Uncertainty, Social Valuation, and Climate Change Policy*

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

Abstract

This paper explores how uncertainty, as it pertains to climate change challenges, operates through multiple channels and has impacts for the timing of responses. We use decision theory to embrace a broad notion of uncertainty and highlight its significance for forming robustly optimal policies. These prudent policies depend on social valuations such as the social cost of climate change and the social value of research and development. Drawing insights from stochastic response theory and asset pricing, we assess when and why enhanced uncertainty concerns have important consequences for social valuation and lead to more proactive policy approaches to climate change.

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 and Pengyu Chen, Bin Cheng, Jiaying (Jessie) Liao, and Zhaoyang Xu for the exceptional research assistance. In addition, participants at the 2023 SITE Conference on Climate Finance, Innovation, and Challenges for Policy, the 2024 NBER Spring Asset Pricing Meeting, the 2024 Texas A&M Young Scholar Finance Consortium, and the 2024 SFS Cavalcade Annual Meeting provided valuable feedback. An 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://climatesocialpolicy.readthedocs.io/en/latest/index.html`.This research was made possible through the generous financial support by 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 fictitious social planner has two important decisions to make each period in addition to the familiar choice of aggregate capital investment: how much carbon to emit and how much to invest in R&D. In addition to the stock of capital, the behavior and valuation of two additional assets are of particular interest: (i) the magnitude of global warming, which we measure using the temperature anomaly, and (ii) the knowledge capital from research and development targeted at developing new green technologies.

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 event. 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. Our motivation for this is consistent with the arguments in Aghion et al. (2022) and Alhamdan et al. (2022), but our aim is to explore the potential benefits within a formal quantitative analysis. Many existing quantitative studies focus primarily on reducing emissions or replacing existing cleaner technologies. For pedagogical simplicity, we also treat this 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 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. These publications formalize sagacious responses to alternative types of uncertainty, including *risk* within a model (unknown outcomes with known probabilities), *ambiguity* across models (unknown priors weighting alternative models or parameter configurations) and potential *model misspecification* (unknown flaws of fully specified probability models). It is the latter two of these constructs that give a way to address the "deep uncertainties" that are often present in policy discussions in typically casual ways. 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 make operational.

Within this modeling environment, we have two main 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 the informational event is realized. As we show in Section 8, uncertainty aversion (preference for robustness) impacts not only modal responses but also dispersion both before and after the event is realized. This is particularly true for investments in R&D. Second, more uncertainty about

the probabilities reflected in the R&D intensity can result in a more proactive investment in R&D.[1] This result does not hold in a monotone way with the degree of uncertainty aversion (concern about robustness). For more extreme aversions, uncertainty undermines the incentive to invest. We characterize the two counteracting forces in play that lead to this non-monotonicity and suggest that the more proactive outcome is applicable over an interesting range of robustness concerns. One force, familiar from one-period investment theory, is that a higher aversion to uncertainty (more concerns about robustness) makes the investment less attractive. In contrast, the consequences of a successful technology discovery that eliminates a significant source of uncertainty have larger preference consequences post discovery when the uncertainty concerns are more pronounced. More specifically, the preference-based payoff relative to the absence of a discovery becomes more pronounced. This gives a rather stark illustration of how increases in aversion to uncertainty can result in a more proactive policy approach and is developed in Sections 9.5 and 10.

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 tool kit for uncertainty quantification featured in many scientific disciplines.

i) uncertainty decomposition by channels

Our climate 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 at the same time. 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 this approach in Section 6.

ii) Uncertainty-adjusted probabilities that support robust decision making

Consistent with the decision theory formulations in Anderson et al. (2003), Hansen and Miao (2018), and Cerreia-Vioglio et al. (2025), preferences are defined via a minimization over possible probabilities

[1] A potential for more broadly-based uncertainty aversion to induce a more proactive policy was demonstrated in a substantially different monetary policy setting by Sargent (1999) and Cogley et al. (2008).

subject to penalization.[2] For the resulting policy problems, one can 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[3] [4] We report these distributions to elucidate some of our main findings and to help assess the plausibility of uncertainty aversion parameters. We report on implied uncertainty-adjusted probabilities in Section 7.

iii) marginal social valuation

We apply novel representations of marginal valuations including for the social cost of climate change and the 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 polices (see Section 4). The representations we describe and apply are in the form familiar to intertemporal asset valuation but applied to social cash flows. They open the door to additive decompositions of the alternative sources of the social cash flows, and they 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."[5] This tool is developed and applied in Section 9.

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

[2]See also Jacobson (1973) for an initial control theory entry into this literature, and Maccheroni et al. (2006) for an axiomatic representation for a general class of preferences that incorporates minimization and robustness concerns.

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

[4]The worst-case probability construction requires and application of the Min-Max Theorem, since the minimization solution implicitly in the preferences depends on the particular decision rule being considered.

[5]See Hansen and Souganidis (2025) for formal derivations and justifications of the formulas we use.

the implications analyzed using our conceptual toolkit.

Although our main conclusions are substantive, applications of the uncertainty decompositions and representations of marginal valuations are original to this paper. They have more generally applicability in uncertainty quantification and robust policy assessment. Our example economies are designed to support the investigation of climate change policy under uncertainty aversion. In particular, they highlight the value of our uncertainty decomposition and marginal social valuation toolkit.

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.

As mentioned before, an important component to our model specifications are two types of Poisson events, an informational or awareness event of the implications for further increases in global warming and the technology discovery. These Poisson events are essentially "metaphors" for more complicated phenomena, but ones that could unfold relatively rapidly. 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 D. Given the stylized nature of our model specification, a more formal and complete calibration is not feasible.[6] As a consequence, the model outcomes we report should be thought of as illustrations of potentially important mechanisms. Since there are com-

[6]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.

peting 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 exchanges 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 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 (the superscripts denote the investment type):

$$\begin{aligned} C_t + I_t^k + I_t^r + \alpha K_t \phi_0 \left(B_t\right)^{\phi_1} &= \alpha K_t \\ C_t + I_t^k + I_t^r &= \alpha K_t \left(1 - \phi_0 \left(B_t\right)^{\phi_1}\right) \end{aligned} \tag{1}$$

for $\phi_1 \geqslant 2$ and $0 < \phi_0 \leqslant 1$. The notation B_t is defined by

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

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. 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.[7] 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.

[7] We elaborate more on this mathematical formulation and interpretation in Appendix C.

This technology is, by design, homogeneous of degree one. For a fixed K_t, the implied production function is flat when emissions exceed the threshold of $\beta\alpha K_t$ and has a zero left derivative at this point. The function equals $1-\phi_0$ when $\mathcal{E}_t = 0$ and increases up the threshold as a concave function with curvature dictated by the parameter ϕ_1. We feature the case in which $\phi_1 = 3$. 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.[8]

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 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.[9] The jump intensity for a new discovery $\mathcal{J}^L(R_t)$ is proportional to

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

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

the knowledge stock R. For notational convenience, we normalize the proportionality coefficient to be one, so that $\mathcal{J}^L(R_t) = R_t$. 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 no-damage-jump model without uncertainty aversion 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 D.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:

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

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 a TCRE obtained from the set Θ of TCRE's implied by alternative climate models. The term ςdW_t captures short timescale fluctuations. Figure 1 gives a smoothed histogram of the θ's that we use in our computations. The inputs are constructed using pulse experiments applied to models of carbon and climate dynamics from the climate science literature. Appendix D.5 describes the histogram construction.[10]

[10] 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 here this simpler specification to avoid including additional state variable.

For baseline probabilities across the alternative climate models, we presume that each model has the same subjective probability. For tractability, we abstract from parameter learning since learning about such parameters has been slow.[11]

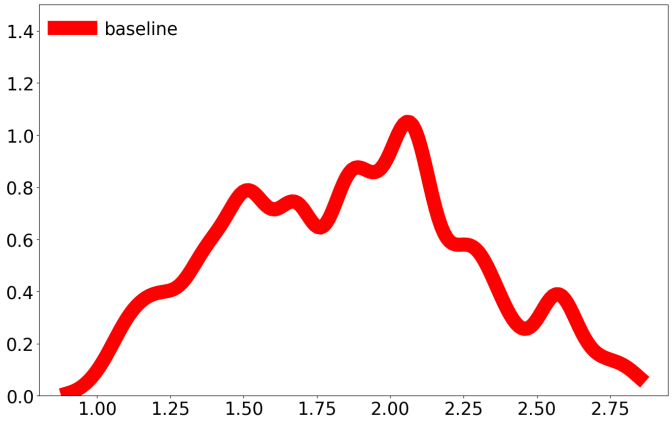

Figure 1: Smoothed density for the exponentially weighted (over horizon) responses of temperature to an emissions pulse based on input from 144 different models denoted by $\theta(\ell)$ for $\ell = 1, 2, ..., 144 = L-1,$. The computations are based on an exponential decay rate matching the discount rate $\delta = .01$.

Rather than choose a particular θ, we will use a weighted averages of the θ's across the different climate model using 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. Our damage specification uses a piecewise log quadratic specification as a function of the temperature anomaly y. We suppose that the derivative of logarithmic damages $\hat{n}$ with respect to the temperature anomaly is

$$\begin{aligned} \frac{d\hat{n}}{dy} &= \lambda_1 + \lambda_2 y & y \leqslant \hat{y} \\ \frac{d\hat{n}}{dy} &= \lambda_1 + \lambda_2 (y - \hat{y} + \overline{y}) + \lambda_3(\ell)(y - \hat{y}) & y > \hat{y} \end{aligned} \tag{5}$$

for $\ell = 1, ..., L-1$, where $\hat{y}$ denotes the temperature at which the damage jump takes place and $\overline{y}$ denotes the upper limit for when the damage jump can take place. Equation (5) has an initial condition $\hat{n}(0) = 0$. Since $\hat{y} \leqslant \overline{y}$, the derivative is positive. The implied damage function is the exponential of

[11] One could imagine that in the future observations on more extreme temperatures result in learning becoming more evident. An aspect of this phenomenon is reflected in our informational-damage jump.

$$\hat{n}(y) = \begin{cases} \lambda_1 y + \frac{\lambda_2}{2} y^2 & 0 \leqslant y \leqslant \hat{y} \\ \lambda_1 y + \frac{\lambda_2}{2}\hat{y}^2 + \frac{\lambda_2}{2}(y - \hat{y} + \overline{y})^2 + \frac{\lambda_3(\ell)}{2}(y - \hat{y})^2 - \frac{\lambda_2}{2}(\overline{y})^2 & y \geqslant \hat{y} \end{cases} \tag{6}$$

Notice from specification (6) that the logarithm of damages, $\hat{n}$, depends on the temperature anomaly, $\hat{y}$, at the time of jump. Potential damages are more extreme when the jump occurs at lower anomalies. Initially, the value of λ_3 is latent, but it is sequently revealed as a *jump revelation of the latent damage severity* as captured by $\lambda_3(\ell)$. 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 jump can occur. The numerical values for the d's are provided in Appendix D along with a plot of the implied intensity function. The d_2 term guarantees that the jump occurs prior to the temperature anomaly reaching $\overline{y}$.[12] Later we will show graphically probabilities induced by our parameter settings. In what follows we will abbreviate the name of this jump type by referring to it as a *damage severity jump*, while remembering what is actually revealed is $\lambda(\ell)$ where $1 \leqslant \ell \leqslant L - 1$.

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

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

We take these probabilities as baseline specifications, but we allow for robust adjustments to them.

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

[12] In effect, it imposes a form of "value-matching" at $y = \overline{y}$.

by Waidelich et al. (2024).

We view this jump an informational or awareness event that reveals the potential economic damages for more severe potential temperature movements. Early damage recognition is "bad news" in our specification

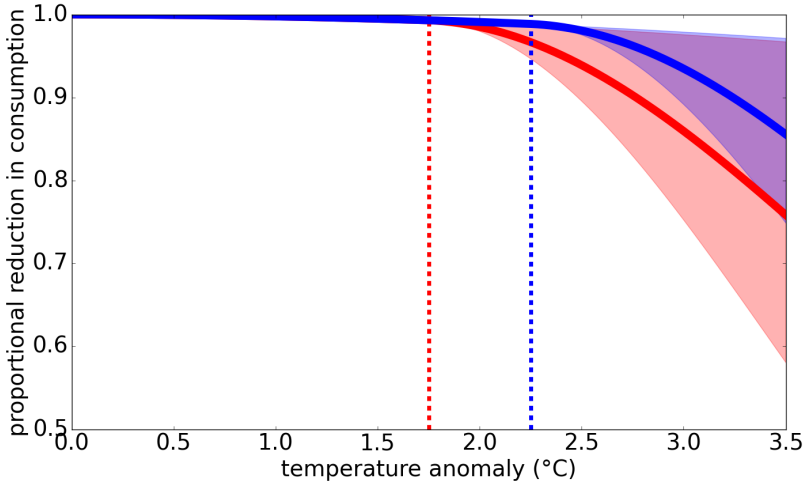

Figure 2: Range of possible damage functions for two cases with different jump thresholds. The solid curves show the average values, and the shaded regions give the range of possible values for $\exp(-n)$, representing the proportional reduction in the economy's productive capacity. We constructed this figure for $\underline{y} = 1.5$ and $\overline{y} = 2.5$. The red curve and region show the damage functions when the jump occurs at $Y_t = \hat{y} = 1.75$. The blue curve and region show the damage functions when the jump occurs at $Y_t = \hat{y} = 2.25$.

with more extreme concavity of the damage function, in contrast with late damage recognition.[13] 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. We intend this as an initial illustrative platform for addressing the layers of important uncertainties in climate damages and prudent society responses.

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).*

[13] We find this more appealing than the so-called carbon-budgeting approach often referred to discussions of climate change policy. A carbon budget constraint imposes a Hotelling-type constraint on emissions to avoid a temperature threshold. Setting a carbon budget in terms of cumulative emissions typically abstracts from the inherent uncertainty in how emissions impact temperature. When it's taken to be a hard constraint, the implied damages when the constraint binds immediately become very substantial in contrast to damage function specifications like ours and others engaged in climate-economic research.

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:

$$\begin{aligned} V_t &= \delta \int_0^\infty \exp(-\delta u)\mathbb{E}\left[(\log C_{t+u} - \log N_{t+u}) \mid \mathfrak{F}_0\right] du \\ &= \delta \int_0^\tau \exp(-\delta u)\mathbb{E}\left[(\log C_{t+u} - \log N_{t+u}) \mid \mathfrak{F}_0\right] du + \exp(-\delta\tau)\mathbb{E}\left(V_{t+\tau} \mid \mathfrak{F}_t\right). \end{aligned}$$

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):

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

with a robust counterpart.

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 E 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:

$$X_t \stackrel{\text{def}}{=} \begin{bmatrix} \widehat{K}_t \\ \widehat{R}_t \\ Y_t \end{bmatrix},$$

where

$$\widehat{K}_t \stackrel{\text{def}}{=} \log K_t \quad \widehat{R}_t \stackrel{\text{def}}{=} \log R_t$$

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

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

Consumption net of damages is then 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:

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

Remark 2.3. *An alternative and similar model that is straightforward to explore has output diminished by damages. In this case,*

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

where C_t^d is consumption inclusive of damages. While we report results with proportional damages to consumption, we have replicated some of our most important results for the computational more challenging

case of proportional damages to output.

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 analysis 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.[14]

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,

[14]See Cerreia-Vioglio et al. (2025) for an axiomatic formulation of misspecification aversion.

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:

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

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:[15]

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

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 that penalizes the search over possible misspecifications. The penalty parameter ξ 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:

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

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

[15] 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'.

choice of $H_t = 0$. More generally, small values of ξ 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 results in drift adjustments with a larger magnitude.[16]

In light of this computation, we include an additional term to the HJB equation:

$$-\frac{1}{2\xi_m}\frac{\partial V}{\partial x'}(x)\sigma(x,a)\sigma(x,a)'\frac{\partial V}{\partial x'}(x)$$

for a hypothetical state, x, and action, a.

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 implicitly assigned equal weights across of 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 member of the set. We allow for time dependence in the baseline probability in our notation to allow for recursive learning, although we will 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, form the model average TCRE:

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

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

[16] While this looks obvious from the solution to Problem 3.1, it is a bit more subtle because the value function implicitly depends on ξ.

evolution drifts

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

is the parametric counterpart to less structured drift distortion, ςH_t under more generic misspecification concerns for the temperature evolution equation (4).

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[17]

Problem 3.2.

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

with a minimizing solution

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

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.

[17] 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). The time dependence in the baseline probability allows for recursive learning, although we will abstract from this learning in our application.

Remark 3.3. *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 jumps depend on endogenously determined intensities that govern the probabilities of the jump realizations. Our specification of these intensities thus induces a corresponding endogeneity in the information structure.

We suppose there is 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.[18] To confront misspecification, we solve:

Problem 3.4.

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

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

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 3.5. *In our example economy, we compute V^ℓ's by solving value functions conditioned on the damage 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 technology or damage jump. 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 see the role of the shadow value of the corresponding asset stock.

