge stock of R&D (SVRD). It is quantified by the marginal value of the knowledge stock and expressed in units of (damaged) consumption.[18]

The first-order conditions for emissions depend on the uncertainty-adjusted dynamics for temperature:

$$\left[\frac{\partial V}{\partial y}(X_t)\right]\tilde{\theta}_t^* + \mathcal{E}_t\left[\frac{\partial^2 V(X_t)}{\partial y^2}\right]|\varsigma|^2 + \left(\frac{\delta}{C_t}\right)\frac{\phi_0\phi_1}{\beta}\left(B_t\right)^{\phi_1-1}\mathbf{1}_{\{\mathcal{E}_t\leqslant\beta\alpha K_t\}} = 0$$

where

$$\tilde{\theta}_t^* = \int_\Theta \theta Q_t^*(\theta) dP_t(\theta)$$

provided there is a nonnegative solution for emissions. Otherwise, emissions are zero.[19] This equation continues to hold if we divide the terms by marginal utility of damaged consumption, $\delta N_t/C_t$. The implied social cost of carbon (SCC) is

$$\left(\frac{C_t}{\delta N_t}\right)\left(\left[-\frac{\partial V}{\partial y}(X_t)\tilde{\theta}_t^*\right] - \mathcal{E}_t\left[\frac{\partial^2 V(X_t)}{\partial y^2}\right]|\varsigma|^2\right), \tag{12}$$

and the social benefit is

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

[18] The division by R_t offsets our use of the logarithm of the knowledge stock as we noted in Section 2.5.

[19] We could capture this outcome formally by including a Karush-Kuhn-Tucker multiplier.

both expressed in terms of damage-adjusted consumption as the numeraire. Going forward, we refer to

$$-\left(\frac{C_t}{\delta N_t}\right)\left[\frac{\partial V}{\partial y}(X_t)\right]$$

as the social cost of global warming (SCGW), which is a central input into the SCC.

These marginal formulas are evaluated at the socially efficient allocation as is required for deducing the Pigouvian taxation, and they depend on the presumed ambiguity aversion. This is in contrast to many empirically-based measures of the SCC. Notice also that the SCC (12) includes an explicit volatility adjustment because emissions in our model alter the local exposure to Brownian motion risk.

We report investment as a ratio to output, where our measure of output is αK for a given capital realization K. This measure deliberately abstracts from the energy input in production. Because energy use is the source of the externality, we report the resulting emissions separately rather than incorporating an abatement adjustment into output.

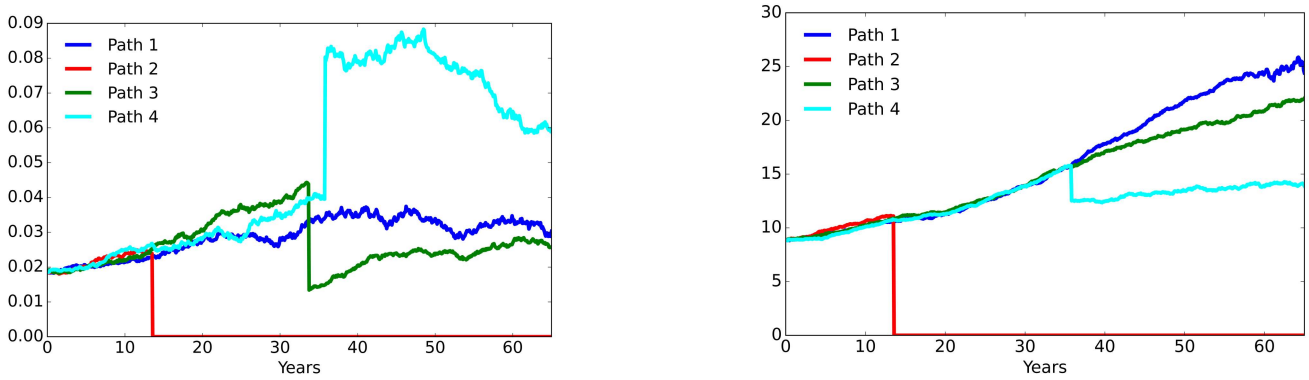
5 Stochastic simulations

Prior to providing a systematic reporting of the quantitative findings, we illustrate model output via some simulations. Figure 3 gives four stochastic pathways for the robustly optimal choice of R&D investment and emissions. The simulations are chosen to illustrate a range of possible outcomes, rather than to represent typical trajectories. The pathways were generated using the baseline probabilities with the investment and emissions decision rules computed under a choice of ξ_m that we refer to as the less averse specification.

Panel (a) illustrates the impacts of the Poisson events on R&D investment trajectories from our stochastic pathways for 65 years. Each case begins with substantial R&D investment and grows stochastically for a while. Neither a damage information event nor a technology discovery occurs along path 1. When a technology jump occurs, as in path 2, R&D investment drops to zero. Damage information events, which reveal the λ_3 governing the damage function, occur first along paths 3 and 4. On path 3, the news about damage curvature at the jump date is good (a low value of λ_3), and the path for R&D investment shows a substantial drop. Along path 4, the news about damage curvature is bad (a high value of λ_3), and there is a prominent increase in R&D investment, presumably because of the forthcoming pronounced damages.

With the exception of the path 2 technological breakthrough, the emissions responses depicted in Panel (b) of Figure 3 are more muted than the R&D investment responses. When there is good news about

damages, as in path 3, we see a negligible impact on emissions. In contrast, for path 4, when there is a bad news realization about damages, there is a notable drop in emissions.



(a) R&D investment-output ratio

(b) Emissions

Figure 3: Illustrative stochastic pathways computed with baseline probabilities and controls given by the $\xi_m = .1$, or less uncertainty aversion, specification. Panel (a) shows the robustly optimal R&D investment choices expressed as ratio of output abstracting from the energy input (αK) for the stochastic simulation pathways. Panel (b) shows the robustly optimal emissions choices for the stochastic simulation pathways. Path 1 (dark blue) is a case with no jumps. Path 2 (red) is a case with a technology jump at year 14. Path 3 (green) is a case with a damage curvature jump ($\lambda_3 = .053$) at year 34. Path 4 (light blue) is a case with a damage curvature jump ($\lambda_3 = .298$) at year 36.

We make a more systematic investigation of some of these findings in the next section.

6 Uncertainty assessments

We implement two complementary quantifications of uncertainty impacts. Since our model has multiple channels by which uncertainty can impact policy outcomes, we first consider the uncertainty consequences of the particular channels. We then formally inspect the uncertainty-adjusted probabilities.

6.1 Uncertainty deconstruction

As we delineate in the introduction, our model has four channels, each of which has its own distinct uncertainty input. We rely on our formulation of decision theory to explore the relative importance of alternative channels by which uncertainty impacts the valuation and alternative policy challenges. We compare actions and outcomes when we activate each uncertainty channel by itself when solving the adversarial minimization problem to the corresponding actions and outcomes when all channels are activated simultaneously. This will open the door to some more focused investigations of the uncertainty impacts.

We used two penalization parameters to represent the uncertainty aversion: ξ_m for misspecification concerns about the stochastic dynamics and ξ_a for ambiguity in how to weight the alternative climate models. It is fair game to explore sensitivity along both dimensions. Since the outcome of the model ambiguity gives a counterpart to a drift distortion in the temperature evolution equation, to limit the scope of the sensitivity analysis, we do the following. For a given value of ξ_m, we choose ξ_a to match the same magnitude of the drift distortion in the initial time period. In much of what follows we will illustrate results for three specifications of ξ_m with the following labels or references — **baseline**: $\xi_m = \infty$; **less aversion**: $\xi_m = .1$; and **more aversion**: $\xi_m = .05$. In addition, reported outcomes are computed at the initial time period for our simulations unless otherwise stated. For results reported across time horizons, scaling factors used to compute the outcomes are based on values from the initial simulation time period.

	SVRD		R&D investment / output	
Uncertainty channel	$\xi_m = .1$	$\xi_m = .05$	$\xi_m = .1$	$\xi_m = .05$
baseline	3.91 (60%)	3.91 (38%)	.0066	.0066
climate uncertainty	4.01 (61%)	4.10 (40%)	.0070	.0073
damage uncertainty	4.06 (62%)	4.20 (41%)	.0072	.0076
productivity uncertainty	3.83 (59%)	3.76 (36%)	.0064	.0061
technology uncertainty	5.93 (91%)	8.26 (80%)	.0153	.0296
all channels	6.52	10.36	.0185	.0465

Table 1: The SVRDs and the R&D investment-output ratios when different uncertainty channels are activated. Columns 1 and 3 show outcomes for the less uncertainty aversion case. Columns 2 and 4 show outcomes for the more uncertainty aversion case. Each row shows the outcomes when activating different uncertainty channels: no uncertainty (row 1), climate model only (row 2), damage model only (row 3), capital productivity only (row 4), R&D technology only (row 5), and all channels (row 6). The numbers in the parentheses for the SVRDs are the values in each row as a percentage of the "all channels" outcomes.