The first-order conditions for investment in new capital are

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

Thus, we obtain the formula for investment:

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

where the term

$$\frac{C_t\left[\frac{\partial V}{\partial \hat{k}}(X_t)\right]}{\delta\, K_t N_t} \tag{11}$$

is the "Q" from the theory of investment, inclusive of damages. The division by K_t occurs because of our choice of $\log K_t$ as a state variable.

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

$$\frac{\partial V}{\partial \hat{r}}(X_t)\psi_0\psi_1\left(I_t^r\right)^{\psi_1-1}\left(R_t\right)^{-\psi_1}-\left(\frac{\delta}{C_t}\right)=0. \tag{12}$$

Thus

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

The term in square brackets, when scaled by $1/N_t$, is the social value of the knowledge stock of R&D expressed in units of (damaged) consumption.

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

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

where

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

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 is

$$\frac{\left[-\frac{\partial V}{\partial y}(X_t)\tilde{\theta}_t^*\right]-\mathcal{E}_t\left[\frac{\partial^2 V(X_t)}{\partial y^2}\right]|\varsigma|^2}{\delta\left(\frac{N_t}{C_t}\right)}, \tag{13}$$

and the social benefit is

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

both expressed in terms of damage-adjusted consumption as the numeraire. Notice that the model ambigu-

ity adjustment contributes to the social cost of carbon. These formulas are evaluated at the socially efficient allocation as is required for deducing the Pigouvian taxation. This is in contrast to many empirically-based measures. Notice also that the social cost of carbon (13) includes an explicit volatility adjustment because emissions in our model alter the local exposure to Brownian motion risk. Model-based measures that abstract from uncertainty influence by emissions outcome omit this term.

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 (Panel A) and emissions (Panel B). The simulations are not intended to be typical but rather are chosen to illustrate some of the possible outcomes. 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 the R&D investment trajectories from our stochastic pathways for 40 years. Each case begins with substantial R&D investment and grows stochastically for a while. Neither a damage severity 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 realization events, which reveal the λ_3 governing the damage function given by formula (6), 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 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.

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

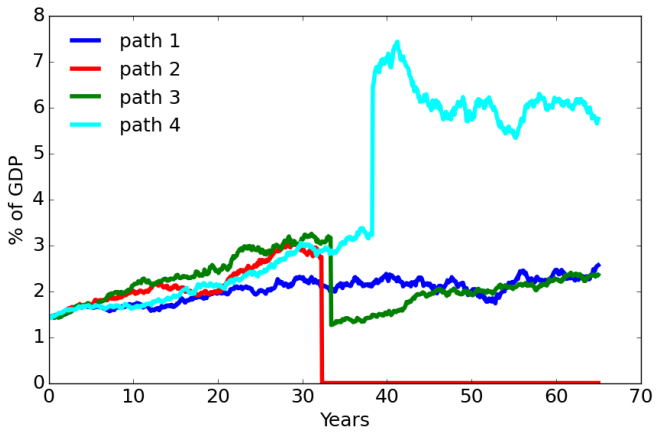

(A) R&D investment

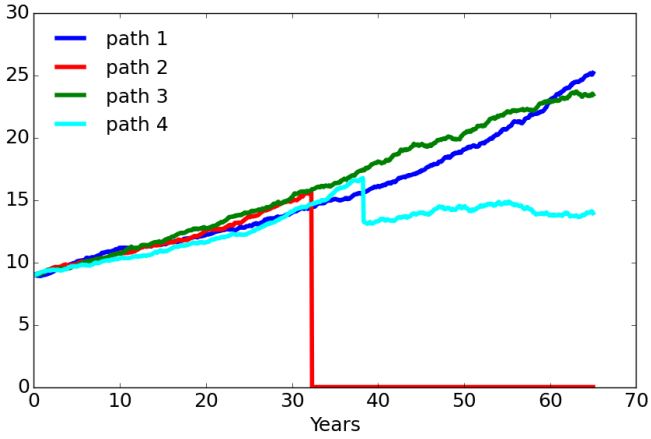

(B) emissions

Figure 3: Illustrative stochastic pathways computed with baseline probabilities and controls given by the $\xi = .1$, or less uncertainty aversion, specification. Panel (A) shows the optimal R&D investment choices for the stochastic simulation pathways. Panel (B) shows the 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 33. Path 3 (green) is a case with a damage jump ($\lambda_3 = .053$) at year 34. Path 4 (light blue) is a case with a damage jump ($\lambda_3 = .298$) at year 39

6 Uncertainty deconstruction

We start our quantitative findings with an assessment of the overall uncertainty contributions using the four channels we delineated in the introduction. Each channel 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 to the door to some more focused investigations of the uncertainty impacts.

We used two penalization parameters, ξ_m and ξ_a to represent the uncertainty aversion, the first 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 three labels or references:

- $\xi_m = \infty$, "baseline"
- $\xi_m = .1$, "less aversion"
- $\xi_m = .05$, "more aversion".

In addition, reported outcomes are computed at the initial time period for our simulations unless otherwise stated. For results reported across time horizons, the scaling factors used to compute the outcomes are based on the values from the initial time period for our simulations.

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 7% reduction in the social value of R&D when activating only the technology channel. This is in contrast to a 30% to 35% 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.

	SVRD		R&D investment / output	
Uncertainty channel	$\xi_m = .1$	$\xi_m = .05$	$\xi_m = .1$	$\xi_m = .05$
baseline	41.3 (66%)	41.3 (43%)	.0063	.0063
climate uncertainty	41.5 (66%)	41.7 (44%)	.0063	.0064
damage uncertainty	42.6 (68%)	43.9 (46%)	.0067	.0071
productivity uncertainty	40.5 (65%)	39.8 (42%)	.0061	.0058
technology uncertainty	57.8 (93%)	77.9 (82%)	.0123	.0223
all channels	62.4	95.2	.0144	.0334

Table 1: Social value of research and development (SVRD) and the R&D investment-output ratios when different uncertainty channels are activated. Columns 1 and 3 show the SVRD and R&D investment-to-output ratio, respectively, for the less uncertainty aversion case. Columns 2 and 4 show the SVRD and R&D investment-to-output ratio, respectively, 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 SVRD are the percent of the total value when not activating all channels simultaneously.

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.[19] 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 10. When we activate only the uncertainty concerns about the productivity channel, R&D investment decreases slightly when going from $\xi_m = .1$ to $\xi_m = .05$. In this case, the planner increases capital investment due to uncertainty in productivity growth.

Table 2 presents the results of the uncertainty decomposition for SCGW and emissions. 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	54.7 (58%)	54.7 (29%)	9.32	9.32
climate uncertainty	54.6 (58%)	54.6 (29%)	9.31	9.31
damage uncertainty	57.1 (60%)	59.5 (32%)	9.29	9.27
productivity uncertainty	54.2 (57%)	53.5 (29%)	9.32	9.33
technology uncertainty	84.2 (89%)	133.4 (71%)	9.09	8.79
all channels	94.6	187.3	9.01	8.48

Table 2: Social cost of global warming and emissions when different uncertainty channels are activated. Columns 1 and 3 show the SCGW and emissions, respectively, for the less uncertainty aversion case. Columns 2 and 4 show the SCGW and emissions, respectively, 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 SCGW are the percent of the total value when not activating all channels simultaneously.

7 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

[19] The contributions are not constructed to be additive.

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.[20] 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.[21] 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.

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 distortion when moving from less to more aversion underscores the sensitivity to 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 technology shock.[22] 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.

[20] 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.

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

[22] In Appendix D.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.

ξ_m	capital	knowledge stock
.05	-.184	-.008
.1	-.096	-.003

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

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).

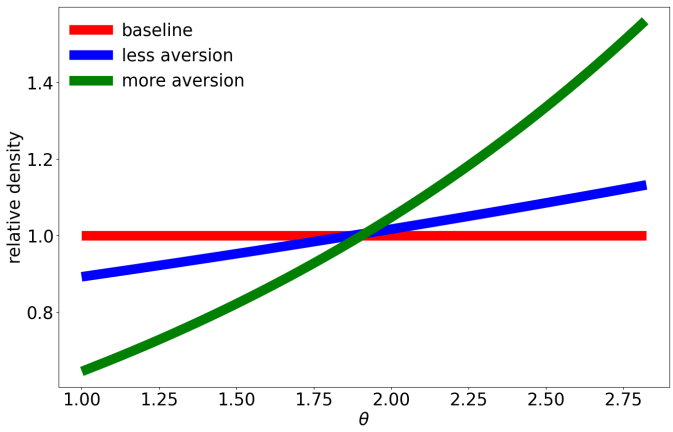

Figure 4: Uncertainty-adjusted densities for θ. The baseline and two uncertainty-adjusted densities are relative to a uniform baseline density. The resulting model averages are 1.86 for $\xi_m = \infty$, 1.89 for $\xi_m = .1$, and 1.95 for $\xi_m = .05$.

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 jumps or a technology jump. By conditioning first on the state variable history, we find that the probability that any

of the jumps has occurred in the time period t to be:

$$F_t = 1 - \exp\left[-\int_0^t \sum_{\ell=1}^{L} G_t^{*\ell} \mathcal{J}^\ell(X_\tau) d\tau\right],$$

which increases in t. We compute the probability for any of the jumps occurring by differentiating F_t with respect to t is:

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

We unpack this density construction so as to apply to each of the separate jump possibilities,

$$f_t^\ell \stackrel{\text{def}}{=} G_t^{*\ell} \mathcal{J}^\ell(X_t) \exp\left[-\sum_{\ell=1}^{L} \int_0^t G_t^{*\ell} \mathcal{J}^\ell(X_\tau) d\tau\right], \qquad \ell = 1, ..., L \tag{14}$$

which can be viewed as density with respect to a uniform measure over the possible time horizons and a uniform measure over discrete states. To obtain the corresponding *ex ante* densities of interest, we take conditional expectations over the future values of the state variables.

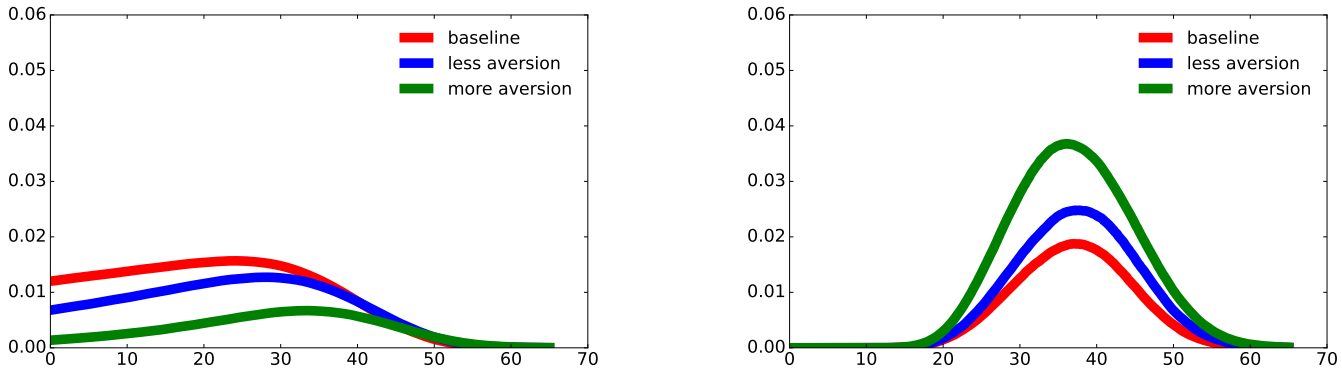

(a) Contribution of technology jump.

(b) Contribution of damage jumps.

Figure 5: Additive decomposition of the jump-time densities under the baseline probability and uncertainty-adjusted probabilities for two specifications of uncertainty aversion. The technology discovery contribution is 60% for the baseline specification, 47% for less aversion, and 22% for more aversion.

Figure 5 plots the *ex ante* density after aggregating over all of the possible damage jumps. The left panel of this figure shows how uncertainty aversion is reflected in the uncertainty-adjusted density of technology contribution, making a delayed success more likely. The delayed prospects for success makes *R&D* investment *less attractive*. Consistent with these probabilistic delays, as we increase the uncertainty aversion, the likelihood of the first jump being a damage jump becomes much higher under the uncertainty-

adjusted probability measures. The modal jump date for a damage jump occurs around 38 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.

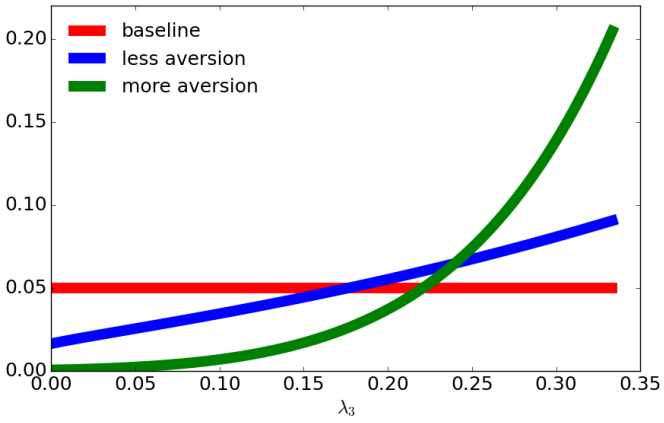

Figure 6: Uncertainty-adjusted probabilities across the alternative λ_3's for the base line probabilities and two specifications of uncertainty aversion. The resulting λ_3 averages are 0.167 for the baseline specification, 0.211 for less aversion, and 0.274 for more aversion.

Since there are multiple damage-curve realizations, 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 $[0, .33]$[23] 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 severity turn out to be much more consequential than the climate model uncertainty.

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

8 Acting now or waiting?

Our damage jump reveals the tail of the damage function associated with more extreme warming of the environment. Responses to the revelation of this information will have the usual impact of conditioning in that the responses are sensitive to what damage curvature is revealed. The responses to this jump 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

[23] We used twenty equally spaced discrete support points in our computations.

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 gives densities for the "before and after" emissions and R&D investment responses to knowledge of any of the collection of potential damage curve realizations. The emissions response shows an aspect of wait for more information as the overall density for emissions is shifted towards lower emissions after the tail of the damage curve is revealed. The median response differences are between $10\% - 12\%$. There is a modest reduction in emissions induced by increased uncertainty concerns prior to a jump, as seen by the small shift to the left in the density as ξ_m decreases. Thus, emissions sensitivity to the reported changes in the robustness parameter, ξ_m, while present, is quite modest.

	R&D investment / output					
	.1 quantile		.5 quantile		.9 quantile	
ξ_m	before	after	before	after	before	after
∞	1.32	0.61	1.57	2.21	1.81	3.34
.1	2.66	0.56	3.11	3.76	3.58	6.55
.05	4.94	0.37	5.91	5.40	6.76	9.52

Table 4: Quantiles for R&D investment before and after the damage realization for different values of ξ. R&D investment is reported as a percent relative to total output. Columns 1-2 report the outcomes for the .1 quintile, Columns 3-4 report the outcomes for the .5 quintile, and Columns 5-6 report the outcomes for the .9 quintile. Columns 1, 3, and 5 report the R&D investment-to-output ratio just prior to a damage jump. Columns 2, 4, and 6 report the R&D investment-to-output ratio just after a damage jump. Outcomes are shown for the three different values of the aversion parameter ξ_m in the respective rows.

As is illustrated in Figure 7, the R&D responses become much more heterogeneous after the Poisson damage event, as the investment decisions become tailored to potential damages induced by more extreme global warming. The spread in the uncertainty-adjusted density increases when we activate the robustness concerns. Table 4 reports quantiles that document this sensitivity where the lower quantiles decrease and the upper quantiles increase right after the jump. Prior to a damage jump, increased concerns about uncertainty lead to a leftward shift in the density for R&D investment and hence an increase in the reported quantiles. These results emphasize the intertemporal implications of uncertainty in our framework related to a more refined view of damage possibilities for the more extreme temperature changes. Moreover,

the findings highlight the amplified sensitivity to changes in uncertainty concerns for R&D investment in comparison to those for emissions.

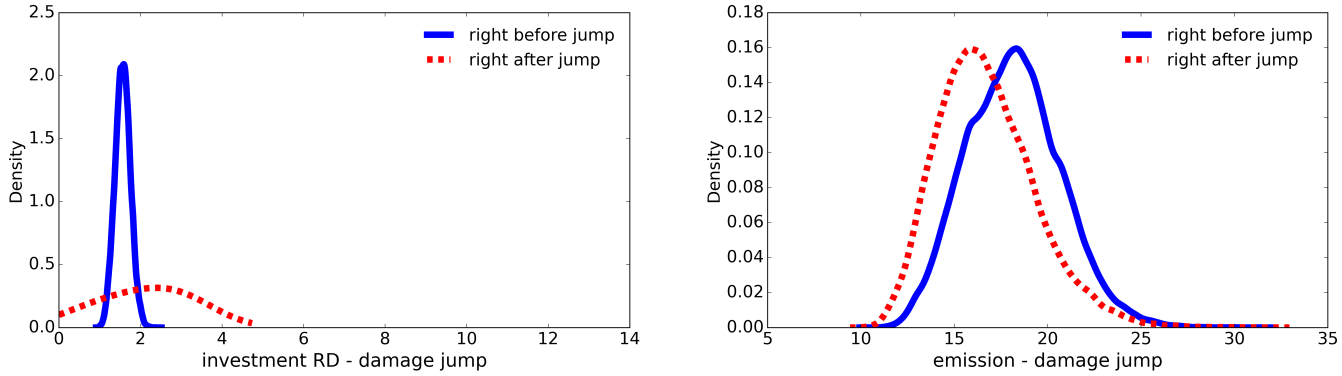

(A) $\xi_m = \infty$

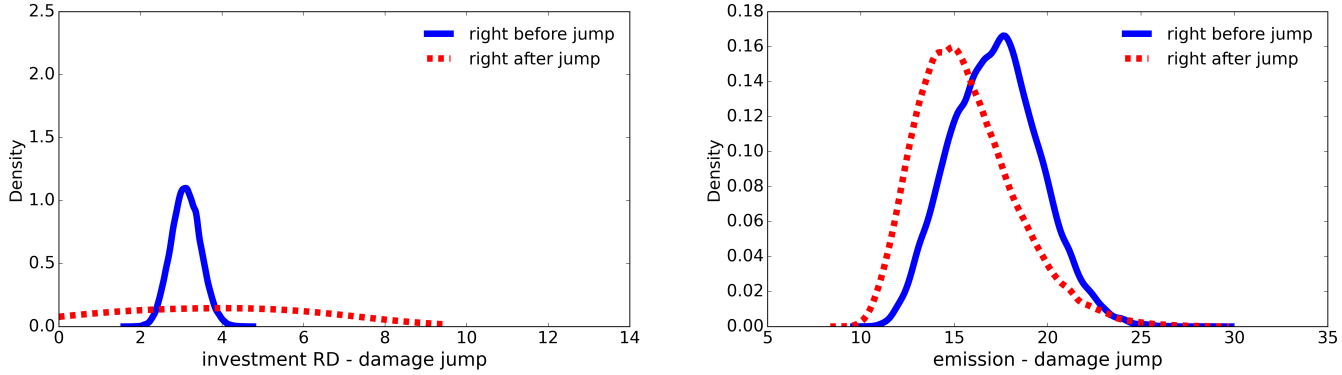

(B) $\xi_m = 0.1$

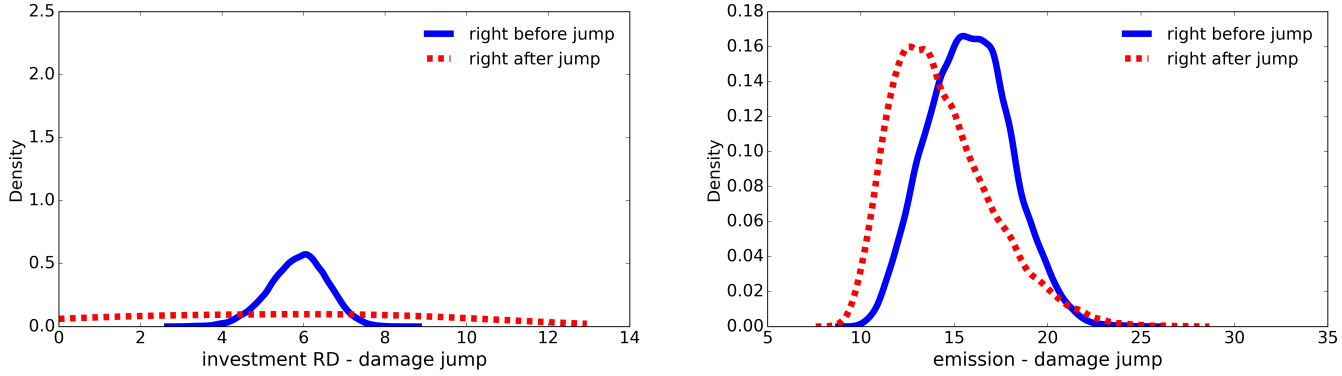

(C) $\xi_m = 0.05$

Figure 7: Densities for actions just before and just after a damage jump. Panel (A) shows results for the baseline specification. Panel (B) shows results for the less aversion specification. Panel (C) shows results for the more aversion specification. The left figure in each panel reports R&D investment and the right figure in each panel reports emissions. The blue line in each figure shows the distribution of outcomes just before a damage jump realization, and the red dashed line in each figure shows the distribution of outcomes just after a damage jump realization. Emissions are reported in gigatons of carbon (GtC),and R&D investment is reported as a percent relative to total output.

9 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 on 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 valuation. We construct our marginal value representations prior to the first jump to the first jump occurring.

As we have shown, this jump could be a damage severity jump 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:

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

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 in these formulas are the uncertainty adjustment factors for the intensities. The adjustment:

$$\exp\left(-\int_0^t \left[\sum_{\ell=1}^{L} G_\tau^{*\ell} \mathcal{J}^\ell(X_\tau)\right] d\tau\right). \tag{15}$$

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 for mathematical convenience.

We next describe a type of local sensitivity analysis that allows us to deconstruct the implications of the model. Key inputs in this and other related forms of sensitivity are marginal impulse response processes, Λ.[24] These processes 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.

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

$$\frac{\partial V}{\partial x}(X_0) \cdot \Lambda_0 = \widetilde{\mathbb{E}}\left[\int_0^\infty Dis_t \left(\Lambda_t \cdot Scf_t\right) \mid X_0, \Lambda_0\right] \tag{16}$$

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 not happening.

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:

$$\Lambda_t \cdot Scf_t.$$

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

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.

9.1 Marginal-utility flow

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

$$\text{flow i} \stackrel{\text{def}}{=} \delta \left(\frac{\partial \log c}{\partial x} - \frac{\partial \hat{n}}{\partial x} \right), \tag{17}$$

which is the date t vector of marginal utilities of damaged consumption and δ is a convenient utility scaling.

Again using the chain rule, we write the term (17) 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 $n = \exp[\hat{n}(y)]$, c/n is the adjusted consumption for damages. The first term on the left 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.

9.2 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.

9.2.1 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:

$$\text{flow ii} = \xi_m \sum_{\ell=1}^{L} g^{*\ell} \mathcal{J}^\ell \left(\frac{\partial \log \mathcal{J}^\ell}{\partial x} \right) \left(\exp\left[\frac{1}{\xi_m} \left(V^\ell - V \right) \right] - 1 \right) \tag{18}$$

where

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

is the robust adjustment to the intensity.[25] Observe that the continuation values are also affected by the magnitude of ξ_m.

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 uncertain damage information jump. 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 (18). 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. The partial derivative of interest in this case multiplies the post-jump continuation value contribution:

$$\xi_m \left(\exp\left[\frac{1}{\xi_m} \left(V^L - V \right) \right] - 1 \right) \geqslant V^L - V. \tag{19}$$

This contribution is pertinent because a technology discovery is very good news to the planner and so $V^L - V$ is typically positive, thus (19) results in a positive payoff to an R&D investment. 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 gives the planner *more incentive* to invest.

[25] Notice that this g^* construction of the uncertainty adjustment to an intensity depends on endogenous state variables. This facilitates solutions and computations. In the representations that follow the decision maker treats the intensity adjustments as exogenous as is formalized and justified in Hansen and Souganidis (2025). For instance, in constructing the flow ii term we do not include partial derivative of $\log g^{*\ell}$ with respect to the state vector.

Since the implied dynamical system has interactions over states, a marginal change in the knowledge stock in the initial period will impact all of the future states inducing cross effects in the valuation.

9.2.2 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:

$$\text{flow iii} \stackrel{\text{def}}{=} \sum_{\ell=1}^{L} g^{*\ell} \mathcal{J}^{\ell} \frac{\partial V^{\ell}}{\partial x}. \tag{20}$$

Notice that this term captures future marginal contributions to valuation as we may deduce post-jump counterparts to the pre-jump formulas depicted here for the partial derivatives: $\frac{\partial V^{\ell}}{\partial x}$ for the alternative ℓ's. Notice, in particular that the prospective technology jump contribution to this term will depend on how the post jump continuation values respond to the technology discovery. Given that a technology discovery happens,

$$\frac{\partial V^{L}}{\partial r} = 0,$$

so there is no direct marginal contribution of this 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 (14) where the these intensities in conjunction with the contribution (15) to the factor, Dis_t, account for the uncertainty-adjusted jump density.

9.3 Additional scaling adjustments

For the two capital stocks, our computations of state variable derivatives are expressed in terms of logarithms. Recall the formula:

$$\begin{aligned} \frac{\partial V}{\partial \log k} &= \left(\frac{\partial V}{\partial k}\right) k \\ \frac{\partial V}{\partial \log r} &= \left(\frac{\partial V}{\partial r}\right) r, \end{aligned}$$

making our calculations counterparts to expenditures. In the results we report, we also divide the valuations by current-period marginal utility of consumption inclusive of damages given by $\delta n/c$. This scaling converts our capital valuations into units of (damaged) consumption, making our computations more comparable

to the familiar intertemporal asset pricing valuation.

9.4 Decompositions

The integral representation (16) 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 is impacted. Together, these provide valuation counterparts to the innovation accounting methods often used in empirical macroeconomics.

Table 21 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 SVRD.

When viewing the decomposition results in Panels A and B, it is helpful to recall that in our model economy, after one of the damage severity jumps occurs, technology jump uncertainty remains an important concern. In contrast, once a technology discovery occurs, the damage severity jump 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 severity 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 to 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 severity jumps. The marginal changes in both future capital and R&D stock contribute to flow iii in roughly similar magnitudes. 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.

ξ_m	flow i	flow ii		flow iii		
		damage	tech	damage	tech	sum
.05	1.22	-0.57	4.04	5.94	0.60	11.22
.1	0.76	-0.22	3.22	2.54	0.74	7.04
∞	0.38	-0.08	2.48	1.26	0.50	4.55

(A) SVRD decomposition by prospective jump types.

ξ_m	flow i		flow ii		flow iii	
	capital	R&D	capital	R&D	capital	R&D
.05	0.63	0.61	0	4.04	2.80	3.78
.1	0.39	0.38	0	3.22	1.78	1.54
∞	0.19	0.19	0	2.48	0.95	0.83

(B) SVRD decomposition by state variable contributions.

Table 5: Decompositions of the three flow contributions to the SVRD delineated in Section 9. Panel (A) decomposes the the flow contributions to the SVRD across jump types. We aggregate the valuation contributions of all of the possible damage coefficients, $\lambda_3(\ell)$, for $1 \leqslant \ell \leqslant L-1$. The outcomes are reported in columns 2 and 4 show the contributions coming from the potential damage severity jump, while columns 3 and 5 show the contributions coming from the prospective technology jump. 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. Columns 1, 3, and 5 show the contributions coming from the capital stock channel, while columns 2, 4, and 6 show the contributions coming from the R&D knowledge stock channel. For both panels, the table shows how these decompositions depend on ξ_m.

9.5 Why might R&D investments respond proactively to enhanced uncertainty concerns?

From the first-order conditions (12) for R&D investment, we see a dependence on the social value of research and development (SVRD) as quantified by the marginal value of the knowledge stock. In this section we have isolated two potentially counteracting contributions to the SVRD.

- The uncertainty-adjusted probabilities of a technology discovery assigns more weight to distant successes relative to more recent ones as the aversion is enhanced.
- The endogenously determined net utility outcome from a potential successful discovery, measured by

$$V^L - V,$$

  increases as aversion increases because of a reduction in the pre-jump value function V. The reduction in V dominates any reduction in V^L because, post-jump, there is less exposure to uncertainty.

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 9.1. *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 more of side show, especially as it induces a less confident view of the prospects of a technological success. But others may disagree. 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 pro-active 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.*

10 A simplified model with only a technology jump

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. To do this, we assume that there is a single known damage curve, i.e. there is no damage severity 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 closed down uncertainty aversion related the diffusion contribution and the alternative climate models to further enhance the focus on the prospective technology discovery.

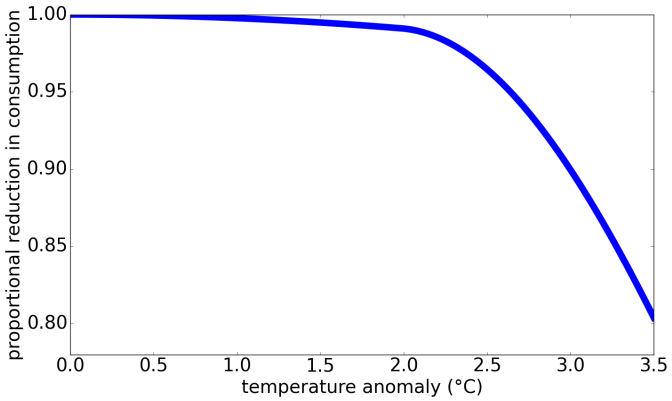

Figure 8: The consumption impact of the single damage function constructed by imposing $\hat{y} = 2$ and the arithmetic average of the $\lambda_3(\ell)'s$. The damages range from no impact on consumption (1.00) for a $0.0^{\circ}C$ temperature anomaly to a nearly 20% reduction in consumption (~ 0.80) for a $3.5^{\circ}C$ temperature anomaly.

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 33 years. The enhanced medians for the less averse and more averse specifications are about 43 years and 66 years, respectively. The uncertainty-adjusted densities make the R&D investment appear *less* attractive from a social valuation perspective.

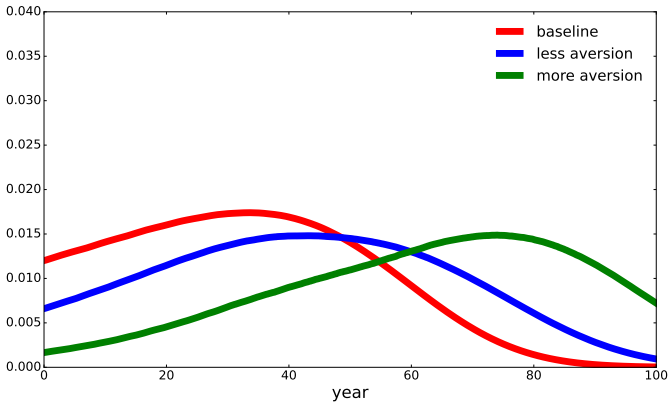

Figure 9: Jump time densities for the technology jump only model, including the baseline specification and uncertainty-adjusted densities for two alternative specifications of uncertainty aversion.

We now explore the impact of uncertainty aversion on social valuation while holding social cash flows fixed. This type of exercise is common in asset pricing analyses of endowment economies due to its pedagogical clarity. To implement it, we temporarily hold the investment decision rules fixed.

Consider first the valuation of flow i. Typically, an initial increase in either capital or the stock of knowledge induces a positive stochastic response in utility even after accounting for damages. Since

uncertainty-adjusted probabilities postpone the likelihood of a technological success, this resulting valuation will increase because the uncertainty-adjusted probabilities for a jump *not happening* are higher for all dates.

The prospect of a technology change *happening* is central to the construction of flows ii and iii. The impact of uncertainty aversion on their contributions to valuation will depend on the trajectories of the discounted stochastic responses to a marginal change in the stock of R&D. For our example economies the responses to a marginal change in the stock of knowledge is front-loaded, as we will see in a subsequent figure. Since the uncertainty-adjusted probabilities place relatively more weight on social payoffs further in the future, this impact alone would lead to a decrease in term ii. It turns out that this same impact can lead to an increase in the term iii contribution.

But this is not the end of the story. The post-jump continuation values also come into play. In this single jump model, the only nonzero entry of the gradient vector, $\frac{\partial \log \mathcal{J}^L}{\partial x}$, is the partial with respect to the knowledge stock. As a consequence, only the prospective marginal change in this state impacts the flow ii term. Reducing the penalization parameter, ξ_m, has two reinforcing impacts on the objective function for the minimization. Not only is $V^L - V$ increased, but the exponential adjustment

$$\xi_m \left(\exp\left[\frac{1}{\xi_m} \left(V^L - V \right) \right] - 1 \right)$$

has bigger implications for smaller values of ξ_m. Consider the counterpart for the flow iii term. The only nonzero entry for the gradient vector, $\frac{\partial V^\ell}{\partial x}$, is the partial with respect to the logarithm of capital since $R\&D$ investments are no longer desirable after the technology discovery. This term is further simplified by the fact that $\frac{\partial V^\ell}{\partial k} = 1$ is independent of ξ_m.

ξ_m	flow i	flow ii	flow iii	sum	R&D investment	capital investment
.05	1.75	6.76	2.17	10.68	.0283	.750
.1	1.02	4.68	1.78	7.48	.0151	.764
∞	0.53	3.47	1.13	5.13	.0075	.773

Table 6: Flow decomposition and initial investments for the one-jump specification. R&D and capital investment are reported as percents relative to total output. Outcomes are shown for the three different values of the aversion parameter ξ_m in the respective rows.