Table 1 gives the uncertainty decompositions for both the uncertainty contributions to the social values of research and development (SVRD) and the corresponding R&D investments for two values of ξ_m. We will say more about the plausibility of these values in a subsequent section. For both the social valuations and the implied actions, technological uncertainty accounts for the bulk of the uncertainty enhancement. Consider the case of a more modest uncertainty aversion. There is only a 9% reduction in the social value of R&D when activating only the technology channel. This is in contrast to an approximately 40% reduction

for the other three channels with a similar reduction under the baseline specification. While technology continues to dominate when $\xi_m = .05$, the percent reductions are larger across the board. Importantly, the technology uncertainty channel alone still captures 80% of the SVRD value.[20]

The R&D investment-output ratios reflect the dominant impact of technology uncertainty, and they *increase* in response to an enhanced aversion to misspecification uncertainty with one exception. Thus, the planner is more proactive in the R&D investment response when the uncertainty concerns are enhanced. We will have more to say about this finding in Section 9. When we activate only the uncertainty concerns about the productivity channel, R&D investment decreases slightly even relative to the baseline case, given the enhanced concerns about the overall productivity.

Table 2 presents the results of the uncertainty decomposition for the SCGW and emissions.[21] The patterns are very similar to those reported in Table 1. The technology channel is again the dominant contributor to the uncertainty responses. Emissions respond proportionately much less than R&D investment in the face of uncertainty concerns.

	SCGW		emissions	
Uncertainty channel	$\xi_m = .1$	$\xi_m = .05$	$\xi_m = .1$	$\xi_m = .05$
baseline	60.5 (51%)	60.5 (23%)	9.26	9.26
climate uncertainty	62.1 (52%)	63.5 (25%)	9.25	9.23
damage uncertainty	63.6 (54%)	66.6 (26%)	9.24	9.21
productivity uncertainty	59.9 (50%)	59.3 (23%)	9.27	9.27
technology uncertainty	103.0 (87%)	174.1 (67%)	8.96	8.58
all channels	118.9	258.9	8.85	8.15

Table 2: The SCGWs and emissions when different uncertainty channels are activated. Columns 1 and 3 show outcomes for the less uncertainty aversion case. Columns 2 and 4 show outcomes for the more uncertainty aversion case. Each row shows the outcomes when activating different uncertainty channels: no uncertainty (row 1), only climate model (row 2), only damage model (row 3), only capital productivity (row 4), only R&D technology (row 5), and all channels (row 6). The numbers in the parentheses for the SCGWs are the values in each row as a percentage of the "all channels" outcomes.

[20]This form of decomposition, sometimes referred to as a first-order index, is not additive. In contrast, a Shapley decomposition is additive and accommodates the cross-dependencies of multiple channels. See Shapley (1953). The computed Shapley value percentages for the technology channel are 85% for $\xi_m = .1$ and 81% for $\xi_m = .05$. See Appendix H for more details.

[21]Multiplying the SCGWs by $\tilde{\theta}^* = .00186$ for $\xi_m = \infty$, $\tilde{\theta}^* = .00189$ for $\xi_m = .1$, and by $\tilde{\theta}^* = .00197$ for $\xi_m = .05$, and scaling by 1000 to convert to units of dollars per ton of carbon, gives primary contributors to the SCCs. See formula (12).

6.2 Uncertainty-adjusted probabilities

We next consider the implied uncertainty-adjusted probabilities that emerge from our analysis for the alternative sources of uncertainty aversion: Brownian shocks, climate model weights, and Poisson events. These probabilities are the solutions to the minimization problem, evaluated at the solutions to the maximization problem. These should not be viewed as "best-guess" distributions; but rather they are the "worst-case" distributions subject to penalization that are a vehicle by which the planner constructs robustly optimal courses of action. Although so far we have featured results for two specifications of the penalty parameter: $\xi_m = .05$ (more aversion) and $\xi_m = .1$ (less aversion), we also later report some results for a larger range of values.[22] It is straightforward to run our solution code for other values of ξ_m.

While it is hard to interpret directly the magnitudes of ξ_m, we find it valuable to adopt an approach commonly used for robust Bayesian methods by inspecting worst-case probability specifications and isolating where the probability adjustments are most prominent.[23] Formally, this is motivated by an application of the Min-max Theorem, which constructs a single probability specification under which robustly optimal decisions are optimal. In addition to confirming plausibility, as suggested by Good (1952), this outcome reveals where the uncertainty concerns are most impactful for the robustly optimal policy. In effect, ξ_m provides a way to assess how consequential the various uncertainty inputs are to the decision problem.

ξ_m	capital	knowledge stock
.1	-.094	-.004
.05	-.178	-.011

Table 3: Drift distortions for capital stock evolution (center column) and the knowledge stock evolution (right column) for less aversion (top row) and more aversion (bottom row).

We start by reporting the drift distortions for the two capital stock processes. These are given in Table 3. The negative drift distortions for capital and the knowledge stock highlight that worst-case concerns tilt towards reduced effectiveness of investment in capital and R&D. The larger relative changes in the drift

[22] Abstracting from the model averaging, there is a mathematical equivalence between robustness and recursive utility risk adjustments for some of the results we report. The corresponding risk aversion parameter settings from recursive utility for less aversion ($\xi_m = .1$) and more aversion ($\xi_m = .05$) are $\gamma = 11$ and $\gamma = 21$, respectively. We only mention this because of the extensive use and familiarity of recursive utility in the asset pricing literature. The robustness interpretation is central to our analysis and is particularly apropos for climate economic applications. This is consistent with the perspective stated in Rising et al. (2022) and elsewhere, and it supports our uncertainty decompositions.

[23] For instance, see Good (1952) and Chamberlain (2000).

distortion when moving from less to more aversion underscore the sensitivity to the uncertainty concerns faced by the planner in setting robustly optimal investment strategies. The negatives of drift distortions for the capital stock are interpretable as implied shadow Sharpe ratios induced by uncertainty aversion to the exposure of capital to Brownian risk, since the reported distortions are multiplied by the volatility exposure to the productivity shock.[24] The magnitudes of the distortions for the knowledge stock are quantitatively small. As we will see, the jump components to the technological uncertainty are a major contributor to the social valuation of the stock of knowledge, in contrast to these reported drift distortions.

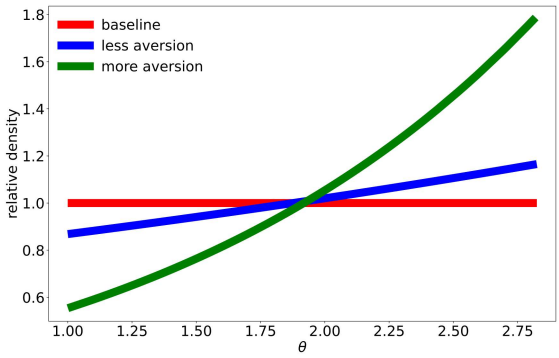


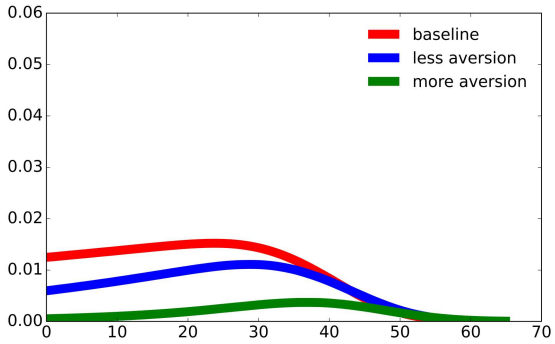
Figure 4: Uncertainty-adjusted densities for θ. The baseline and two uncertainty-adjusted densities are relative to a uniform baseline density. The resulting model averages are 1.86 for $\xi_m = \infty$, 1.89 for $\xi_m = .1$, and 1.97 for $\xi_m = .05$.

Figure 4 reports the baseline and uncertainty-adjusted relative probability density functions for two different values of ξ_m relative to baseline probabilities that equally weight each of the 144 different model outcomes. By construction, the baseline relative density is constant at a value of one. The two relative densities of interest slope upward to the right, placing more probability on the higher values of θ. This implies an increase in the mean. While the relative densities show substantive divergence from unity for values of θ near 1 and 2.5, these turn out to be very *atypical* values of θ as we displayed in Figure 1. As a result, the mean increases are quite modest, as reported in Figure 4, especially for $\xi_m = .1$ (less aversion).