Since we entertain production, investments will also respond to changes in uncertainty concerns and impacts social valuation. Table 6 reports the three flow contributions to the SVRD for alternative aversions to uncertainty, and provides the initial robustly optimal investments for R&D and capital. In this simplified model, we continue to see more R&D investment when uncertainty concerns are greater. All three contributions to the SVRD increase with uncertainty concerns, including the flow iii term. In the absence of damage jumps, the flow ii term is now the major contributor to the SVRD, accounting for almost two-thirds of overall valuations. In addition, the flow i term is substantially more prominent as seen by comparing Table 6 to the previous Table 5.

Previously we explained why flows i and ii would *increase*, while term iii would *decrease*, if we held fixed the investment choices of the policy-maker. However, as we see from Table 6, flow iii actually increases. This outcome necessarily reflects changes in the stochastic responses of capital to a marginal change in the initial knowledge stock. This state interaction is confirmed in the upper-left panel of Figure 10. Figure 10A reports the conditional expectations of the discounted stochastic responses of (log) capital under uncertainty-adjusted jump probabilities for alternative horizons. The intertemporal contributions of flow iii are proportional to the corresponding expected responses because, as we argued previously, $\frac{\partial V^L}{\partial \hat{k}} = 1$. This can be seen by comparing the plots in Figures 10A and 10B. Due to the investment tradeoff in the allocation of output, capital initially increases, although it does eventually decline. The location 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 integrated (over the time horizon) flow iii term contribution increases, consistent with the results reported in Table 6.

The second row of Figure 10 reports the counterpart results for the flow ii term. The horizon decomposition provides a rather different accounting than for flow iii. As we see in Figure 10C, the knowledge stock responses move in the opposite direction from the capital responses, but without the peak behavior. The ordering of the initial responses is to be expected given the output constraint on the fictitious planner. The trajectories become flatter for smaller values of ξ_m (greater uncertainty aversion). These reduced responses when we increase the uncertainty aversion give a force for *reducing* the SVRD. As we noted previously, an

[26]To adjust for economically interpretable units, we formed the Panel (B) and Panel (D) trajectories by multiplying the Panel (A) and Panel (C) trajectories by initial period damaged consumption and dividing by the stock of knowledge. In addition, to measure the social payoffs of interest, we constructed the date t contributions of the Panel (B) trajectories by multiplying by $\frac{\partial V}{\partial \hat{k}}$ and similarly the date t contributions to the Panel (D) trajectories by multiplying by $\frac{\partial \log \mathcal{J}^L}{\partial \hat{r}}(X_t)\xi_m \left[\exp\left(\frac{1}{\xi_m}\left[V^\ell(X_t) - V(X_t)\right]\right) - 1\right]$.

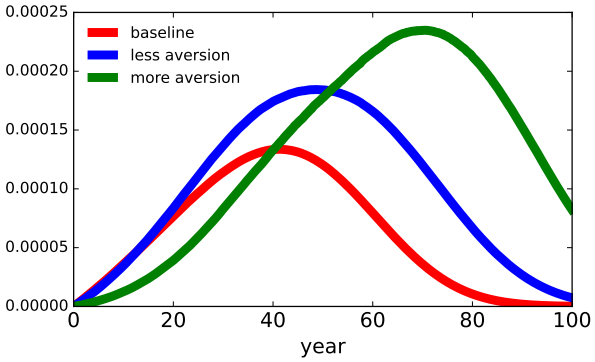

(A) Horizon decomposition of the expected discounted capital responses

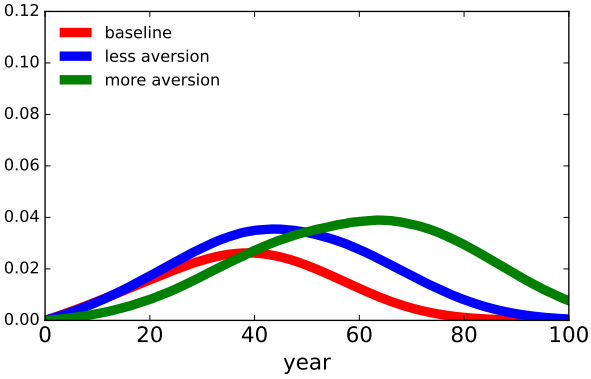

(B) Horizon decomposition of the flow iii contribution to the SVRD

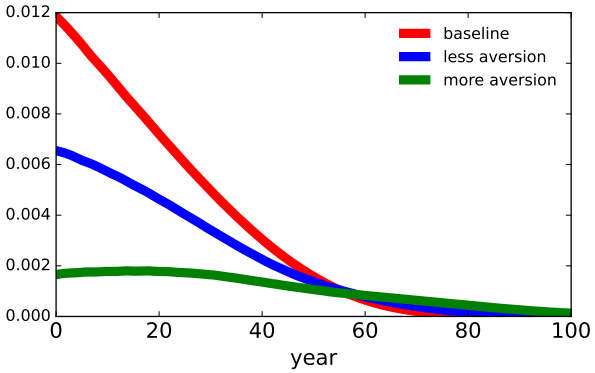

(C) Horizon decomposition of the expected discounted knowledge stock responses

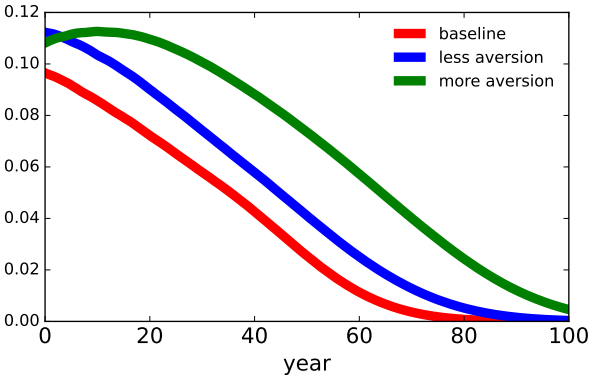

(D) Horizon decomposition of the flow ii contribution to the SVRD

Figure 10: Horizon decompositions of state responses and flow contributions for the SVRD. Panels (A) and (C) report expectations of the horizon-dependent $Dis_t \Lambda_t G_t^{*L} \mathcal{J}^{\ell}(X_t)$ conditioned on the initial time period information for the (logarithms of) capital and knowledge stock entries of Λ_t. Panels (B) and (D) report the corresponding horizon-dependent term iii and term ii contributions to the SVRDs.[26]Each panel plots trajectories under the baseline ($\xi_m = \infty$) specification as well as two alternative specifications of uncertainty aversion.

increase in the difference $V^L - V$ gives a counteracting force. The latter contribution dominates for the values of ξ_m in the tables and figures so far. This more than offsetting impact is evident by comparing Figures 10C and 10D. The decomposition in Figure 10D includes the contribution of the difference in the value functions after and prior to a hypothetical discovery. As we saw in Table 6, flow iii is the major contributor to the enhanced uncertainty response of the SVRD and the corresponding pro-active response of investment to uncertainty about the successful discovery of a new technology.

ξ_m	SVRD	R&D investment / output
∞	5.01	.0075
.10	7.12	.0151
.05	9.74	.0283
.01	9.77	.0288
.009	8.99	.0243
.008	8.14	.0200
.007	7.17	.0155

Table 7: Social value of R&D (technology jump only) as a function of the robustness parameter, ξ. Column 1 shows the knowledge stock valuation. Column 2 shows the R&D investment-to-output ratio. Rows range from the baseline ($\xi_m = \infty$) in row 1 to an extreme uncertainty aversion ($\xi_m = .007$) in row 7.

In summary, we describe two competing forces that contribute to the R&D investment outcome. 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. Table 7 shows that for very high misspecification aversions (very low values of ξ_m), the first force can dominate, leading instead to a reduction in the R&D investment.

11 Some final thoughts

Before concluding, we highlight some practical and conceptual points and suggest a couple of literature connections that are potentially worthy of further consideration. This discussion arises in part from questions about where our research might fit into a broader agenda.

In presenting our results, we have explored decompositions, sensitivity analysis, and different model specifications when framing our results. For practical reasons and space constraints, we have limited the scope of our sensitivity analysis, although the formulation and computational methods allow for further changes in the parameter configurations. Appendix E provides additional sensitivity checks on our main results. The first set of results focuses on separately changing the abatement technology scaling parameter (ϕ_0), the subjective rate of discount (δ), and the intertemporal elasticity of substitution (IES) ($1/\rho$). The

second set of results examines changing combinations of the R&D knowledge stock scaling parameter (χ), the knowledge stock volatility (σ_r), and the R&D investment effectiveness (ψ_0). The final sensitivity analysis test is to explore a two-step technological innovation framework, with the possibility of reaching an intermediate technology jump state before reaching the terminal innovation jump state. Although each set of sensitivity analysis exercises generates some quantitative differences in results, the qualitative implications for our uncertainty analysis persist across all settings: increased uncertainty concerns lead to more proactive investment in R&D prior to the informational event and more nuanced policy responses after the revelation of the severity of the damage.

There is a rather mature literature on investment posed as a decision about when to exercise a real option. The real options problem and our R&D investment problem have some similarities, but there are also some important differences. Our decision maker can influence the timing of the R&D payoff through investment, but there is not a specific choice of when literally to exercise the option. In our environment, there is an implicit investment cost due to our output constraint. Arguably, this plays a somewhat similar role to the cost of exercising an option, but the latter cost is only incurred on the date the option is exercised. In Appendix F, we explore the ways in which our findings have and do not have counterparts to insights from a real options example economy tailored from Miao and Wang (2007) in which uncertainty aversion and the form of the option payment come 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.

Finally, we address two speculations about the broader relevance of our analysis. First, in terms of pragmatic policy challenges, we suggest that government investment will play an important role in the development of new green technologies, at least until economic viability becomes more self-evident. A common concern for taking strong immediate action by skeptics is that the costs of using inefficient government approaches undermine the attractiveness of public solutions as opposed to market solutions. A blunt way to put this objection is, "This problem will not be solved by the government throwing money at it." While there are good reasons to be skeptical about political processes undermining the attractiveness of collective action, our robustness calculations suggest that the uncertainties may be large enough to push for efforts to overcome, at least partially, such political distortions. Although our analysis, so far, does not include a more serious investigation of the political economy of large-scale public investment projects, our analysis features the potential for stimulating early stage R&D investment in truly novel technologies

rather than through subsidies that create inefficiencies and special interest rent seeking.

Second, a commonly expressed concern regarding climate policy is the danger of focusing exclusively on one class of potential catastrophic events and omitting others, say future pandemics, nuclear war, political chaos, and even others that might catch us by surprise. While 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 others simultaneously. Thus, we chose 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 tries to look at tradeoffs in societal resources allocated to a menu of disasters. For such an ambition, we would find it prudent to embrace the potential interactions between climate change and other potential disasters. For instance, global warming can be associated with other precipitation events, such as extreme droughts (see, for example, Singh et al. (2022)), and could engender violence or more general group conflict (see, for example, Burke et al. (2024) and Hsiang (2025)). Our refinements to uncertainty quantification will be valuable in cutting a broader swath of extreme events.

12 Conclusions

We investigate how broad-based concerns about uncertainty impact robustly optimal policy responses to climate change. Our substantive findings emphasize the importance of including endogenous R&D in quantitative assessments of socially prudent actions along with current and near-term future carbon reductions. Our calculations expose the limitations of commonly proposed policy solutions that involve a gradual decrease in emissions to a net zero target without explicitly featuring an endogenous role for R&D investments or future informational "tipping points" and speculating on their potential successes.

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 impact prudent policies that address climate change. Of course, a more realistic policy en-

vironment requires additional modeling complexity, including differences in (i) how vulnerable alternative regions around the world are to climate change, (ii) the incentives multiple policy makers face, and (iii) the alternative forms of potential new green technologies potentially worthy of investment even with their concurrent doubts. 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 incentives for more proactive policies with enhanced uncertainty concerns remain a serious possibility.

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 our emphasis in our paper on investing in longer term technology success even if we remain unsure about its prospects.

Appendices

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

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

A HJB Equations

As part of dynamic programming backward induction, we start by giving the post jump HJB equations. The post jump value functions are then inputs into the pre-jump HJB equations. Recall there are L possible jump realizations from the pre-jump setting: $L-1$ potential damage curve realizations and one technology jump realization.

A.1 A numerically convenient transformation

Prior to constructing post damage jump value functions, we describe two numerically convenient transformations. Note that the damage function in this setting depends upon the temperature at which the damage jump occurred, $\hat{y}$. This dependence has a special structure that we exploit in computations. Define a new state variable transforming the temperature anomaly:

$$z \stackrel{\text{def}}{=} y - \hat{y} + \overline{y}.$$

At the time of the jump this new state variable will be initialized at $\overline{y}$. In (6), we constructed the log damage function $\hat{n}$. Note that this construction depends implicitly on both ℓ and $\hat{y}$, dependence that we now make explicit by writing $\hat{n}(y;\ell,\hat{y})$. We transform this log damage function for $y \geqslant \hat{y}$ by constructing:

$$\begin{aligned}
\hat{m}(z;\ell) &\stackrel{\text{def}}{=} \hat{n}(y;\ell,\hat{y}) - (\lambda_1\hat{y} + \frac{1}{2}\lambda_2\hat{y}^2) \\
&= \lambda_1(y-\hat{y}) + \frac{1}{2}\lambda_2(y-\hat{y}+\overline{y})^2 + \frac{1}{2}\lambda_3(\ell)(y-\hat{y})^2 - \frac{1}{2}\lambda_2(\overline{y})^2 \\
&= \lambda_1(z-\overline{y}) + \frac{1}{2}\lambda_2 z^2 + \frac{1}{2}\lambda_3(\ell)(z-\overline{y})^2 - \frac{1}{2}\lambda_2(\overline{y})^2
\end{aligned}$$

Importantly, the $\hat{m}$ is constructed so as to depend on z, but not separately on y and $\hat{y}$.

We will construct some post technology jump value functions by first using z and $\hat{m}$ and then transforming these value functions by adding back the adjustment:

$$-\left(\lambda_1 \hat{y} + \frac{1}{2}\lambda_2 \hat{y}^2\right).$$

This two-step approach gives us a numerically convenient way to capture dependence of value functions on the temperature anomaly realization, $\hat{y}$ when the damage severity jump takes place.

A.2 Post damage and technology jumps

We start with HJB equation after both a technology jump and one of the damage severity jumps has taken place. Denote the resulting value function $V^{\ell,L,\hat{y}}$ for $1 \leqslant \ell \leqslant L-1$ and the value function counterpart constructed using z and $\hat{m}$ as $W^{\ell,L}$. Importantly, $W^{\ell,L}$ does not depend on $\hat{y}$.

Since we are conditioning on technology jump having already occurred, $\phi_0 = 0$, and there is no incentive for a further R&D investment targeted for the discovery of clean technology. The resulting output technology is:

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

For the relevant state space, control set, and distortion set we use

$$x = \{z, \hat{k}\}, \qquad \Phi = \{i^k\}, \qquad \Gamma = \{h^k\}$$

Observe that temperature no longer varies over time. Since we use the transformed temperature variable z, we also use $\hat{m}$ instead of $\hat{n}$ in the HJB equation for $W^{\ell,L}$.

To compute $W^{\ell,L}$, we solve the HJB equation

$$\begin{aligned} 0 = \max_{i^k} \min_{h^k} \; & \delta \log(\alpha - i^k) + \delta \hat{k} - \delta \hat{m}(\cdot;\ell) - \delta W^{\ell,L} + \xi_m \frac{\left|h^k\right|^2}{2} \\ & + \frac{\partial W^{\ell,L}}{\partial \hat{k}} \left[\mu_k + i^k - \frac{\kappa}{2}\left(i^k\right)^2 - \frac{|\sigma_k|^2}{2} + \sigma_k h^k\right] + \frac{\partial^2 W^{\ell,L}}{\partial \hat{k}^2} \frac{|\sigma_k|^2}{2} \end{aligned}$$

Using guess and verify, the robustly optimal choice variables are given by:

$$\begin{aligned} h^{k*} =& -\frac{1}{\xi_m}\sigma_k \\ i^{k*} =& \frac{(1+\alpha\kappa) - \sqrt{(1+\alpha\kappa)^2 - 4\kappa(\alpha-\delta)}}{2\kappa}, \end{aligned}$$

and the value function by:

$$W^{\ell,L}\left(y,\hat{k}\right) = \log\left(\alpha - i^{k*}\right) - \frac{1}{\delta 2\xi_m}|\sigma_k|^2 + \frac{1}{\delta}\left[\mu_k + i^{k*} - \frac{\kappa}{2}\left(i^{k*}\right)^2 - \frac{|\sigma_k|^2}{2}\right] + \hat{k} - \hat{m}(y;\ell).$$

As a second step, we then construct:

$$V^{\ell,L,\hat{y}} = W^{\ell,L} - \left(\lambda_1\hat{y} + \frac{1}{2}\lambda_2\hat{y}^2\right),$$

which follows since the damage adjustment, $\hat{n}$, enters negatively into the discounted objective of the planner.