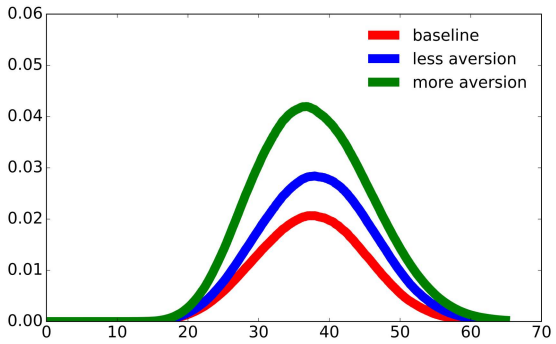
Our model allows for multiple jumps with state-dependent intensities. We compute the relative contributions for the timing of the first jump. This first jump could be either one of the damage curvature jumps or a technology jump. Using the formulas provided in Appendix D, we report additive decompositions of the jump-time densities (relative to uniform densities) for the damage information event and technology

[24] In Appendix E.7, we discuss further the (shadow) local asset pricing implications of our model. The implications of global asset pricing depend on the consumption allocations deduced from prudent policy analysis. This ambition is different from deducing positive predictions for suboptimal allocations.

jump contributions. The damage information event aggregates over all possible damage curvature jumps.



(a) Contribution of technology jump.



(b) Contribution of the damage information event.

Figure 5: Additive decomposition of the jump-time densities under the baseline probability and uncertainty-adjusted probabilities for two specifications of uncertainty aversion. By design, the density contributions across the two panels sum to one. The technology discovery contribution is 60% for the baseline specification, 42% for less aversion, and 12% for more aversion.

Since there are multiple potential damage curves, we report how misspecification aversion shifts the weights of the unknown damage curve parameter, λ_3. The baseline ($\xi_m = \infty$) distribution for λ_3 is uniform $[0, .33]$. We used twenty equally spaced discrete support points in our computations. As we see in Figure 6 the uncertainty-adjusted probabilities are tilted to the right, with a quite extreme shift when $\xi_m = .05$. The implied averages for λ_3 are given in the caption to Figure 6. The adjustment $\xi_m = .1$, while more modest, is still quite notable. The distributional shifts for the potential damage curvature turn out to be much more consequential than the climate model uncertainty.

Figure 5 plots the *ex ante* density contributions for the first jump time. The left panel of the figure displays the contribution coming from a technology jump and the right panel from the damage information event. The jump time density is the sum of these two contributions. Uncertainty aversion pertaining to the technology contribution makes a successful discovery appear more likely to be delayed, reflected in the uncertainty-adjusted density in the left panel. The delayed prospects for success make R&D investment *less attractive*. Consistent with these probabilistic delays, as we increase the uncertainty aversion, the contribution of the damage information event becomes much higher under the uncertainty-adjusted probability measures. The modal jump date for a damage information event occurs around 37 years, about the same with or without ambiguity aversion. Under a "business as usual" vantage point, the counterpart density would have a mode that occurred much earlier with a much greater contribution from damages.

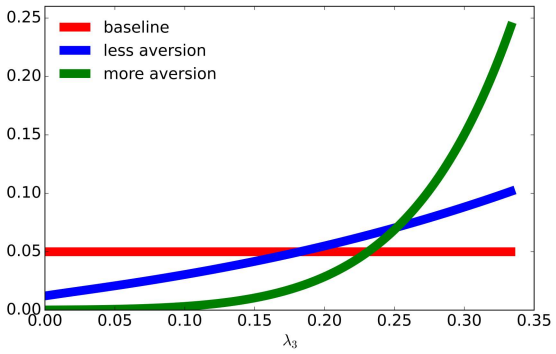


Figure 6: Uncertainty-adjusted probabilities across the alternative λ_3's for the baseline probabilities and two specifications of uncertainty aversion. The resulting λ_3 averages are .167 for the baseline specification, .221 for less aversion, and .285 for more aversion.

In results that follow, the Poisson jump uncertainty will play a central role in the policy analysis.

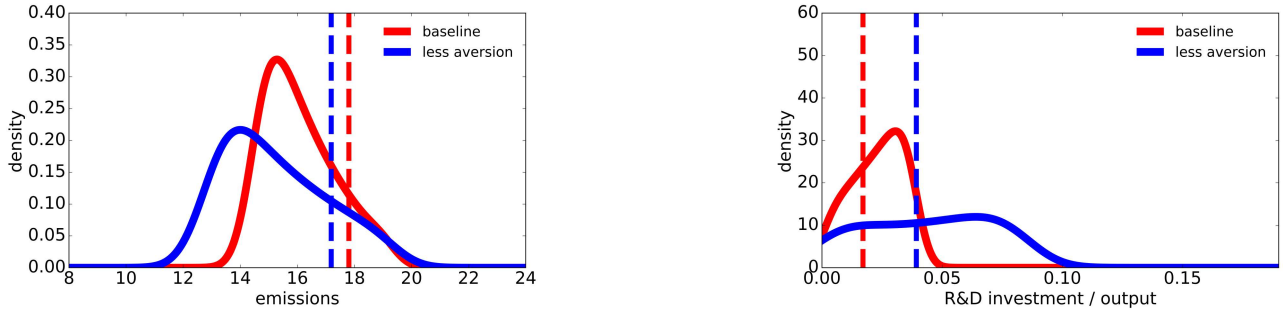
7 Acting now or waiting?

Our damage information event reveals the tail of the damage function associated with more extreme warming of the environment. As usual with conditioning, responses depend on which damage curvature is revealed. The responses to this event reflect the uncertainty tradeoff between acting now versus waiting until there is more information. The damages we confront in the future could be very severe or possibly modest, giving a motive for waiting. On the other hand, if they turn out to be severe, it may be very costly to delay all action and an initial response based on this possibility could be prudent.

Figure 7 and Table 4 give the "before and after" emissions and R&D investment responses to a damage information event. For the sake of illustration, the figure and table consider two different state configurations at the time when the event occurs. In both cases we assume the event occurs at a temperature anomaly of two degrees Celsius ($\hat{y} = 2$). We set capital stock and the knowledge stock to be the medians assuming no jump of any kind takes place in prior periods. We report results for the baseline specification and the less aversion specification ($\xi_m = .1$). While the capital stock is very similar, the knowledge stock is substantially higher for the second configuration. We report results imposing the corresponding initial conditions when depicting baseline and less aversion results.[25] The dashed lines give the pre-event emissions and R&D investments, and smoothed densities show post-event responses across potential damage

[25]In Appendix I we provide additional results where we hold fixed the state variable initialization at both the baseline and the less-aversion median values and look at the outcomes for both the baseline and less-aversion robustly optimal policy rules.

curvatures. Table 4 reports the pre-event values and the post-event median value and the minimum and maximum values across potential damage curvatures for R&D investment and emissions responses.



(a) Emissions

(b) R&D investment

Figure 7: Robustly optimal emissions and R&D investment choices before and after a damage information event. Each case uses the median capital and knowledge stock values for no-jump simulations at $\hat{y} = 2^{\circ}$C under the respective uncertainty aversion specification. The median capital and knowledge stock values for the baseline case are $k = 1505$ and $r = 35$. The median capital and knowledge stock values for the less-aversion case are $k = 1508$ and $r = 56$. The post-damage information event policies are evaluated at $z = 2.5$. Red curves depict the baseline responses and blue curves depict the less-aversion responses. The solid curves are Gaussian kernel densities of the 20 post-damage curvature jump outcomes, with equal weights and Scott's bandwidth multiplied by 0.8. The dashed vertical lines give the outcomes before the damage curvature jump. Emissions are reported in gigatons of carbon (GtC). R&D investment is reported as a ratio relative to output abstracting from the energy input (αK).

Panel (a) of Figure 7 and columns 2–3 of Table 4 show the pre- and post-event emissions responses. Not surprisingly, there is considerable heterogeneity in the post-event emissions responses as the reactions become tailored to potential damages induced by more extreme global warming. With the uncertainty concerns ($\xi_m = .1$), Figure 7 shows that the response density shifts to the left. As reported in Table 4, the median post-event responses are substantially lower than the pre-event counterparts showing that there is indeed a "wait-and-see" aspect to our findings for emissions.

The pre- and post-event R&D investment responses are shown in panel (b) of Figure 7 and columns 4–5 of Table 4. As both the table and figure show, there is a very substantial amount of heterogeneity in the post-event responses, even more so than with emissions. Not only are the median post-event responses higher than the pre-event counterparts, but Figure 7 also shows that the spread in post-event responses is substantially amplified when the planner has uncertainty concerns ($\xi_m = .1$).

	emissions (GtC)		R&D investment / output	
ξ_m	pre-event	post-event	pre-event	post-event
∞	17.8	15.8	.017	.024
		(14.6, 18.9)		(.002, .038)
.1	17.2	14.9	.039	.045
		(12.9, 19.0)		(.001, .083)

Table 4: Robustly optimal emissions and R&D investment choices before and after a damage information event. Each case uses the median capital and knowledge stock values for no-jump simulations at 2°C under the respective uncertainty aversion specification. The median capital and knowledge stock values for the baseline case are $k = 1505$ and $r = 35$. The median capital and knowledge stock values for the less-aversion case are $k = 1508$ and $r = 56$. The post-damage information event policies are evaluated at $z = 2.5$. Parentheses below each median report the minimum and maximum of the 20 equally weighted post-event, pre-technology policy values. Emissions are reported in gigatons of carbon (GtC). R&D investment is reported as a ratio relative to output abstracting from the energy input (αK).

8 Marginal valuations as asset prices

We next describe a type of local sensitivity analysis that helps us interpret some of our main findings using an asset pricing perspective. In terms of an asset pricing analogy, we take an intertemporal cash flow rather than a one-period return perspective when deducing contributions and uncertainty implications to valuations. This intertemporal valuation perspective seems particularly relevant because of the long-term implications of our model specification and Poisson events that we entertain over long horizons.

We draw on an approach for representing intertemporal marginal values developed and justified in Hansen and Souganidis (2025). Our interest is in using the resulting representations to interpret and deconstruct model implications. The representations of marginal valuation closely resemble intertemporal cash flow valuation under uncertainty. The formulas we use are designed to incorporate broadly-based uncertainty concerns and a prominent role of Poisson events with endogenous intensities in our framework, both of which are central to our application.

Distinct marginal social valuations emerge in the three sets of first-order conditions for the planner controls: the value of capital $\left(\frac{\partial V}{\partial k}\right)$, the value of R&D $\left(\frac{\partial V}{\partial r}\right)$, and the cost of global warming $\left(-\frac{\partial V}{\partial y}\right)$. While much of the climate-economics literature has focused on variants of the social cost of global warming, we are also particularly interested in the social value of R&D and how uncertainty aversion impacts this valuation in order to understand why our planner is proactive in response to enhanced aversion to uncertainty. The asset-pricing type representations of the marginal valuations, expressed in terms of discounted social cash

flows, allow us to inspect and deconstruct the contributions to these social valuations. We construct our marginal value representations prior to the first jump occurring.

As we have shown, this jump could be one of the damage curvature jumps or a technology discovery jump. The initial jump plays the role of a state-dependent terminal condition with an endogenously determined continuation value conditioned on the jump occurring. This initial jump could be any of the L possible jump types. Given this perspective, we adjust the discounting to incorporate jump probabilities. The resulting discount rate we use for time τ is:

$$\delta + \sum_{\ell=1}^{L} G_\tau^{*\ell} \mathcal{J}^\ell(X_\tau).$$

with the date t, state-dependent discount factor given by:

$$Dis_t \stackrel{\text{def}}{=} \exp\left(-\int_0^t \left[\delta + \sum_{\ell=1}^{L} G_\tau^{*\ell} \mathcal{J}^\ell(X_\tau)\right] d\tau\right).$$

Recall that the $G_\tau^{*\ell}$'s are the uncertainty adjustment factors for the intensities. The adjustment:

$$\exp\left(-\int_0^t \left[\sum_{\ell=1}^{L} G_\tau^{*\ell} \mathcal{J}^\ell(X_\tau)\right] d\tau\right). \tag{13}$$

to the subjective discounting gives the probability of there not being a jump until date t. We include this term in the discount factor Dis_t so that we can continue to exploit diffusion-based characterizations of marginal valuation.

We next describe a type of local sensitivity analysis that allows us to deconstruct the implications of the model. The key inputs in this and other related forms of sensitivity analysis are marginal impulse response processes Λ,[26] which give the impacts on the future state vector X of a marginal change in one of the initial states. When the state dynamics are nonlinear, Λ is stochastic. We refer to it as the stochastic response process, as it is a generalization of the linear impulse responses commonly used in applied macroeconomics. By construction, the process Λ has the same dimensionality as the number of components of X. It is built so that when we initialize it at a coordinate vector, it gives the responses to an initial change in the corresponding state variable. See, for instance, Fournie et al. (1999), Borovička et al. (2014) and Hansen and Souganidis (2025) for the stochastic evolution of Λ.

[26] In the applied mathematics literature, they are also called first variational processes.

We use the stochastic responses and state-dependent discounting in the marginal-valuation formula:

$$\frac{\partial V}{\partial x}(X_0) \cdot \Lambda_0 = \widetilde{\mathbb{E}}\left[\int_0^\infty Dis_t \left(\Lambda_t \cdot Scf_t\right) dt \mid X_0, \Lambda_0\right] \tag{14}$$

where the expectation $\widetilde{\mathbb{E}}$ reflects the diffusion dynamics incorporating the minimizing drift distortions, H^*, implied by robustness. Importantly, this expectation only features the diffusion dynamics and conditions on a jump that does not happen.

In this asset-pricing type formula, the stochastic flow process measures the impact of an initial change in a state variable as a product of two marginal calculations, loosely motivated by the chain rule:

$$\Lambda_t \cdot Scf_t.$$

In this calculation, Scf_t is a three-dimensional vector with each entry measuring a marginal contribution of a date t state variable to date t cash flow. The dot product with Λ_t reflects the stochastic response at date t to an initial change in one of the state variables, or more generally a linear combination of state variables depending upon the initialization, Λ_0. In light of this construction, we may divide $\Lambda_t \cdot Scf_t$ into contributions from each of the states since the dot product sums across vector entries. This decomposition shows how the state interdependencies are reflected in the marginal valuations.

There are three sources of date t contributions to the stochastic flow vector Scf_t :

i) marginal utility;

ii) marginal impact of a jump;

iii) marginal impact when you jump.

8.1 Marginal-utility flow

Flow i measures the impact of a marginal change in a state variable on future contributions to utility after adjusting for damages:

$$\text{flow i} \stackrel{\text{def}}{=} \delta\left(\frac{\partial \log c}{\partial x} - \frac{\partial \hat{n}}{\partial x}\right), \tag{15}$$

which is the date t vector of partial derivatives of the utility of damaged consumption with respect to each state variable, where δ is a convenient utility scaling.

Again using the chain rule, we write the term (15) as the following product:

$$\text{flow i} = \delta \left[\frac{\partial \log(c/n)}{\partial (c/n)} \right] \left[\frac{\partial (c/n)}{\partial x} \right] = \delta \left(\frac{n}{c} \right) \left[\frac{\partial (c/n)}{\partial x} \right].$$

where $n = \exp[\hat{n}(y)]$, and c/n is the adjusted consumption for damages. The first term on the right is the marginal utility of (damaged) consumption, which appears as the familiar construction of a stochastic discount factor process in asset pricing, while the second term captures the dependence of (damaged) consumption on the state vector. This type of adjustment commonly occurs in production-based asset pricing models when considering real (as opposed to financial) investments.

8.2 Prospective jump contributions

Prospective Poisson jumps also introduce cash-flow contributions to this valuation. These additional terms account for state-dependent responses operating through uncertainty in the jump intensity functions and the continuation value functions, as outlined below.