A.3 Post-damage but pre-technology

We next compute each of the post-damage, pre-technology jump values functions $V^{\ell,\hat{y}}$ where only a damage jump has been realized for $\ell = 1, ..., L-1$.

Again, we start with the HJB for the value functions W^ℓ, $1 \leqslant \ell \leqslant L-1$ using $\hat{m}$. The modified state vector x, control set Φ, and distortion set Γ are

$$x = \{\hat{k}, z, \hat{r}\}, \qquad \Phi = \{i^k, i^r, e\}, \qquad \Gamma = \{h^k, h^z, h^r, g^L\}.$$

The function W^ℓ solves:

$$\begin{aligned}
0 \;=\; & \max_{i^k,i^r,e}\min_{h,q,g^L} \delta\log\left[\alpha - i^k - i^r - \alpha\phi_0\left(1-\frac{e}{\beta\alpha k}\right)^{\phi_1}\right] + \delta\hat{k} - \delta\hat{m}(z;\ell) - \delta W^\ell(\hat{k},z,\hat{r}) \\
& + \frac{\partial W^\ell}{\partial \hat{k}}(\hat{k},z,\hat{r})\left(-\mu_k + i^k - \frac{\kappa}{2}\left(i^k\right)^2 - \frac{|\sigma_k|^2}{2} + \sigma_k h\right) + \frac{\partial^2 W^\ell}{\partial\hat{k}\,\partial\hat{k}'}(\hat{k},z,\hat{r})\frac{|\sigma_k|^2}{2} \\
& + \frac{\partial W^\ell}{\partial y}(\hat{k},z,\hat{r})e\left(\sum_{\theta\in\Theta}\theta q(\theta)\right) + \frac{\partial^2 W^\ell}{\partial y\,\partial y'}(\hat{k},z,\hat{r})\frac{|\varsigma|^2}{2}e^2 \\
& + \frac{\partial W^\ell}{\partial \hat{r}}(\hat{k},z,\hat{r})\left(-\zeta + \psi_0(i^r)^{\psi_1}\exp\left(-\psi_1\left(\hat{r}-\hat{k}\right)\right) - \frac{|\sigma_r|^2}{2} + \sigma_r h\right) + \frac{\partial^2 W^\ell}{\partial\hat{r}\,\partial\hat{r}'}(\hat{k},z,\hat{r})\frac{|\sigma_r|^2}{2} \\
& + \mathcal{J}^L g^L\left[W^{\ell,L}(\hat{k},z,\hat{r}) - W^\ell(\hat{k},z,\hat{r})\right] + \xi_m\mathcal{J}^L\left[1 - g^L + g^L\log g^L\right] + \frac{\xi_m}{2}h'h + \xi_a\sum_{\theta\in\Theta}q(\theta)\log q(\theta).
\end{aligned}$$

where, as an input, we impose function $W^{\ell,L}$ constructed in Section A.2.

As a second step we compute:

$$V^{\ell,\hat{y}}(\hat{k},y,\hat{r}) = W^{\ell,L}(\hat{k},y-\hat{y}+\bar{y},\hat{r}) - \left(\lambda_1\hat{y} + \frac{1}{2}\lambda_2\hat{y}^2\right).$$

Importantly,

$$V^{\ell,L,\hat{y}} - V^{\ell,\hat{y}} = W^{\ell,L} - W^\ell,$$

which is needed to make the HJB equation for W^ℓ to be relevant for the construction of $V^{\ell,\hat{y}}$.

A.4 Post-technology but pre-damage

Let V^L denote the value function. Once the technology jump has happened the subsequent damage severity jumps are inconsequential. This is true because temperature anomaly no longer increases in response to economic production since $\phi_0^L = 0$ and remains at a level for which damage curvature uncertainty is inconsequential.

In this case, we directly construct V^L restricting the temperature anomaly to reside in $0 \leqslant y < \hat{y}$. The function, V^L can be computed in essentially the same way $W^{\ell,L}$ as described in Section A.2. The HJB

equation for V^L is

$$0 = \max_{i^k} \min_{h^k} \delta \log(\alpha - i^k) + \delta \hat{k} - \delta \hat{n}(\cdot;\ell) - \delta V^L + \xi_m \frac{|h^k|^2}{2} + \frac{\partial V^L}{\partial \hat{k}} \left[\mu_k + i^k - \frac{\kappa}{2} \left(i^k \right)^2 - \frac{|\sigma_k|^2}{2} + \sigma_k h^k \right] + \frac{\partial^2 V^L}{\partial \hat{k}^2} \frac{|\sigma_k|^2}{2},$$

which again can be solved using guess and verify, applying a close counterpart to the solution in Section A.2. Note, from construction (6) that $\hat{n}$ does not depend on ℓ over the temperature anomaly domain $0 \leqslant y < \hat{y}$.

A.5 Pre-technology and pre-damage

Finally, we present the HJB equation for the value function V prior to any jump realization. We need as inputs to this equation the post jump value functions: $V^{\ell,\hat{y}}$ for $1 \leqslant \ell \leqslant L-1$ and V^L. Importantly, we only need the function $V^{\ell,\hat{y}}$ for $\hat{y} = y$ as prior to jumps being realized, the recursive equation compares the current continuation value, V, to what happens hypothetically at temperature anomaly y if a jump happens to occur at y. Thus we define:

$$V^\ell \stackrel{\text{def}}{=} V^{\ell,y}$$

for $\ell = 1, 2, ...L-1$. The complete HJB equation is:

$$\begin{aligned}
0 \;=\; & \max_{i^k, i^r, e} \min_{h, q(\theta), g^\ell} \delta \log \left[\alpha - i^k - i^r - \alpha \phi_0 \left(1 - \frac{e}{\beta \alpha k} \right)^{\phi_1} \right] + \delta \hat{k} - \delta \hat{n}(y) - \delta V(\hat{k}, y, \hat{r}) \\
& + \frac{\partial V}{\partial \hat{k}}(\hat{k}, y, \hat{r}) \left[-\mu_k + i^k - \frac{\kappa}{2} \left(i^k \right)^2 - \frac{|\sigma_k|^2}{2} + \sigma_k h \right] + \frac{\partial^2 V}{\partial \hat{k}\, \partial \hat{k}'}(\hat{k}, y, \hat{r}) \frac{|\sigma_k|^2}{2} \\
& + \frac{\partial V}{\partial y}(\hat{k}, y, \hat{r}) e \left[\sum_{\theta \in \Theta} \theta q(\theta) \right] + \frac{\partial^2 V}{\partial y\, \partial y'}(\hat{k}, y, \hat{r}) \frac{|\varsigma|^2}{2} e^2 \\
& + \frac{\partial V}{\partial \hat{r}}(\hat{k}, y, \hat{r}) \left[-\zeta + \psi_0 (i^r)^{\psi_1} \exp(-\psi_1 \hat{r}) - \frac{|\sigma_r|^2}{2} + \sigma_r h \right] + \frac{\partial^2 V}{\partial \hat{r}\, \partial \hat{r}'}(\hat{k}, y, \hat{r}) \frac{|\sigma_r|^2}{2} \\
& + \sum_{\ell=1}^{L} \mathcal{J}^\ell g^\ell \left[V^\ell(\hat{k}, y, \hat{r}) - V(\hat{k}, y, \hat{r}) \right] + \xi \sum_{\ell=1}^{L} \mathcal{J}^\ell \left[1 - g^\ell + g^\ell \log g^\ell \right] + \frac{\xi}{2} h' h + \sum_{\theta \in \Theta} q(\theta) \log q(\theta).
\end{aligned}$$

B Social planner preferences: IES not equal to one

We adopt a recursive representation of preferences in continuous time for the planner extended to allow $\rho > 1$. We start by forming the continuation value for each calendar date as follows:

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

where $\exp(V)$ gives an ordinally equivalent representation of preferences since $\exp(\cdot)$ is an increasing function. These preferences are dynamically consistent with a recursive representation.[27] The particular representation is homogeneous of degree one in consumption and $\exp(V)$.

The following differential equation gives the local representation expressed in terms of V:

$$\lim_{\epsilon \downarrow 0} \frac{1}{\epsilon} \left(V_{t+\epsilon} - V_t\right) = -\frac{\delta}{1-\rho} \left[\left(\frac{(C_t)^{1-\rho}}{\exp\left[(1-\rho)V_t\right]}\right) - 1\right] \tag{21}$$

which is a backward recursion linking future continuation values and current consumption to the current continuation value.

We introduce stochasticity under a unitary risk specification, which we implement by using conditional expectation for the recursive representation of V :

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

As in case where $\rho = 1$, we use the adjustments derived in Section 3 to replace the local means of the continuation value process with the robust counterparts.

C Production function interpretation of abatement

The type of mathematical formulation used to represent abatement in our framework has been applied in the literature on climate-economics many times since the work of Nordhaus. See, for example, Nordhaus (2017). One possible interpretation is as follows. Consider a fixed proportions technology for energy and capital. In such a setting, a reduction in dirty energy less than the required proportion of capital must be

[27] This consistency is evident by raising both sides of the equation V_t to the power $1-\rho$ and scaling by $\frac{1}{1-\rho}$ to increase the transformation.

replaced by a clean alternative subject to a convex cost. The resulting modification in the output equation would be:

$$C_t + I_t^k + I_t^r + \tilde{\phi}_0 \left(K_t\right)^{1-\phi_1} \left(\beta\alpha K_t - \mathcal{E}_t\right)^{\phi_1} \mathbf{1}_{\{\mathcal{E}_t < \beta\alpha K_t\}} = \alpha K_t$$

for $\tilde{\phi}_0 \stackrel{\text{def}}{=} \phi_0 \beta^{-\phi_1} \alpha^{1-\phi_1}$. The term:

$$\tilde{\phi}_0 \left(K_t\right)^{1-\phi_1} \left(\beta\alpha K_t - \mathcal{E}_t\right)^{\phi_1} \mathbf{1}_{\{\mathcal{E}_t < \beta\alpha K_t\}}$$

is often referred to as an abatement cost, perhaps even a cost that is external to the firm.

Although we find this representation to be substantively interesting, in our framework we may also view the right-hand side of (1) as a production relation without reference to an abatement cost. This, of course, is a matter of interpretation, but it also impacts how we think of output and of plausible calibrations of ϕ_0 and ϕ_1. Both have their attractive features and limitations. The production function interpretation allows for curvature that is missed in a fixed proportion technology but assumes that all of the energy inputs are dirty. Observe that when $\mathcal{E}_t = 0$ the output or output net of the abatement cost is $\alpha K_t(1-\phi_0)$, depending on how one views the technology and cost structure.

We now verify that the desired mathematical properties that support a production function interpretation. We show that the first derivatives are positive and that the second-derivative matrix is negative semi-definite, when we interpret the following:

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

for

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

as a production function.

Notice first that the candidate production function is homogeneous of degree one in (e, k). Then consider

the partial derivatives with respect to e:

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

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

Next, we consider derivatives with respect to k:

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

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

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

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

The negative semi-definite Hessian matrix follows since

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

Next, we compute the implied energy demand price elasticity. The first-order conditions for emissions are:

$$\frac{\partial V}{\partial y}(\bar{\theta}+\varsigma h)+\frac{\partial^2 V}{\partial y\,\partial y'}|\varsigma|^2 e+\delta\,(c)^{-1}\,\frac{\phi_0\phi_1}{\beta}\left(1-\frac{e}{\beta\alpha k}\right)^{\phi_1-1}\mathbf{1}_{\{e<\beta\alpha k\}}=0.$$

Divide by the marginal utility of damaged consumption to get:

$$\left(\frac{1}{\delta n}\right)\frac{\partial V}{\partial y}c(\bar{\theta}+\varsigma h)+\left(\frac{1}{\delta n}\right)\frac{\partial^2 V}{\partial y\,\partial y'}c|\varsigma|^2 e+\frac{\phi_0\phi_1}{\beta n}\left(1-\frac{e}{\beta\alpha k}\right)^{\phi_1-1}\mathbf{1}_{\{e<\beta\alpha k\}}=0.$$

View the first term on the left to be the shadow price of energy demand. Then use the implicit function theorem to compute: $\frac{de}{dp}$.

D Parameter values for the example economy

We quantitatively discipline our analysis by choosing model parameters based on (i) external empirical estimation and measurement of relevant economic and asset pricing outcomes, (ii) direct calibration using simplified versions of the model to match steady-state implications to empirical macroeconomic moments, and (iii) indirect calibration comparing outcomes from the simplified model settings to observable empirical macroeconomic and asset pricing moments. We validate these calibrated parameter choices by examining how the calibrated model fits additional macroeconomic and asset pricing empirical moments not included in the calibration. In addition, we perform sensitivity analysis for key parameters and specifications related to preferences, abatement technology, and climate damage as an additional validation of our quantitative results. In what follows, we provide the details for calibration, validation, and sensitivity analysis.

The model parameters chosen for preferences and the dynamics of the capital stock, knowledge stock, climate and climate damages are given in Tables 8 – 12. Steady-state values that are adapted to empirical

moments for capital dynamics calibration are given in Table 13. The initial values and ranges for the state variables are provided in Tables 14 – 15. The descriptions of these choices are given below.

Parameter	Value
α	.115
μ_k	.043
κ	6.67
σ_k	[0, 0.01]

Table 8: Capital dynamics

Parameter	Value
ζ	0
ψ_0	.1
ψ_1	.5
σ_r	[0, 0.0078]
χ	.0013

Table 9: Knowledge dynamics

Parameter	Value
$\bar{\beta}$	.12
ϕ_0	.5
ϕ_1	3
$\bar{\theta}$	1.86 / 1000
ς	$1.2 \times \bar{\theta}$

Table 10: Climate dynamics

Parameter	Value
δ	.01
ρ	1
ξ_m	$\{.05, .1, \infty\}$
ξ_a	$\{.00175, .0035, \infty\}$

Table 11: Preferences

In constructing Table 11, we set the model ambiguity parameter ξ_a so that at the initial date, the implied increase in the model average matched a run with model ambiguity aversion replaced by model misspecifaction concerns for all of the evolution equations, including the temperature evolution equation. We performed the computation for each choice of ξ_m. The implied values of ξ_a are reported in the last row.

Parameter	Value
λ_1	.00017675
λ_2	$2 \times .0022$
$\{\lambda_3(\ell)\}_{\ell=1,...,L-1}$	$\left\{\frac{1}{3}\frac{\ell-1}{L-2}\right\}_{\ell=1,...,L-1}$
d_0	1.5
d_1	0.36
d_2	0.0001
$\underline{y}$	1.5
$\bar{y}$	2.5

Table 12: Climate damages

Variable	Value
Investment/capital[a]: i^k	.090
Growth rate of capital[a]: η	.020
Marginal value of capital[a]: π	2.50

Table 13: Imposed steady states for the model specification without climate impacts.

State variables	values
K_0	739
Y_0	1.1
R_0	11.2

Table 14: State Variable Initial Values

State variables	range
$\log(K)$	$[4, 7]$
Y	$[0, 4]$
$\log(R)$	$[1, 6]$

Table 15: State variable ranges

D.1 Preferences

The subjective discount rate is set to a value of $\delta = 0.01$, consistent with the value used by others in the macroeconomics and asset pricing literature, including Barnett et al. (2020, 2022) and Barrage and Nordhaus (2023). The baseline choice of the IES is set to $\rho = 1$, which is a convenient baseline

configuration. As a sensitivity analysis, we examine outcomes for $\rho = 2/3$ and $\rho = 3/2$. As Hansen et al. (2024) emphasize, the specification of ρ has an important impact for the overall consumption-savings choice and for this reason could interact with calibration choices for the productivity of capital. In our baseline analysis we choose $\xi \in \{.05, .1\}$. These values provide probability distortions that we view as reasonable based on the outcomes reported in the results for our numerical example and useful in illustrating the effects of the model uncertainty mechanism in our framework. Since the numerical magnitude of ξ reflects social preferences, we are very receptive to considering implications for alternative values of ξ. Our computations and code are easily implemented for such explorations.

D.2 Capital dynamics

The choices of parameters for the productivity and evolution of productive capital follow Barnett et al. (2022), who use an undamaged version of the consumption capital model to calibrate the economic growth rate to a value of 2%, a marginal value of capital of 2.5, and an investment-capital ratio of .09, consistent with empirical values from the BEA and World Bank databases. Specifically, the calibration approach uses externally estimated model parameters, as well as the growth rate η and steady state values for (π, i_k), as inputs for the relevant non-climate-state evolution equations and FOCs. We then invert the system of equations to derive the remaining model parameters of interest, exploiting the tractable recursive structure of this problem as follows:

a) From the first-order investment conditions, solve for κ given (i^k, π):

$$\pi = \left(1 - \kappa i^k\right)^{-1}$$

b) From the growth equation, solve for $\overline{\mu}_K$ given η, κ, and i^k:

$$\eta = -\overline{\mu}_K + \frac{\kappa}{2}\left(i^k\right)^2,$$

c) Given (c, i^k), we determine α from the output constraint:

$$\alpha = i^k + \delta\pi.$$

For our chosen set of inputs, the resulting parameter values are given by $\kappa = 6.667, \mu_k = .043$ and $\alpha = .115$. Finally, the capital volatility is set to $\sigma_k = [0, .01]$, matching annual percent changes in the time series of GDP from the World Bank database.