8.2.1 Marginal impact of a jump

A marginal change in the state vector in the current time period alters the jump intensities, contributing the following term to the marginal valuation:[27]

$$\text{flow ii} = \xi_m \sum_{\ell=1}^{L} \mathcal{J}^\ell \left(1 - \exp\left[-\frac{1}{\xi_m} \left(V^\ell - V \right) \right] \right) \left(\frac{\partial \log \mathcal{J}^\ell}{\partial x} \right). \tag{16}$$

The dependence of the intensities on endogenous states differentiates our setup from most "rare event" models in asset pricing. Specifically, the intensities $\mathcal{J}^\ell$ for $\ell = 1, 2, ..., L-1$ capture the stochastic dynamics of the prospective damage curvature jumps. These intensities depend only on the temperature, which simplifies the contribution of $\frac{\partial \log \mathcal{J}^\ell}{\partial x}$ to the marginal impact formula (16). Analogously, $\mathcal{J}^L$ depends only on the knowledge stock. This latter contribution is important for understanding the impacts of uncertainty on investment in R&D. Since the implied dynamical system has interactions over states, a marginal change

[27] In the following formula, the term $\mathcal{J}^\ell \left(\frac{\partial \log \mathcal{J}^\ell}{\partial x} \right)$ can be replaced by $\frac{\partial \mathcal{J}^\ell}{\partial x}$, which is pertinent when intensities are zero. Our damage intensities are zero for $y < \underline{y}$ as are their derivatives. The intensity derivatives, however, have a simple discontinuity at $y = \underline{y}$ with the left limits being zero and the right limits being positive. Since the temperature increment has a Brownian component, this simple discontinuity turns out to be inconsequential because temperature will not accumulate at this discontinuity over time.

in the initial period knowledge stock will impact all future states, inducing cross effects in the valuation.

8.2.2 Marginal impact when there is a jump

A marginal change in the state vector in the current time period alters the continuation values conditioned on a jump happening:

$$\text{flow iii} \stackrel{\text{def}}{=} \sum_{\ell=1}^{L} g^{*\ell} \mathcal{J}^{\ell} \frac{\partial V^{\ell}}{\partial x} \tag{17}$$

where the $g^{*\ell}$'s are the robust adjustments to the intensities given by:

$$g^{*\ell} \stackrel{\text{def}}{=} \exp\left[-\frac{1}{\xi_m}\left(V^{\ell} - V\right)\right] \text{ where } G_t^{*\ell} = g^{*\ell}(X_t).$$

Notice that this term captures future marginal contributions to valuation as we may deduce post-jump marginal valuation counterparts as reflected in the partial derivatives: $\frac{\partial V^{\ell}}{\partial x}$ for the alternative ℓ's. In particular, observe that this term depends on how the continuation values respond to the state variables in the post-jump world. Once a technology discovery happens, there is no direct marginal contribution of the flow iii term to valuation.

Notice that the flow ii and flow iii terms include a scaling by the uncertainty-adjusted intensities. This inclusion is consistent with formula (23) given in Appendix D where these intensities in conjunction with the contribution (13) to the factor, Dis_t, account for the uncertainty-adjusted jump density.

8.3 Decompositions

The integral representation (14) gives a horizon-based decomposition of marginal values. The flow terms open the door to three additional types of additive decompositions: one based on the flow type, another based on the jump type ℓ, and a third on which future values of states are impacted. Together, these provide valuation counterparts to the innovation accounting methods often used in empirical macroeconomics.

Table 5 shows the quantitative importance of the three flow terms for SVRD as well as two additional refinements. The contributions of flow ii, the marginal impact of a jump, and flow iii, the marginal impact when you jump, dominate the direct marginal utility contribution given by flow i. This reflects how important the longer-term contributions of the uncertain jump processes are to the SVRDs.

When viewing the decomposition results in Panels (a) and (b) of Table 5, it is helpful to recall that in our

ξ_m	flow i	flow ii		flow iii	
		damage	tech	damage	tech
∞	.31	-.06	2.03	1.20	.40
.1	.73	-.22	2.72	2.71	.62
.05	1.02	-.58	3.35	6.63	.30

(a) SVRD decomposition by potential jump types.

ξ_m	flow i		flow ii	flow iii	
	capital	R&D	R&D	capital	R&D
∞	.17	.14	2.03	.82	.80
.1	.39	.35	2.72	1.70	1.65
.05	.56	.47	3.35	2.39	4.57

(b) SVRD decomposition by state variable contributions.

Table 5: Decompositions of the three flow contributions to the SVRD delineated in Section 8. Panel (a) decomposes the flow contributions by jump types. For the columns labeled "damage," we aggregate the contributions of all of the possible damage coefficients, $\lambda_3(\ell)$, for $1 \leqslant \ell \leqslant L-1$. We report the technology jump contributions in columns labeled "tech." Panel (b) decomposes the flow contributions to the SVRD across state variable channels. The temperature anomaly contributions are small and are omitted from this table. The columns labeled "capital" report the contributions operating through the prospective capital stocks, while columns labeled "R&D" report the contributions operating through the prospective R&D knowledge stocks. For both panels, the table shows how these decompositions depend on ξ_m.[28]

model economy, after the damage information event, technology jump uncertainty remains an important concern. In contrast, once a technology discovery occurs, the damage information event is no longer relevant. As we see from Panel (a), the flow ii impact on valuation is primarily due to the prospective technology jump. In fact, the prospective damage curvature jumps contribute negatively to flow ii, albeit in a small way. As shown in Panel (b), the flow ii contribution reflects the marginal impact of an initial R&D stock change only on the future R&D stock and not on the future capital stock. The story is entirely different for flow iii. The flow iii contribution to valuation mostly reflects contributions from the prospective damage curvature jumps. While the marginal changes in both future capital and R&D stock contribute to flow iii, R&D increases in prominence relative to capital as ξ_m decreases. This state interaction comes from the output tradeoff when considering the two types of investments. Although flow ii increases with stronger concerns about robustness, such concerns have an even more dramatic impact on flow iii.

8.4 Why do enhanced uncertainty concerns induce a proactive R&D response?

From the first-order conditions for R&D investment in (11), we see a dependence on the social value of research and development (SVRD) as quantified by the marginal value of the knowledge stock. Consider flow ii of the SVRD as given by (16) to see how uncertainty aversion alters optimal R&D investment. The partial derivative of interest multiplies the robustly adjusted post-jump continuation value contribution:

[28]The implied SVRD totals from this table differ from those reported in Table 1 by between .6 and 3.4 percent because of our use of simulation methods to compute the flow decompositions.

$$\mathcal{J}^L \xi_m \left(1 - \exp\left[-\frac{1}{\xi_m}\left(V^L - V\right)\right]\right) \leqslant \mathcal{J}^L\left(V^L - V\right). \tag{18}$$

This inequality is a direct implication of the solution to Problem 3.3. It reflects the planner's cautious view of the potential success of a green technological innovation being realized and reflects the planner's *reduced* incentive to invest in R&D.

Uncertainty concerns, however, also impact the planner's view of a technology success should it happen. This is manifested in the robust adjustments to the continuation values, V^L and V. Although increasing the uncertainty aversion (decreasing ξ_m) decreases both V^L and V, it has a much greater impact on V as there is a substantial decrease in the exposure to uncertainty once the new clean technology is in place. This boost in the planner's payoff to the R&D investment *increases* the planner's incentive to invest.

To summarize, we have isolated two potentially counteracting contributions to the SVRD.

- The uncertainty-adjusted probabilities of a technology discovery assign more weight to distant successes relative to more recent ones as the aversion is enhanced. This carries over to the robust adjustment to the HJB equation as reflected in inequality 18.

- The endogenously determined net utility outcome from a potential successful discovery, measured by $V^L - V$, increases as aversion increases due to a reduction in the pre-jump value function V. The reduction in V dominates any reduction in V^L since, post-jump, there is less exposure to uncertainty. We explain this further in Section 9.

The first of these makes the R&D investment less attractive and the second one more attractive. Quantitatively, the second channel dominates in our calculations presented so far. Since matters are more complicated because of other contributions to the SVRD, in the next section, we investigate further the two contributions in a simplified version of our model.

Remark 8.1. *Preferences for robustness, as we have implemented in this paper, are known to imply a preference for early resolution of uncertainty, see Strzalecki (2013). This construct is formally defined to be purely preference based without entertaining investment responses to the early resolution. In an important paper, Kreps and Porteus (1978) feature the early resolution property of preferences as a fundamentally interesting construct using a purely risk-based formulation. Within a robustness setting, we view it as more of a side show, especially as it induces a less confident view of the prospects of a technological success,*

although others may weigh it differently. Importantly, the implied preference for early resolution cannot be the entire explanation. Our planner has multiple investment vehicles and, as we will illustrate, it is possible for either of the two forces we described to dominate. A proactive R&D investment response only applies for a range of the uncertainty aversion parameter ξ_m. It does not follow as a necessary outcome of our specification of preferences for robustness.