D.3 Knowledge stock dynamics

For the R&D investment parameters, we make the following choices. We assume the depreciation of R&D stock is given by $\zeta = 0$ and set the returns-to-scale value to be $\psi_1 = .5$, each for simplicity and computational tractability purposes. We then set the R&D investment cost scaling to be $\psi_0 = .1$ and set the multiplicative factor that translates US R&D stock to a global value in units of the arrival rate for our Poisson jump process for technological change to be $\chi = .0011$, such that the simplified no-damage-jump model without uncertainty version generates R&D investment values that peak below .5% of GDP, in line with the peak for major U.S. R&D investment programs as estimated by Stine (2008), and estimates for returns to R&D investment from Lucking et al. (2019) and Bloom et al. (2019). Moreover, these parameter values also produce an expected arrival time for the green technological innovation consistent with proposed policy timelines of a carbon neutral transition occurring between 2050 and 2080. Alhamdan et al. (2023) similarly cite major U.S. R&D investment programs such as the Manhattan project as an anticipated timeline for one type of major green energy innovation, nuclear fusion development. The volatility of knowledge stocks is set to $\sigma_r = [0, .0078]$, matching the annual percent changes in the time series of U.S. R&D capital stock from the BLS database.

D.4 Climate dynamics and damages

Our choices for the emissions component of the production technology (i.e., the abatement cost parameters following the interpretation of Nordhaus and others) are as follows. We set $\phi_1 = 3$, similar to the values of the estimated parameter of Cai and Lontzek (2019) and Barrage and Nordhaus (2023). Although the value of ϕ_0 is highly uncertain, we choose $\phi_0 = .5$ as a reasonable benchmark for the fraction of lost output to achieve zero emissions. We also consider $\phi_0 = .1$, consistent with Barrage and Nordhaus (2023), for a sensitivity analysis comparison. The value of the intensity of the output emissions $\beta = .12$ comes from the implied emissions intensity value for 2020 from Cai and Lontzek (2019), which together with the chosen subjective discount rate generates annual carbon emissions consistent with estimates by Figueres et al. (2018) of around 10-11 GtC in a simplified version of the model without uncertainty aversion or jumps.

The climate dynamics parameter values are set as follows. Specifically, $\bar{\theta}$ is the average value across 144 climate model outcomes constructed from pulse experiments provided by Joos et al. (2013) and Geoffroy et al. (2013), and matches the values reported in Masson-Delmotte et al. (2021). The value for ς is chosen based on a parameter uncertainty interpretation of TCRE specification. Specifically, the chosen value is based on the implied standard deviation associated with the coefficient of the Matthew's approximation for a constant emissions path near the current value.

As noted previously, for climate damages, the functional form and values of λ_1, λ_2, and $\lambda_3(\ell)$ are roughly consistent with the range of climate damage specifications from the literature, including Nordhaus (2019), Weitzman (2012), and the very recent study by Waidelich et al. (2024). We choose baseline values for our intensity function of $\underline{y} = 1.5$ and $\overline{y} = 2.5$, motivated by the literature on climate thresholds and tipping points, which we elaborate on in Remark D.1. We set $\mathsf{d}_0 = 1.5$ and $\mathsf{d}_2 = 0.0001$, and then impose $\mathsf{d}_1 = 0.36$, such that the modal point of the density for the jump intensity is $y_m = 2.0$ based on a simplified model of the temperature pathway.[28]

Remark D.1. *The choices of $\underline{y}$ and $\overline{y}$, used for the range over which climate damage function jump realizations occur, are motivated by the literature on "climate tipping points" and "climate thresholds." Specifically, climate scientists, economists, and others are concerned about potentially drastic shifts in climate-carbon dynamics and realized economic damages resulting from continued climate change. Drijfhout et al. (2015) and Armstrong McKay et al. (2022) enumerate and characterize various potential thresholds of potential consequence, with the latter stating that "[t]he Earth may have left a safe climate state beyond $1^{\circ}C$ global warming. A significant likelihood of passing multiple climate tipping points exists above $\sim 1.5^{\circ}C$, particularly in major ice sheets. The probability of a tipping point increases further in the Paris range of 1.5 to $< 2^{\circ}C$ warming." Rogelj et al. (2018) and Rogelj et al. (2019) suggest a $1.5^{\circ}C$ target to limit damages from such events, while noting $2.0^{\circ}C$ as a potentially more plausible target that also includes risks of more severe damage consequences.*

However, consensus on the plausibility and timeline for such thresholds is far from decided. Hansen et al. (2025) notes that while tipping point concerns are real, "[m]any tipping point processes are reversible if Earth cools, but the recovery time varies and may be long for some feedbacks" and that the most threatening and catastrophic thresholds are still likely to occur far into the future. Ritchie et al. (2021) conclude that

[28] For sensitivity analysis, we also consider cases with $\overline{y} = 2.0$ and $\overline{y} = 3.0$. The values for $\underline{y}$, r_0, and r_2 remain fixed at the baseline values for these cases, while the value of r_1 is adjusted so that the modal point of the density for the jump intensity is $y_m = 1.75$ or $y_m = 2.25$, respectively.

"the point of no return" for climate thresholds is highly uncertain based on their analysis using recent developments in dynamical systems theory.

We note that our damage jump specification is not an immediate drop in output or a significant realization of climate damage. Instead, our damage jumps represent information revelation about the severity of climate damage function curvature going forward, including the possibility of more severe consequences similar to tipping points or thresholds. Thus, our damage function specification allows for concerns about potential tipping points or threshold in such a way that we believe is still consistent with the more moderate dynamics emphasized by Hansen et al. (2025) and Ritchie et al. (2021), rather than a stark "falling off the cliff" specifications used in some other settings.

D.5 Climate model uncertainty

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

D.6 Initial values

The initial value of capital is set so that our initial GDP matches the 2020 World GDP value of \$85 trillion estimated by the World Bank National Accounts data. With our choice of $\alpha = 0.115$, we end up with $K_0 = 739.13$. The initial value of knowledge capital is set to $R_0 = 11.2 \times \chi$, which converts and scales the value for current US R&D capital stock in the BLS database to a global value such that the expected arrival time of a breakthrough green technological change in the simplified no-damage-jump model without uncertainty aversion is around 30-50 years. The initial value of atmospheric temperature anomaly is set to $Y_0 = 1.1$ degrees Celsius to match recent estimates from the IPCC AR6 (Masson-Delmotte et al. 2021).

D.7 Some additional asset-pricing implications

For additional context about the model's calibration and quantitative implications, we examine (shadow) local asset pricing moments in the form of the implied risk-free rate from the model solution. We examine the outcomes under the baseline ($\xi = \infty$), less aversion ($\xi = 0.1$), and more aversion ($\xi = 0.05$) specifications. We calculate numerically the one-month risk-free rate using simulations that accounts for all

relevant uncertainty components. The simulation-based expression for the monthly, risk-free rate expressed in terms of annualized time units is given by

$$r_{f,t} \quad = \quad \delta - \frac{1}{\epsilon} \log \left(\tilde{\mathbb{E}} \left[\frac{C_t/N_t}{C_{t+\epsilon}/N_{t+\epsilon}} \mid \mathfrak{F}_t \right] \right)$$

where $\tilde{\mathbb{E}}\left[\cdot \mid \mathfrak{F}_t\right]$ denotes the expectation under the uncertainty-adjusted probability distribution, $N_t = \exp\left[\hat{n}(Y_t)\right]$. and $\epsilon = 1/12$ is the monthly specification. The results are provided in Table 16. Importantly, the values for each uncertainty aversion specification are within the range of reasonable values used in the asset pricing literature based on empirical estimates of the risk-free rate. Moreover, the increasing values of the risk-free rate as the uncertainty aversion increases reflects increasing concerns of negative outcomes that will impact future consumption. Note that because we are examining equilibrium solutions under social optimality, these do not represent market-based prices from the decentralized solution. Nevertheless these values provide context for the quantitative plausibility of our model's calibration.

	baseline ($\xi = \infty$)	less aversion ($\xi = 0.1$)	more aversion ($\xi = 0.05$)
risk-free rate ($r_{f,0}$)	2.3%	2.6%	2.8%

Table 16: Instantaneous risk-free rate for the full model computed under the baseline ($\xi = \infty$), less aversion ($\xi = 0.1$), and more aversion ($\xi = 0.05$) specifications. The values are computed at the initial time period using stochastic simulations under the corresponding (uncertainty-adjusted) probability distribution.

E Model sensitivity

We now discuss sensitivity analysis exercises that provide further insight into the economic mechanisms underlying our quantitative implications. We first outline the implications for changing the subjective rate of discount (δ). We then explore the effects of changing the intertemporal elasticity of substitution (IES) (ρ). Next, we examine the consequences of changing the abatement technology scaling parameter (ϕ_0). Additionally, we compare the effect of changing the combination of the R&D knowledge stock scaling parameter (χ) and the R&D investment effectiveness parameter (ψ_0), while briefly noting implications for changing the knowledge stock volatility (σ_r) as well. Finally, we explore the impacts of featuring a two-step ladder technological innovation framework. We provide a subset of results and intuition in what follows, with a complete set of detailed analysis and computational results provided in the Online Appendix.

E.1 Subjective rate of discount

The prior environmental economics literature has explored the sensitivity of the SCC to changes in the "discount rate." Often these analyses feature the implied discount rate used in computing present values abstracting from stochastic discounting. In our setting with misspecification, stochastic discounting and endogenously determined changes in the probability measure are central ingredients in valuation, as is expected from our understanding of asset pricing.

In Table 27 we report computations for both the social value of R&D and the social cost of global warming. Our results show that increasing δ leads to drops in valuations, confirming a sensitivity often noted in the literature on environmental economics. Since the reported numbers are in a logarithmic scale, we see drops in the valuations by about 30% by increasing δ from .01 to .015, and by an additional 30% to nearly 40% when increasing δ to .02.

	SCRD			SCGW		
Uncertainty aversion	$\delta = .02$	$\delta = .015$	$\delta = .01$	$\delta = .02$	$\delta = .015$	$\delta = .01$
More Aversion	0.030	0.041	0.051	-0.060	-0.081	-0.100
Less Aversion	0.021	0.027	0.033	-0.035	-0.043	-0.051
Neutrality	0.014	0.018	0.022	-0.023	-0.026	-0.029

Table 17: Initial robust actions for two alternative initial specifications of ϕ_o.

In Online Appendix D we report the social value of R&D and the social cost of global warming counterparts for each of the uncertainty aversion scenarios: uncertainty neutral ($\xi_m = \infty$), less aversion ($\xi_m = .05$), and more aversion ($\xi_m = .01$). In each uncertainty aversion case, we similar drops in the valuations as δ increases. In addition, we note that increasing uncertainty continues to amplify the valuations across all values of δ considered.

E.2 Intertemporal elasticity of substitution (IES)

We also consider alternative values of the IES using the preferences defined in Appendix B. Much of the asset pricing literature has studied the consequences of changing the IES on asset valuation within the setting of an endowment economy. However, our economy is a production economy and changing the IES has a big impact on production outcomes. As expected from growth models, investments in both types of capital relative to output are higher when the elasticity is greater (ρ is smaller) and the valuations of the corresponding capital stocks move in the opposite way, as expected. The observed growth sensitivities in

emissions prior to any jumps are consistent with the investment differences since higher investment helps support more initial growth in output.[29]

	R&D-output ratio			Investment-output ratio		
Uncertainty aversion	$\rho = 3/4$	$\rho = 1$	$\rho = 4/3$	$\rho = 3/4$	$\rho = 1$	$\rho = 4/3$
more aversion	5.6%	3.4%	1.9%	83.4%	74.6%	68.8%
less aversion	2.6%	1.4%	0.8%	86.5%	76.6%	69.9%
neutrality	1.3%	0.6%	0.3%	88.0%	77.4%	70.3%

Table 18: Initial investment to output ratios for different specifications of the IES.

IES	$\log SVRD$	$\log SCGW$
$\rho = 3/4$	0.0371	-0.0551
$\rho = 1$	0.0334	-0.0506
$\rho = 4/3$	0.0275	-0.0436

Table 19: Social values at the initial time period for less aversion to misspecification uncertainty.

Table 18 shows the results for two alternative specifications of the IES that differ from unity: $\rho = 3/4$ and $\rho = 4/3$. As noted previously in relation to growth models, the investment in both types of capital relative to output are higher when the elasticity is greater (ρ is smaller). Table 19 shows the (log) social value of the R&D stock and the (log) social cost of global warming for $\rho = 3/4$, $\rho = 1$, and $\rho = 4/3$.

E.3 Abatement technology sensitivity

As we noted in Appendix C, economic models of climate change often make reference to an "abatement cost" with ϕ_0 representing the magnitude of the cost of full abatement and ϕ_1 capturing the convexity of the cost function. Under this alternative interpretation, previous work led by Nordhaus use an alternative calibration of $\phi_0 = .1$. Table 20 shows the R&D-to-output ratio and emissions for our baseline value of $\phi_0 = 0.5$, the Nordhaus-motivated alternative of $\phi_0 = .1$, and an intermediate value of $\phi_0 = .25$. We find that decreasing the value of ϕ_0 leads to i) substantially less R&D as a fraction of output, ii) a more modest but notable proportional reduction in emissions, and iii) the same qualitative response to model uncertainty implications. These responses reflect the reduced output loss induced by a reduction in emissions would

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

for lower values of ϕ_0.

	R&D-output ratio			emissions		
Uncertainty aversion	$\phi_0 = .5$	$\phi_0 = .3$	$\phi_0 = .1$	$\phi_0 = .5$	$\phi_0 = .3$	$\phi_0 = .1$
More Aversion	3.16%	2.01%	0.64%	8.53	8.34	7.95
Less Aversion	1.38%	1.02%	0.45%	9.04	8.81	8.25
Neutrality	0.62%	0.53%	0.31%	9.32	9.10	8.47

Table 20: Initial robust actions for two alternative initial specifications of ϕ_o.

Online Appendix C shows the numerical outcomes for changing the value of ϕ_0 for different choices of the R&D knowledge stock related parameters, namely χ and ψ_0. Importantly, the qualitative implications for changing ϕ_0 and for the impact of uncertainty aversion persist in these alternative parameter specifications.

E.4 R&D knowledge stock parameter sensitivity

Given the novel implications of model uncertainty emphasizing proactive R&D investment early prior to the revelation of damage severity and targeted responses after the severity is fully revealed, we also provide sensitivity analysis relative to the specification of R&D investment and knowledge stock evolution in our model. Specifically, we examine model outcomes for different combinations of the R&D knowledge stock scaling parameter (χ), the knowledge stock volatility (σ_r), and the R&D investment effectiveness parameter (ψ_0). First, in terms of quantitative implications, the most significant impacts come from variation in χ and ψ_0, whereas changes in σ_r have a trivial impact. As such, we do note report results for changes in σ_r, but they are available upon request. Importantly, the trivial impact of shifting the knowledge stock volatility value further highlights the significance of the technology jump in our model, and the role of uncertainty about this jump.

ψ_0	χ	emissions	R&D-output ratio	median jump yr.
0.07	500	9.35	0.0043	27.49
0.07	750	9.20	0.0070	31.72
0.07	1000	9.10	0.0091	34.13
0.11	500	9.45	0.0040	25.89
0.11	750	9.32	0.0062	30.05
0.11	1000	9.24	0.0081	32.47
0.15	500	9.51	0.0036	24.71
0.15	750	9.41	0.0056	28.69
0.15	1000	9.33	0.0072	30.98

(A) Simulation outputs under baseline.

ψ_0	χ	emissions	R&D-output ratio	median jump yr.
0.07	500	8.95	0.0139	32.82
0.07	750	8.62	0.0247	37.23
0.07	1000	8.37	0.0330	39.64
0.11	500	9.26	0.0078	28.76
0.11	750	9.04	0.0138	33.09
0.11	1000	8.86	0.0192	35.59
0.15	500	9.40	0.0057	26.58
0.15	750	9.23	0.0096	30.69
0.15	1000	9.10	0.0132	33.18

(B) Simulation outputs under less aversion.

ψ_0	χ	emissions	R&D-output ratio	median jump yr.
0.07	500	8.24	0.04070	36.71
0.07	750	7.82	0.0532	38.17
0.07	1000	7.60	0.0536	38.10
0.11	500	8.92	0.0180	31.98
0.11	750	8.53	0.0316	35.70
0.11	1000	8.26	0.0401	37.07
0.15	500	9.22	0.0102	28.59
0.15	750	8.94	0.0185	32.77
0.15	1000	8.72	0.0254	34.99

(C) Simulation outputs under more aversion.