9 A simplified model with only technology jumps

As a simplified illustration of why the social planner's R&D investment response is so sensitive to uncertainty concerns, we consider a version of our model focused solely on the potential technology discovery. From Table 2, we saw that this technology channel had the most pronounced contribution in our uncertainty decomposition. In this simplified setting, we assume that there is a single known damage curve, i.e. there is no damage information event. The constructed damage curve instead assumes that the additional curvature begins at a temperature anomaly of $2^{\circ}C$ with a value of λ_3 being the baseline average across the set of possible $\lambda_3(\ell)$'s. The resulting damage function is shown in Figure 8. We also close down uncertainty aversion related to the diffusion contribution and the alternative climate models to further enhance the focus on the prospective technology discovery. We consider two variations of our previous computations. First, we entertain considerably larger specifications of uncertainty aversion (smaller values of ξ_m). Second, we allow for the technology discovery to occur only partially or in two stages instead of one. Our discussion in Section 8.4 applies more directly and transparently to the simplified model with a single jump.[29]

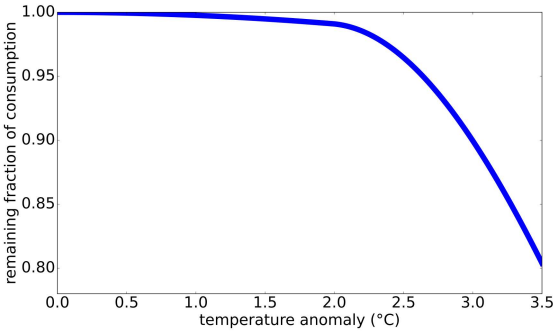


Figure 8: The consumption impact of the single damage function constructed by imposing $\hat{y} = 2$ and the arithmetic average of the $\lambda_3(\ell)$'s. Formally, the vertical axis measures $\exp(-\hat{n})$. The damages range from no impact on consumption (1.00) for a $0.0^{\circ}C$ temperature anomaly to a nearly 20% reduction in consumption ($\sim .80$) for a $3.5^{\circ}C$ temperature anomaly.

[29] Appendix J gives a more formal and explicit characterization of the two counteracting forces reflected in the SVRD's.

9.1 Enhanced uncertainty aversion

Figure 9 gives the baseline and two uncertainty-adjusted densities for the timing of the technology jump. This figure shows the delay in the prospective success under the uncertainty-adjusted probability measures. For example, the median in the baseline case is about 34.6 years. The enhanced medians for the less averse and more averse specifications are about 49.4 and 78.8 years, respectively. The uncertainty-adjusted densities make the R&D investment appear *less* attractive from a social valuation perspective.

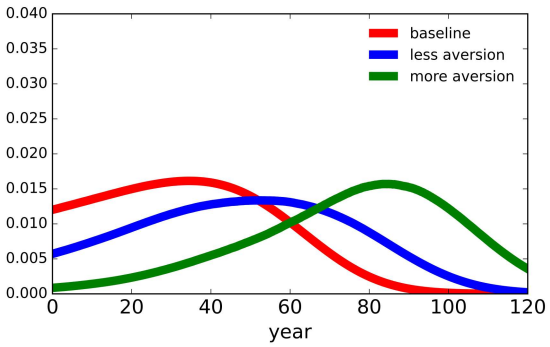


Figure 9: Jump time densities for the technology jump only model, including the baseline specification and uncertainty-adjusted densities for two alternative specifications of uncertainty aversion.

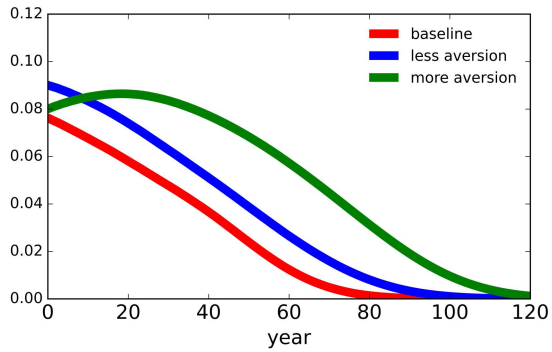
Table 6 reports the three flow contributions to the SVRD for alternative aversions to uncertainty, along with the robustly optimal initial investments for R&D and capital. In this simplified model, we continue to see more R&D investment for greater uncertainty concerns. All three contributions to the SVRD increase with uncertainty concerns, including the flow ii term. In the absence of a damage information event, the flow ii term is now the major contributor to the SVRD, accounting for about two-thirds of overall valuations. In addition, as seen by comparing Tables 6 and 5, the flow i term is substantially more prominent.

Marginal changes in both capital states contribute to the outcomes reported in Table 6 for SVRD. As we have seen, the marginal valuations can be decomposed by horizon, analogous to the measurements from impulse-response analyses. We display such horizon-based marginal valuation decompositions in Figure 10. The left panel reports flow ii decompositions, and the right panel displays flow iii decompositions. As shown in Table 6, flow ii is the major contributor to the enhanced uncertainty responses of the SVRDs.

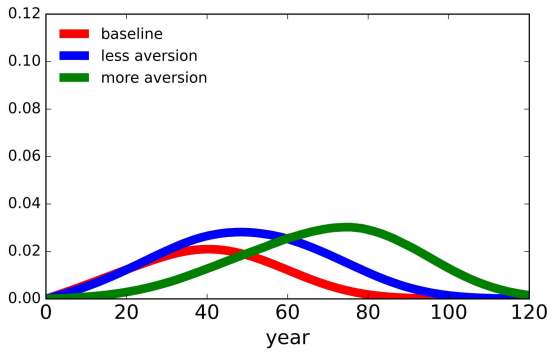
ξ_m	flow i	flow ii	flow iii	sum	R&D investment	capital investment
∞	.45	2.90	.96	4.32	.0079	.771
.1	.99	4.10	1.54	6.64	.0180	.760
.05	1.65	6.07	1.68	9.40	.0344	.742

Table 6: Flow decomposition of the SVRD and initial investments for the one-jump specification. R&D and capital investment are reported as ratios relative to output abstracting from the energy input (αK). Outcomes are shown for the three different values of the aversion parameter ξ_m in the respective rows.

From the left panel of Figure 10, we see that smaller values of ξ_m (greater uncertainty aversion) lead to greater horizon-specific contributions to the SVRDs except sometimes at very short horizons. As noted in Section 8.4, an increase in the difference $V^L - V$ gives a counteracting force to the delays in the jump time densities reported in Figure 9. Indeed, the decomposition in the left panel of Figure 10 includes the contribution of the difference in value functions after and prior to a hypothetical discovery. This latter contribution turns out to dominate the SVRD differences for the values of ξ_m presented thus far.



(a) Horizon decomposition of the flow ii contribution to the SVRDs



(b) Horizon decomposition of the flow iii contribution to the SVRDs

Figure 10: Horizon decompositions of the flow contributions for the SVRDs. To convert the marginal value contributions in Panel (a) from $\hat{r}$ to r and put them in units of the damaged consumption numeraire, we scaled the Panel (a) trajectory by the initial period ratio of damaged consumption to the stock of knowledge. We did the analogous conversion for the Panel (b) trajectory. Each panel plots trajectories under the baseline ($\xi_m = \infty$) specification as well as two alternative specifications of uncertainty aversion.

The horizon decompositions in the right panel of Figure 10 provide a rather different accounting for flow iii. The marginal continuation-value contribution, $\frac{\partial V^L}{\partial \hat{k}} = 1$, is independent of ξ_m due to the additive separability of V^L. The distorted jump intensity, however, remains an important contributor to this term.

Due to the investment tradeoff in the allocation of output, the flow iii contribution initially increases, although it does eventually decline. The locations of the peaks in the responses reflect the uncertainty-adjusted jump probabilities, which give more pessimistic assessments to the timing of an R&D success when uncertainty aversion is enhanced. The resulting flow iii term contribution, obtained by integrating over horizons, increases with the uncertainty aversion as we reported in Table 6.

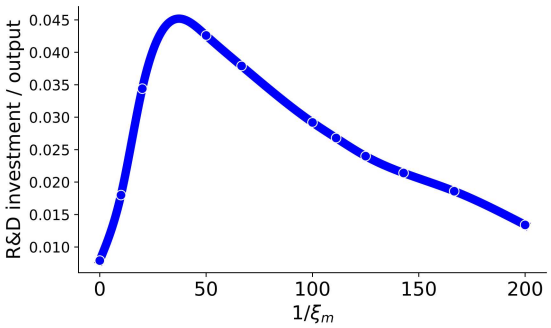


Figure 11: R&D investment as a function of the inverse of the robustness penalty parameter, $1/\xi_m$, for the technology jump only model. In constructing this figure, we first computed R&D investment-to-output ratios for a discrete set of ξ_m's, and then we formed the blue curve by taking a smoothed cubic spline fit to the discrete set of computed outcomes.