Table 21: Simulation outputs for the full model for varying R&D knowledge stock parameter configurations. Column 1 shows the value of ψ_0 used. Column 2 shows the value of χ used. Column 3 shows the emissions value. Column 4 shows the R&D investment-to-output ratio value. Column 5 shows the median year of a technology jump. Rows 1 through 3 assume a low R&D investment efficiency parameter. Rows 4 through 6 assume a medium R&D investment efficiency parameter. Rows 7 through 9 assume a high R&D investment efficiency parameter. Rows 1, 4, and 7 assume a small knowledge stock (inverse) scaling. Rows 2, 5, and 6 assume a medium knowledge stock (inverse) scaling. Rows 3, 6, and 9 assume a large knowledge stock (inverse) scaling. Panel (A) shows outcomes for the baseline uncertainty neutral case. Panel (B) shows outcomes for the less uncertainty aversion case. Panel (C) shows outcomes for the more uncertainty aversion case.

Table 21 shows the 9 variations of the R&D related parameter values we explore: $\psi_0 \in \{0.07, 0.11, 0.15\}$ and $\psi_0 \in \{500, 750, 1000\}$. First, we can see that increasing ψ_0, which improves the efficiency of R&D investment, leads to increased R&D investment which lowers the median year of a technology jump occurring. Also, increasing χ, which decreases the likelihood arrival rate of a technology jump occurrence for a given level of the knowledge stock, leads to increased R&D investment as well but with the median year of a technology jump occurring being higher. The central takeaway when comparing outcomes across uncertainty aversion scenarios is that the qualitative results are essentially identical to the results reported in the main text across all R&D knowledge stock parameters combinations we explore. As mentioned before, Online Appendix C shows the numerical outcomes for changing the R&D knowledge stock related parameters χ and ψ_0 across different choices of ϕ_0. Again, the qualitative implications for the impact of uncertainty aversion persist in these alternative parameter specifications.

F Comparison to a real options formulation

As we noted in Section 11, there are similarities, but also differences, in our R&D investment problem and a real options problem.

As a reminder, in standard option pricing, payoff volatility is valued because the investor can exercise the option when the payoff is sufficiently positive. Increases in volatility enhance the right tail of the payoff distribution. The experiment that we explore, however, changes uncertainty aversion and does not increase the right tail probability. Given the ability to hedge, there is no direct role for risk aversion in the familiar options pricing formula.

Since our social planner does not have the ability to replicate the uncertainty of the R&D investment, we find that the pedagogically revealing real options analysis of Miao and Wang (2007) provides interesting benchmarks for comparison. Two of the specifications they consider essentially solve single-agent decision problems, and standard replication arguments do not apply. While they explore risk aversion with exponential utility, we obtain analogous conclusions with very similar supporting computations starting with linear utility and incorporating robustness concerns. We use their functional forms for pedagogical simplicity.

We compare two specifications of the payoff to exercising the option. In both cases, there is uncertainty in the payoff dynamics. The decision-maker decides when to exercise the option with a known cost of doing so. Other than access to this stochastic payoff, the decision-maker can invest in a riskless technology. In the first specification, exercising the option provides a discrete payoff. In the second specification, exercising the option entitles the decision maker with access to a stochastic flow. To support our investigation, we make the two payoffs comparable. Specifically, we scale the flow so that the present values of the discrete payoff and flow payoff process are the same in the absence of uncertainty about the payoff dynamics. In the case of the discrete payoff, the uncertainty concerns are resolved when the option is exercised, in contrast to the case with a flow payoff as the uncertainty concerns about the transient dynamic persist.

As we will now show, uncertainty has obverse impacts for the two payoffs. This aversion makes the investor more bullish on the investment for the discrete payoff specification in the sense that there is a lower investment threshold. The outcome gets reversed with the stochastic flow payoff and the investor is less bullish on the investment and exercises the option at a higher threshold with enhanced aversion.

Consider a robust counterpart to Model I in Miao and Wang (2007). We refer to this as the flow-payoff

model. We let the utility be linear in consumption and the income process be:

$$dY_t = \alpha dt + \sigma dW_t.$$

For purposes of illustration, we set $\alpha = .1$ and $\sigma = .3$.[30] When the option is exercised, we assume a payoff equal to the discounted expected value minus a prespecified cost I. In our illustration, we set $I = 10$. We use the discounted expected value for comparability to a second model that we consider. The risk-free rate and subjective rate of discount are the same, which we set as .02. We refer to this as the *discrete-payoff model*.

The post exercise value function satisfies:

$$\max_c -\delta V^0 + \delta c + V_w^0(rw - c) = 0,$$

which can be solved by setting $V^0(w) = \delta w$. The value function, V, prior to exercising the option solves the HJB equation:

$$\max_c \min_h -\delta V + \delta c + V_w(\delta w - c) + V_y \alpha + \sigma h V_y + \frac{\xi}{2} h^2 + \frac{\sigma^2}{2} V_{yy}.$$

Guess that $V(w, y) = \delta[w + G(y)]$. Then

$$0 = -\delta G + G_y \alpha + \frac{\sigma^2}{2} G_{yy} - \frac{\sigma^2 \delta}{2\xi} (G_y)^2 \tag{22}$$

where the minimizing h solves:

$$h^* = -\frac{\sigma}{\xi} V_y = -\frac{\sigma \delta}{\xi} G_y. \tag{23}$$

This equation for G essentially matches Miao and Wang (2007)'s G function in their Proposition 1. The discounted expected value with income realization, y, is

$$\frac{1}{\delta} y + \frac{\alpha}{\delta^2}$$

[30] Miao and Wang (2007) use these numbers in the illustration they report in their Figure 4.

obtained by solving the Feynman-Kac equation:

$$0 = -\delta F + y + \alpha F_y + \frac{\sigma^2}{2} F_{yy}.$$

Thus, value matching gives:

$$\delta G(\bar{y}) = \bar{y} + \frac{\alpha}{\delta} - \delta I,$$

and smooth pasting gives

$$\delta G_y(\bar{y}) = 1.$$

We next change the payoff on the investment by letting it be continuous given by the Y process. We refer to this as the *flow-payoff model.* The value function, V^o, after exercising satisfies the HJB equation:

$$\max_c \min_h -\delta V^o + \delta c + V^o_w(\delta w + y - c) + \alpha V^o_y + \sigma h V^o_y + \frac{1}{2}\sigma^2 V^o_{yy} + \frac{\xi}{2} h^2 = 0.$$

We guess that $V^o(w, y) = \delta[w + F(y)]$. With this guess, we solve

$$0 = -\delta F + y + F_y \alpha + \frac{\sigma^2}{2} F_{yy} - \frac{\sigma^2 \delta}{2\xi} (F_y)^2$$

where

$$h^* = -\frac{\sigma}{\xi} V^o_y = -\frac{\delta \sigma}{\xi} F_y.$$

The solution for F is:

$$F(y) = \frac{1}{\delta}\left[y + \frac{\alpha}{\delta} - \frac{\sigma^2}{2\xi\delta}\right],$$

implying that $h^* = -\sigma/\xi$.

As with the discrete-payoff specification, the value function prior to the option being exercised solves (22) with $V(w, y) = \delta[w + G(y)]$. The minimizing h^* is given by formula (23). The value-matching condition is:

$$\delta G(\bar{y}) = \delta F(\bar{y} - I) = \bar{y} + \frac{\alpha}{\delta} - \frac{\sigma^2}{2\xi\delta} - \delta I$$

and the smooth-pasting condition is:

$$\delta G_y(\bar{y}) = \delta F_y(\bar{y}) = 1.$$

Notice that the value matching condition is the same for both models in the limiting case in which $\xi \to \infty$. In both models, the exercise threshold $\bar{y}$ satisfies: $\delta G_y(\bar{y}) = 1$ even for $\xi < \infty$, but, of course, the G functions are different. Figure 11 plots the derivatives for each model specification and for alternative values of the robustness parameter, ξ. Recall that larger values of the penalty parameter imply less aversion. For the discrete-payoff model, the derivatives increase with ξ, resulting in lower values of y when the option is exercised. For the stochastic flow-payoff model the derivatives decrease with ξ, resulting in higher values of the threshold $\bar{y}$.

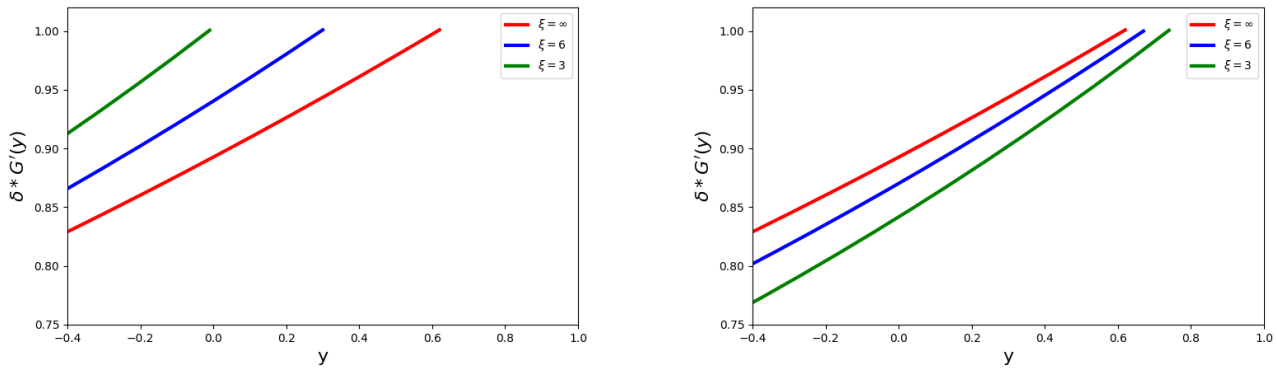

(A) discrete-payoff model (B) flow-payoff model

Figure 11: Plots of δG_y for different values of ξ.

Consisitent with this finding, we report in Figure 12 uncertainty aversion has opposite implications for the threshold $\bar{y}$ for the two payoff specifications. This figure reports the investment thresholds for alternative values of the penalty parameter, ξ. The implied uncertainty aversion for low values of ξ aversion makes the investor more bullish on the investment for the discrete payoff specification in the sense that there is a lower investment threshold. The outcome gets reversed with the stochastic flow payoff: the investor is less bullish on the investment and exercises the option at a higher threshold with enhanced aversion. By design, the two curves converge when $\xi \to \infty$.[31]

[31] These results have counterparts to the risk aversion investigation reported in Miao and Wang (2007), albeit without the explicit present-value link between the two payoffs.

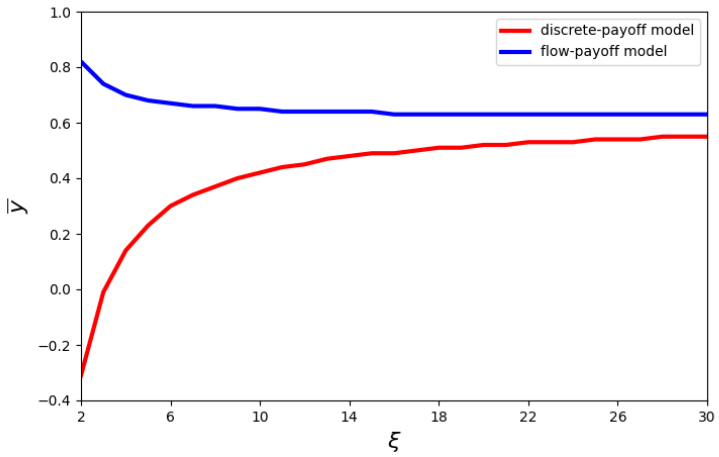

Figure 12: Solution of exercise threshold $\overline{y}$ for each model as functions of the robustness parameter, ξ.

We plot the pre-exercise drift distortions h^* in Figure 13. The post-jump distortions are zero for the discrete-payoff specification, and, by the smooth pasting condition, equal to the distortion at the time the option is exercised for the flow-payoff specification. Recall that the h^*'s are expressed as drifts in a standard Brownian increment. They are not extreme for these illustrations, as the magnitude of the local mean distortion never exceeds ten percent of the standard deviation. The drift distortions prior to exercise are a bit larger in magnitude for the discrete-payoff specification.

Why the difference? In both cases, robustness concerns induce more cautious assessments of the payoff prospects. In the discrete-payoff specification, the uncertainty is resolved as soon as the payoff is exercised, increasing the valuation; while in the flow-payoff specification, the post exercise valuation is diminished under robustness concerns because of the continued uncertainty about future payoffs.

Although there are intriguing similarities between this comparison of real options in different environments and our findings, the investment opportunities are quite different. Rather than an exercise decision, our social planner can make investments that increase the likelihood of a discovery. Uncertainty aversion makes the planner skeptical about the payoff horizon in our analysis. These differences make our uncertainty aversion impact not monotone as displayed in Table 7 and in contrast to the real options findings reported in Figure 12.

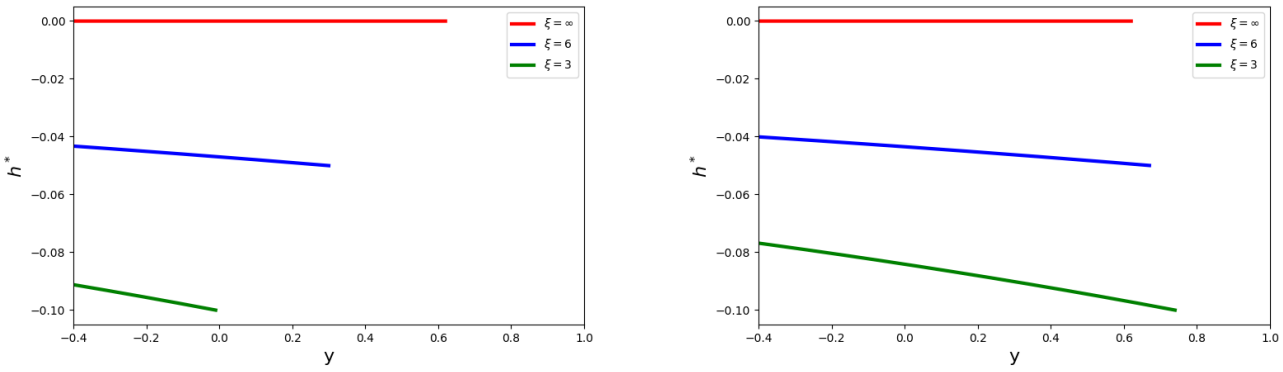

(A) discrete-payoff model

(B) flow-payoff model

Figure 13: Plots of h^* for different values of ξ.

G Computational method

We provide an outline of the computational algorithm used to solve the HJB equations in what follows, and direct the reader to our online appendix for full details.

G.1 Policy iteration

For simplicity, we denote the state space, control set, and distortion set by:

$$X = \{\log k, y, \log r\}, \qquad \Phi = \{i^k, i^j, \mathcal{E}\}, \qquad \Gamma = \{h^k, h^y, h^r, g^\ell\}.$$

The value function $\hat{v}(X)$ depends on the state X.

A pseudo-code of the numerical algorithm for solving the HJB equation via policy iteration is given in Algorithm 1. The algorithm involves an iterative process that alternates between policy improvement steps and policy evaluation steps. The subsequent sections then explain further details for each step in the algorithm.

Algorithm 1 Solving the HJB equation via policy iteration.

Input: Initial guess for value function $\hat{v}^0$, $\epsilon = 10^{-7}$

Output: Optimal value function $\hat{v}^*$

$\hat{v} \leftarrow \hat{v}^0$

Repeat

Step 1: Compute optimal actions Φ^* from the first-order condition of the maximization problem using $\hat{v}$, Φ, Γ

Step 2: Compute optimal probability distortions Γ^* from the first-order condition of the minimization problem using $\hat{v}$, Φ^*, Γ

Step 3: Update the value function $\hat{v}^*$ by solving (conditionally) the HJB PDE with a false transient method. $\hat{v}$, Φ^*, Γ^* are used in the PDE.

Step 4: $\hat{v} \leftarrow \hat{v}^*$ $\quad \Phi \leftarrow \Phi^*$ $\quad \Gamma \leftarrow \Gamma^*$

Until $|\hat{v}^* - \hat{v}| < \epsilon$

G.2 Updating rules for Φ and Γ at policy improvement steps

In solving HJB equations, we often encounter complex, highly non-linear equations for our optimal control choices that do not admit analytical solutions. To address this challenge, we implement an iterative numerical specification which we call a "cobweb" algorithm to approximate the optimal control variables. The cobweb algorithm works as follows:

1. Start with an initial guess for the control variable.
2. Compute the corresponding values in the equations.
3. Update the control variable based on the discrepancies observed.
4. Repeat the process until the control variable converges to a stable value.

This cobweb algorithm is an interior loop within our iterative algorithm 1. The tolerance criteria for the cobweb algorithm is determined by testing the trade-off between necessary accuracy for stability and speed of deriving solutions.