So far, in this section we have characterized more fully the two competing forces that contribute to the R&D investment outcome that we described in Section 8.4. Uncertainty in the success date makes the investment less attractive. This force is more than offset by the net payoff to discovery, which increases because an R&D success eliminates an important source of uncertainty. Although the second force dominates in our calculations, this finding turns out to be sensitive to the range of aversions that are considered. Figure 11 shows that for very high misspecification concerns (high values of $1/\xi_m$), the first force can dominate, leading instead to a reduction in the R&D investment. The SVRD follows a similar pattern, which is to be expected given the first-order conditions (11) for R&D investment.[30]

9.2 Two-step technology discovery

Using this simplified model without a damage information event, we next explore the quantitative impact of altering our one-step technology advancement specification. We first entertain a two-step counterpart with three alternative jump intensity functions.[31] In all cases, the first technology jump goes only part of

[30]The first-order conditions, (11), for $\psi_1 = 0.5$ imply that the robustly optimal R&D investment to output ratio is a (state dependent) quadratic function of the SVRD.

[31]Appendix K and Online Appendix O provide additional results and full details for the model HJB equations.

the way to a successful green transition by setting ϕ_0 in the output constraint, equation (1), to half of its original value. Only after the second technology jump is the full green transition achieved by setting ϕ_0 to zero. Continuing to denote the one-jump model jump intensity as $\mathcal{J}^L$, we explore the following scenarios:

i) we use $2\mathcal{J}^L$ as the intensity for each of the jumps;

ii) we use $\mathcal{J}^L$ as the intensity for each of the jumps;

iii) we use $2\mathcal{J}^L$ as the intensity for each jump, but we reinitialize R_t for the second jump at R_0.

In the third experiment, the initial advancement of the knowledge stock is assumed to be relevant only for the first technology jump, in contrast to the other two two-step specifications. We also include a fourth specification, which uses a one-step technology specification that allows for only a partial success:

iv) we use $\mathcal{J}^L$ as the intensity with a single jump of ϕ_0 to half of its original value;

In Table 7, we report the SVRDs and R&D investments relative to total output for each of these alternative jump specifications using ξ_m equal to .05 and .1 as in our full model. We replicate the qualitative implications of our one-jump specifications, but quantitative magnitudes differ. We view two-jump specification i) as the closest counterpart to the one-jump versions of the model, and find that the baseline R&D investment/output ratios are very similar. While uncertainty aversion continues to lead to higher R&D investments than the baseline case, the increases are more modest than in the one-jump model. Nevertheless, even under less aversion, the R&D investment/output ratio increases by about fifty percent. The decreasing intensity imposed in specification ii) results in more R&D investment across the board. When we reinitialize the R&D knowledge stock in specification iii), the R&D investment/output ratios increase with more uncertainty aversion.

Finally, consider specification iv) with a single jump to a partial technology success imposing the same initial jump intensity as in specification ii). In comparison to specification ii), the SVRDs and the implied investments are substantially reduced for specification iv). This outcome is to be anticipated because in the absence of a second jump, the rewards to R&D investment are greatly reduced. Nevertheless, the qualitative responses to increased uncertainty aversion remain for the aversions we report. As ξ_m decreases, the SVRDs and the R&D investments are amplified.

	(i)		(ii)		(iii)		(iv)	
ξ_m	SVRD	R&D inv.	SVRD	R&D inv.	SVRD	R&D inv.	SVRD	R&D inv.
∞	4.39	.0084	5.94	.0153	4.11	.0074	2.71	.0032
.1	5.40	.0127	7.40	.0238	6.17	.0166	3.29	.0047
.05	6.62	.0190	8.81	.0337	8.63	.0323	4.07	.0072

Table 7: Social valuations and initial investments for the alternative technology specifications. R&D investment is reported as a ratio relative to output abstracting from the energy input (αK). Outcomes are shown for the three different values of the aversion parameter ξ_m in the respective rows.

10 Some concluding thoughts

We investigated how concerns about uncertainty affect robustly optimal policy responses. Within our modeling framework, our calculations demonstrate that R&D investments and future informational "tipping points" in our understanding of damages can materially change prudent climate policy. Most notably, endogenous R&D emerges as an important complement to near-term carbon reductions.

To support these findings, we use an expanded toolkit for uncertainty quantification. Although uncertainty quantification is familiar to most scientific disciplines, by framing this challenge with a formal decision problem, we implement revealing refinements. Central to our approach is the construction of an uncertainty-adjusted probability measure that is conceptually distinct from, but mathematically similar to, the risk-neutral probability construction used in market asset valuation. More generally, familiar representations of market prices of intertemporal cash flows have valuable counterparts for social valuations.

The models in our paper are deliberately stark along some dimensions to illustrate how uncertainty concerns can influence prudent policies addressing climate change. Of course, a more realistic policy environment requires additional modeling complexity, including differences in (i) the vulnerability of alternative regions around the world to climate change, (ii) the incentives multiple policy makers face, and (iii) the alternative forms of new green technologies that may be worthy of investment despite ongoing doubts about their prospects. The uncertainty trade-offs we explore will remain present in such alternative models, and hence our novel uncertainty quantification methods should remain of interest. We suspect that stronger uncertainty concerns will continue to create incentives for more proactive policies.

In presenting our results, we have featured extended uncertainty quantification that includes revealing

decompositions and sensitivity to the planner's aversion to uncertainty. For practical reasons and space constraints, we limit the scope of our sensitivity analysis in the main text and provide additional results in Appendix F and the Online Appendix. We briefly summarize the key insights here. First, in a wide range of alternative parameter configurations for the R&D parameters (ψ_0, χ, σ_r), preference parameters (δ, ρ), and abatement technology parameter (ϕ_0), the qualitative implications of our analysis persist and often the differences are modest. Second, we find notable quantitative impacts for the magnitude of ρ, the inverse being the IES. As is common in production economies, changing the IES alters investments in a substantial way. Specifically, investments in both capital and R&D relative to output increase with the IES (decrease with ρ). Third, decreasing ϕ_0 to magnitudes similar to other integrated assessment models leads to substantially less R&D as a fraction of output and to a more modest but notable decrease in emissions.

The familiar real options problem and our R&D investment problem have some similarities, but also some important differences. Our decision maker can influence the timing of the R&D payoff through investment; however, there is not a specific choice of when to exercise the option. In our environment, there is an implicit investment cost due to our output constraint. In Appendix L, we explore the ways in which our findings do and do not have counterparts to insights from a real options example economy adapted from Miao and Wang (2007), in which uncertainty aversion comes into play. We contrast the more typical options formulation where the payoff to exercising is discrete with one in which the payoff is access to a flow subject to continued uncertainty.

Although our analysis does not include a formal investigation of the political economy of large-scale public investment projects, our findings illustrate that early-stage R&D investment in genuinely novel technologies can be valuable. Both public and private contributions could plausibly play a role, although our model does not speak directly to their appropriate division.[32] Determining who should make these investments, and at which stages of development, remains an important task for future research.

A commonly expressed concern regarding climate policy is the danger of focusing exclusively on one class of potential catastrophic events while omitting others, say future pandemics, nuclear war, political chaos, or others that might catch us by surprise. Although we are glad to acknowledge this point, as a research program we find value in addressing climate change without the burden of addressing all of the

[32]In the case of nuclear fusion, Whyte et al. (2026) use a simplified cost-benefit framework to explore ways to move from scientific success to economic viability. While they abstract from a formal assessment of uncertainty *within* their decision framework, they examine parameter sensitivity *outside* of this framework, which provides a form of robustness check on their computations.

other potential disasters simultaneously. Thus, we choose not to succumb to the view that "you have to do everything to do something." That said, we agree that there could be value to research that looks at tradeoffs in societal resources allocated to a menu of possible catastrophes. For such an ambition, we would find it important to embrace the potential interactions between climate change and other disasters. For example, global warming could be associated with changes in precipitation patterns, including extreme droughts (see Singh et al. (2022)), and could engender violence or more general group conflict (see Burke et al. (2024) and Hsiang (2025)). Our refinements to uncertainty quantification will remain valuable when cutting a broader swath of extreme events.

Finally, we take some comfort in the conclusions of a National Intelligence Council (2021) report stating:

> "Most countries will face difficult economic choices and probably will count on technological breakthroughs to rapidly reduce their net emissions later."

For us, this NIC perspective reinforces the emphasis in our paper on investing in longer-term technological success even if we remain unsure about its prospects.

References

Acemoglu, Daron, Ufuk Akcigit, Douglas Hanley, and William Kerr. 2016. Transition to Clean Technology. *Journal of Political Economy* 124 (1):52–104.

Aghion, Philippe, Lena Boneva, Johannes Breckenfelder, Luc Laeven, Conny Olovsson, Alexander Popov, and Elena Rancoita. 2022. Financial markets and green innovation. ECB Working Paper 2686, European Central Bank (ECB).

Alhamdan, Abdullah, Zachery Halem, Irene Hernandez, Andrew W Lo, Manish Singh, and Dennis Whyte. 2023. Financing fusion energy. *Journal of Investment Management* 21 (1).

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.

Armstrong McKay, David I, Arie Staal, Jesse F Abrams, Ricarda Winkelmann, Boris Sakschewski, Sina Loriani, Ingo Fetzer, Sarah E Cornell, Johan Rockström, and Timothy M Lenton. 2022. Exceeding 1.5 C global warming could trigger multiple climate tipping points. *Science* 377 (6611):eabn7950.

Balay, Satish, Shrirang Abhyankar, Mark F. Adams, Steven Benson, Jed Brown, Peter Brune, Kris Buschelman, Emil Constantinescu, Lisandro Dalcin, Alp Dener, Victor Eijkhout, Jacob Faibussowitsch, William D. Gropp, Václav Hapla, Tobin Isaac, Pierre Jolivet, Dmitry Karpeev, Dinesh Kaushik, Matthew G. Knepley, Fande Kong, Scott Kruger, Dave A. May, Lois Curfman McInnes, Richard Tran Mills, Lawrence Mitchell, Todd Munson, Jose E. Roman, Karl Rupp, Patrick Sanan, Jason Sarich, Barry F. Smith, Hansol Suh, Stefano Zampini, Hong Zhang, and Junchao Zhang. 2025. PETSc/TAO Users Manual. Tech. Rep. ANL-21/39 - Revision 3.23, Argonne National Laboratory.

Ball, Phillip. 2023. What is the Future of Fusion Energy? *Scientific American* .

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. 2024. Policies, Projections, and the Social Cost of Carbon: Results from the DICE-2023 Model. *Proceedings of the National Academy of Sciences* 121 (13):e2312030121.

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.

Burke, Marshall, Joel Ferguson, Solomon M. Hsiang, and Edward Miguel. 2024. New Evidence on the Economics of Climate and Conflict. Working Paper 33040, National Bureau of Economic Research.

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. 2026. Making Decisions under Model Misspecification. *Review of Economic Studies* 93 (2):822–925.

Chamberlain, Gary. 2000. Econometric applications of maxmin expected utility. *Journal of Applied Econometrics* 15 (6):625–644.

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

Daniel, Kent D., Robert B. Litterman, and Gernot Wagner. 2019. Declining CO2 Price Paths. *Proceedings of the National Academy of Sciences* 116 (42):20886–20891.

Dasgupta, Partha and Karl-Göran Mäler. 2001. Net National Product, Wealth, and Social Well-Being. *Environment and Development Economics* 5 (1):69–93.

Drijfhout, Sybren, Sebastian Bathiany, Claudie Beaulieu, Victor Brovkin, Martin Claussen, Chris Huntingford, Marten Scheffer, Giovanni Sgubin, and Didier Swingedouw. 2015. Catalogue of abrupt shifts in Intergovernmental Panel on Climate Change climate models. *Proceedings of the National Academy of Sciences* 112 (43):E5777–E5786.

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.

Good, Irving J. 1952. Rational Decisions. *Journal of the Royal Statistical Society. Series B (Methodological)* 14 (1).

Grossman, Gene M and Elhanan Helpman. 1993. *Innovation and growth in the global economy*. MIT press.

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, James E, Pushker Kharecha, Makiko Sato, George Tselioudis, Joseph Kelly, Susanne E Bauer, Reto Ruedy, Eunbi Jeong, Qinjian Jin, Eric Rignot, et al. 2025. Global Warming Has Accelerated: Are the United Nations and the Public Well-Informed? *Environment: Science and Policy for Sustainable Development* 67 (1):6–44.

Hansen, Lars Peter and James J. Heckman. 1996. The Empirical Foundations of Calibration. *Journal of Economic Perspectives* 10 (1):87–104.

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 Panagiotis E. Souganidis. 2025. Stochastic Responses and Impulse Response Functions. *Proceedings of the National Academy of Sciences* 122 (48).

Hansen, Lars Peter, Paymon Khorrami, and Fabrice Tourre. 2024. Comparative Valuation Dynamics in Production Economies: Long-Run Uncertainty, Heterogeneity, and Market Frictions. *Annual Review of Financial Economics* 16 (Volume 16, 2024):1–38.

Hsiang, Solomon. 2025. The Global Economic Impact of Climate Change: An Empirical Perspective. Working Paper 34357, National Bureau of Economic Research.

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.

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.

Klibanoff, Peter, Massimo Marinacci, and Sujoy Mukerji. 2005. A smooth model of decision making under ambiguity. *Econometrica* 73 (6):1849–1892.

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

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

Lucking, Brian, Nicholas Bloom, and John Van Reenen. 2019. Have R&D spillovers declined in the 21st century? *Fiscal Studies* 40 (4):561–590.

Maccheroni, Fabio, Massimo Marinacci, and Aldo Rustichini. 2006. Dynamic Variational Preferences. *Journal of Economic Theory* 128 (1):4–44.

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

Matthews, H Damon, Nathan P Gillett, Peter A Stott, and Kirsten Zickfeld. 2009. The proportionality of global warming to cumulative carbon emissions. *Nature* 459 (7248):829–832.

Miao, Jianjun and Neng Wang. 2007. Investment, consumption, and hedging under incomplete markets. *Journal of Financial Economics* 86 (3):608–642.

National Intelligence Council. 2021. Climate Change and International Responses Increasing Challenges to US National Security Through 2040. National Intelligence Estimate NIC-NIE-2021-10030-A, Office of the Director of National Intelligence, Washington, DC.

Nordhaus, William. 2019. Economics of the disintegration of the Greenland ice sheet. *Proceedings of the National Academy of Sciences* 116 (25):12261–12269.

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):124002.

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

Ritchie, Paul DL, Joseph J Clarke, Peter M Cox, and Chris Huntingford. 2021. Overshooting tipping point thresholds in a changing climate. *Nature* 592 (7855):517–523.

Rogelj, Joeri, Alexander Popp, Katherine V Calvin, Gunnar Luderer, Johannes Emmerling, David Gernaat,

Shinichiro Fujimori, Jessica Strefler, Tomoko Hasegawa, Giacomo Marangoni, et al. 2018. Scenarios towards limiting global mean temperature increase below 1.5 C. *Nature Climate Change* 8 (4):325–332.

Rogelj, Joeri, Daniel Huppmann, Volker Krey, Keywan Riahi, Leon Clarke, Matthew Gidden, Zebedee Nicholls, and Malte Meinshausen. 2019. A new scenario logic for the Paris Agreement long-term temperature goal. *Nature* 573 (7774):357–363.

Romer, Paul M. 1990. Endogenous technological change. *Journal of political Economy* 98 (5, Part 2):S71–S102.

Shapley, Lloyd S. 1953. A Value for n-Person Games. In *Contributions to the Theory of Games*, edited by Harold W. Kuhn and Albert W. Tucker, no. 28 in Annals of Mathematics Studies, 307–317. Princeton, NJ: Princeton University Press.

Singh, Jitendra, Moetasim Ashfaq, Christopher B. Skinner, Weston B. Anderson, Vimal Mishra, and Deepti Singh. 2022. Enhanced risk of concurrent regional droughts with increased ENSO variability and warming. *Nature Climate Change* 12 (2):163–170.

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.

Strzalecki, Tomasz. 2013. Temporal Resolution of Uncertainty and Recursive Models of Ambiguity Aversion. *Econometrica* 81 (3):1039–1074.

van der Vorst, H. A. 1992. Bi-CGSTAB: A Fast and Smoothly Converging Variant of Bi-CG for the Solution of Nonsymmetric Linear Systems. *SIAM Journal on Scientific and Statistical Computing* 13 (2):631–644.

Waidelich, Paul, Fulden Batibeniz, James Rising, Jarmo S Kikstra, and Sonia I Seneviratne. 2024. Climate damage projections beyond annual temperature. *Nature Climate Change* 14 (6):592–599.

Weitzman, Martin L. 2012. GHG Targets as Insurance Against Catastrophic Climate Damages. *Journal of Public Economic Theory* 14 (2):221–244.

Whyte, D. G., A. Lo, Rachel Bielajew, M. Hancock, R. Moeykens, and G. Shaw. 2026. Criteria for the Economic Viability of Fusion Power Plants. *Journal of Fusion Energy* .