We note that every probability distortion from misspecification aversion has an analytical expression as a function of the value function and it's derivatives. These can be computed and plugged in directly to

the HJB expressions, and the iterative process for the HJB equation solution updates these values as the value function is updated until convergence is achieved.

G.3 Solving the linear PDE equation at the policy evaluation step

Updating value functions, given the state variables, controls, and distortions, is done by solving the linear PDE implicitly. To mitigate the potential instability of the non-linear HJB, we add a false transient (time) dimension and iterate the solution with an implicit time stepping procedure until the false time derivative converges to zero.In addition, we solve the linear system using a stabilized Bi-Conjugate Gradient solver with the ILU preconditioner provided by PETSc, a widely used software library developed by researchers at Argonne National laboratory to efficiently solve sparse linear/nonlinear systems[32]. We outline how we construct the linear system from the HJB equation in what follows.

First, note that we can write the HJB equation from our model using the following form:

$$A\hat{v} + B_{\hat{k}}\frac{\partial \hat{v}}{\partial \hat{k}} + B_y\frac{\partial \hat{v}}{\partial y} + B_{\hat{r}}\frac{\partial \hat{v}}{\partial \hat{r}} + C_{\hat{k}}\frac{\partial^2 \hat{v}}{\partial \hat{k}\,\partial \hat{k}'} + C_y\frac{\partial^2 \hat{v}}{\partial y\,\partial y'} + C_{\hat{r}}\frac{\partial^2 \hat{v}}{\partial \hat{r}\,\partial \hat{r}'} + D = 0.$$

More generally, our system can be expressed in the conditionally linear form:

$$\begin{aligned} 0 \;=\; & V_t(x) + \mathbb{A}(x; V, V_x, V_{xx})V(x) + \mathbb{B}(x; V, V_x, V_{xx})V_x(x) \\ & + \frac{1}{2}tr[\mathbb{C}(x; V, V_x, V_{xx})V_{xx}(x)\mathbb{C}(x; V, V_x, V_{xx})] + \mathbb{D}(x; V, V_x, V_{xx}) \end{aligned}$$

where x is our state variable vector, $V_x(x) = \frac{\partial V}{\partial x}(x)$ and $V_{xx}(x) = \frac{\partial^2 V}{\partial x \partial x'}(x)$ are used for notational simplicity. The solution is obtained by finding a $V(x)$ such that the HJB equality, first-order condition, and $V_t(x) = 0$ hold. Conditional on an initial guess, $V^0(x)$, we calculate the coefficients $\mathbb{A}, \mathbb{B}, \mathbb{C}$ and $\mathbb{D}$ based on the flow utility and the evolution of the state variables. We substitute the calculated coefficients, rearrange the terms, and apply a backward finite difference to $V_t^0(x)$ to obtain the final expression

$$\begin{aligned} \hat{V}(x) \;=\; & V^0(x) + [\mathbb{A}(x, V^0, V_x^0, V_{xx}^0)V^0(x) + \mathbb{B}(x, V^0, V_x^0, V_{xx}^0)V_x^0(x) \\ & + \frac{1}{2}tr[\mathbb{C}(x, V^0, V_x^0, V_{xx}^0)V_{xx}^0(x)\mathbb{C}(x, V^0, V_x^0, V_{xx}^0)] + \mathbb{D}(x, V^0, V_x^0, V_{xx}^0)]\Delta t. \end{aligned}$$

[32] We particularly thank the research professionals from UChicago's Macro-Finance Research Program - Bin Cheng, Pengyu Chen, and Zhaoyang Xu - and our coauthor Hong Zhang, who worked together to integrate the PETSc software (see Balay et al. (2025)) into our computational algorithm.

We solve numerically for $\hat{V}(x)$ from this linear system using the Preconditioned Bi-Conjugate Gradient Stabilized method from van der Vorst (1992).

G.4 Finite-differences scheme

The finite difference scheme used in our computations is described as follows. At interior points in the state space, we apply central differences for both first- and second-order derivatives. At boundary points corresponding to the minimum state values, we use forward differences, while at the maximum boundary values, we apply backward differences again for both first- and second-order derivatives.

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 M. Halem, Irene Hernandez, Andrew W. Lo, Manish Singh, and Dennis Whyte. 2022. Financing Fusion Energy. Technical Report Revision 13 December 2022, MIT Sloan School of Management.

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

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.

Cerreia-Vioglio, Simone, Lars Peter Hansen, Fabio Maccheroni, and Massimo Marinacci. 2025. Making Decisions under Model Misspecification. *Review of Economic Studies* forthcoming.

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* .

Cogley, Timothy, Riccardo Colacito, Lars Peter Hansen, and Thomas J. Sargent. 2008. Robustness and U.S. Monetary Policy Experimentation. *Journal of Money, Credit and Banking* 40 (8):1599–1623.

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.

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

Joos, F., R. Roth, J. S. Fuglestvedt, G. P. Peters, I. G. Enting, W. Von Bloh, V. Brovkin, E. J. Burke, M. Eby, N. R. Edwards, T. Friedrich, T. L. Frölicher, P. R. Halloran, P. B. Holden, C. Jones, T. Kleinen, F. T. Mackenzie, K. Matsumoto, M. Meinshausen, G. K. Plattner, A. Reisinger, J. Segschneider, G. Shaffer, M. Steinacher, K. Strassmann, K. Tanaka, A. Timmermann, and A. J. Weaver. 2013. Carbon Dioxide and Climate Impulse Response Functions for the Computation of Greenhouse Gas Metrics: A Multi-Model Analysis. *Atmospheric Chemistry and Physics* 13 (5):2793–2825.

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.

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):1–8.

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.

Sargent, Thomas. 1999. *Policy Rules for Open Economies (Discussion of Laurence Ball)*, chap. 3, 141–154. University of Chicago 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.

Online Appendices

Below we provide additional details and derivations for various results, which will be made available online in our online notebook:

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

A ξ_a version

A.1 corresponding to $\xi_m = \infty$

	$\frac{\partial V}{\partial \hat{r}}$	$\frac{\partial V}{\partial y}$	I^r ratio	I^k ratio
$\rho = 0.75$	0.0252	-0.0314	0.0126	0.8803
$\rho = 1$	0.0221	-0.0293	0.0062	0.7743
$\rho = 1.33$	0.0181	-0.0264	0.0033	0.7029

A.2 corresponding to $\xi_m = 0.1$

	$\frac{\partial V}{\partial \hat{r}}$	$\frac{\partial V}{\partial y}$	I^r ratio	I^k ratio
$\rho = 0.75$	0.0371	-0.0551	0.0261	0.8645
$\rho = 1$	0.0334	-0.0506	0.0143	0.7657
$\rho = 1.33$	0.0275	-0.0436	0.0077	0.6992

A.3 corresponding to $\xi_m = 0.05$

	$\frac{\partial V}{\partial \hat{r}}$	$\frac{\partial V}{\partial y}$	I^r ratio	I^k ratio
$\rho = 0.75$	0.0563	-0.1120	0.0555	0.8342
$\rho = 1$	0.0510	-0.1004	0.0335	0.7460
$\rho = 1.33$	0.0422	-0.0829	0.0190	0.6881

B ξ_m version

B.1 $\xi_m = \infty$

	$\frac{\partial V}{\partial \hat{r}}$	$\frac{\partial V}{\partial y}$	I^r ratio	I^k ratio
$\rho = 0.75$	0.0252	-0.0314	0.0126	0.8803
$\rho = 1$	0.0221	-0.0293	0.0062	0.7743
$\rho = 1.33$	0.0181	-0.0264	0.0033	0.7029

B.2 $\xi_m = 0.1$

	$\frac{\partial V}{\partial \hat{r}}$	$\frac{\partial V}{\partial y}$	I^r ratio	I^k ratio
$\rho = 0.75$	0.0371	-0.0551	0.0261	0.8645
$\rho = 1$	0.0334	-0.0506	0.0143	0.7657
$\rho = 1.33$	0.0275	-0.0436	0.0077	0.6992

B.3 $\xi_m = 0.05$

	$\frac{\partial V}{\partial \hat{r}}$	$\frac{\partial V}{\partial y}$	I^r ratio	I^k ratio
$\rho = 0.75$	0.0564	-0.1122	0.0557	0.8341
$\rho = 1$	0.0510	-0.1004	0.0335	0.7460
$\rho = 1.33$	0.0423	-0.0831	0.0191	0.6880

C sensitivity analysis (ϕ_0, ψ_0, χ)

C.1 Full model

C.1.1 corresponding to $\xi_m = \infty$

Table 22: Simulation outputs under full model for fixed $\xi_k = \infty$ (total rows = 27).

ϕ_0	ψ_0	χ	e_0	$I_0^r/(\alpha K(1-\phi_0(A)^{\phi_1}))$	$I_0^r/(\alpha K)$	Median year of tech jump
0.1	0.07	500	8.544791	0.00190043	0.00189954	29.690
0.1	0.07	750	8.319915	0.00287562	0.00287368	35.033
0.1	0.07	1000	8.165679	0.00353878	0.00353577	37.999
0.1	0.11	500	8.675003	0.00210789	0.00210712	27.940
0.1	0.11	750	8.474502	0.00314220	0.00314055	32.989
0.1	0.11	1000	8.331799	0.00388442	0.00388184	35.879
0.1	0.15	500	8.770612	0.00213869	0.00213803	26.473
0.1	0.15	750	8.590094	0.00314622	0.00314487	31.361
0.1	0.15	1000	8.458318	0.00389123	0.00388913	34.249
0.3	0.07	500	9.127565	0.00356786	0.00356642	28.062
0.3	0.07	750	8.954474	0.00569124	0.00568773	32.564
0.3	0.07	1000	8.831682	0.00733472	0.00732879	35.117
0.3	0.11	500	9.242916	0.00342008	0.00341908	26.431
0.3	0.11	750	9.095188	0.00532738	0.00532505	30.721
0.3	0.11	1000	8.987540	0.00683624	0.00683234	33.157
0.3	0.15	500	9.320805	0.00319191	0.00319118	25.119
0.3	0.15	750	9.191028	0.00488346	0.00488181	29.292
0.3	0.15	1000	9.094596	0.00623912	0.00623639	31.749
0.5	0.07	500	9.346626	0.00430257	0.00430105	27.489
0.5	0.07	750	9.202543	0.00695191	0.00694811	31.720
0.5	0.07	1000	9.099786	0.00906155	0.00905501	34.127
0.5	0.11	500	9.446284	0.00395478	0.00395380	25.891
0.5	0.11	750	9.324384	0.00623008	0.00622772	30.052
0.5	0.11	1000	9.235200	0.00807132	0.00806732	32.470
0.5	0.15	500	9.512329	0.00360305	0.00360236	24.712
0.5	0.15	750	9.405777	0.00557013	0.00556853	28.691

C.1.2 corresponding to $\xi_m = 0.1$

Table 23: Simulation outputs under full model for fixed $\xi_k = 0.1$ (total rows = 27).

ϕ_0	ψ_0	χ	e_0	$I_0^r/(\alpha K(1-\phi_0(A)^{\phi_1}))$	$I_0^r/(\alpha K)$	Median year of tech jump
0.1	0.07	500	8.277944	0.00316531	0.00316302	34.005
0.1	0.07	750	7.959286	0.00480987	0.00480444	39.593
0.1	0.07	1000	7.745289	0.00578310	0.00577459	42.651
0.1	0.11	500	8.513817	0.00294117	0.00293972	30.664
0.1	0.11	750	8.245921	0.00449523	0.00449183	36.044
0.1	0.11	1000	8.052351	0.00559653	0.00559095	39.106
0.1	0.15	500	8.661668	0.00271605	0.00271502	28.430
0.1	0.15	750	8.431484	0.00410777	0.00410545	33.615
0.1	0.15	1000	8.259595	0.00515543	0.00515161	36.640
0.3	0.07	500	8.750300	0.00918965	0.00918086	33.036
0.3	0.07	750	8.405408	0.01567337	0.01564560	37.833
0.3	0.07	1000	8.151270	0.02048339	0.02043006	40.629
0.3	0.11	500	9.052710	0.00603669	0.00603374	29.109
0.3	0.11	750	8.807751	0.01016989	0.01016125	33.996
0.3	0.11	1000	8.615968	0.01370903	0.01369213	36.536
0.3	0.15	500	9.204302	0.00470462	0.00470309	26.995
0.3	0.15	750	9.012328	0.00769669	0.00769255	31.413
0.3	0.15	1000	8.860020	0.01031389	0.01030604	34.076
0.5	0.07	500	8.952408	0.01388577	0.01387138	32.820
0.5	0.07	750	8.620097	0.02466227	0.02461192	37.227
0.5	0.07	1000	8.372014	0.03298653	0.03288381	39.637
0.5	0.11	500	9.259758	0.00784915	0.00784552	28.756
0.5	0.11	750	9.036770	0.01377334	0.01376168	33.088
0.5	0.11	1000	8.857447	0.01916032	0.01913588	35.587
0.5	0.15	500	9.402066	0.00570290	0.00570123	26.576
0.5	0.15	750	9.233665	0.00963049	0.00962568	30.686
0.5	0.15	1000	9.097085	0.01324278	0.01323316	33.184

C.1.3 corresponding to $\xi_m = 0.05$

Table 24: Simulation outputs under full model for fixed $\xi_k = 0.05$ (total rows = 27).

ϕ_0	ψ_0	χ	e_0	$I_0^r/(\alpha K(1-\phi_0(A)^{\phi_1}))$	$I_0^r/(\alpha K)$	Median year of tech jump
0.1	0.07	500	7.906052	0.00529371	0.00528731	38.613
0.1	0.07	750	7.504891	0.00742563	0.00741125	43.788
0.1	0.07	1000	7.271944	0.00801720	0.00799741	46.117
0.1	0.11	500	8.299302	0.00419822	0.00419529	33.621
0.1	0.11	750	7.948408	0.00643042	0.00642306	39.135
0.1	0.11	1000	7.701533	0.00784195	0.00782980	42.153
0.1	0.15	500	8.522988	0.00351626	0.00351455	30.527
0.1	0.15	750	8.229392	0.00543074	0.00542653	35.875
0.1	0.15	1000	8.008477	0.00684695	0.00683972	39.072
0.3	0.07	500	8.096348	0.02331330	0.02324766	37.864
0.3	0.07	750	7.626086	0.03217178	0.03200856	40.706
0.3	0.07	1000	7.366942	0.03358916	0.03336338	41.578
0.3	0.11	500	8.736463	0.01172224	0.01171071	32.429
0.3	0.11	750	8.341442	0.02012263	0.02008316	36.786
0.3	0.11	1000	8.045163	0.02622347	0.02614436	38.957
0.3	0.15	500	9.027171	0.00746962	0.00746574	29.135
0.3	0.15	750	8.734585	0.01298813	0.01297533	33.603
0.3	0.15	1000	8.496258	0.01783395	0.01780678	36.214
0.5	0.07	500	8.239226	0.04070448	0.04054891	36.714
0.5	0.07	750	7.816186	0.05319197	0.05283267	38.170
0.5	0.07	1000	7.604366	0.05364911	0.05318412	38.097
0.5	0.11	500	8.915995	0.01800642	0.01798616	31.979
0.5	0.11	750	8.532106	0.03156315	0.03148777	35.697
0.5	0.11	1000	8.264544	0.04005090	0.03990365	37.070
0.5	0.15	500	9.218543	0.01015119	0.01014588	28.585
0.5	0.15	750	8.940857	0.01845026	0.01843067	32.771
0.5	0.15	1000	8.719903	0.02544468	0.02540168	34.993

D sensitivity analysis δ

D.1 Full model

D.1.1 corresponding to $\xi_m = \infty$

Subjective discount rate	$\frac{\partial V}{\partial \hat{r}}$	$\frac{\partial V}{\partial y}$
$\delta = .010$	0.022146	-0.02932477
$\delta = .015$	0.01827176	-0.02647366
$\delta = .020$	0.01406391	-0.0231746

Table 25: Social values at the initial time period for **baseline** to misspecification uncertainty.

D.1.2 corresponding to $\xi_m = 0.1$

Subjective discount rate	$\frac{\partial V}{\partial \hat{r}}$	$\frac{\partial V}{\partial y}$
$\delta = .010$	0.03344168	-0.05065515
$\delta = .015$	0.02746952	-0.04343948
$\delta = .020$	0.02055801	-0.03520589

Table 26: Social values at the initial time period for **less aversion** to misspecification uncertainty.

D.1.3 corresponding to $\xi_m = 0.05$

Subjective discount rate	$\frac{\partial V}{\partial \hat{r}}$	$\frac{\partial V}{\partial y}$
$\delta = .010$	0.05094414	-0.10018285
$\delta = .015$	0.04145785	-0.08099132
$\delta = .020$	0.03001312	-0.05968805

Table 27: Social values at the initial time period for **more aversion** to misspecification uncertainty.

E HJB Equations with 2 Technology Steps

There are L possible jump realizations from the pre-jump setting: $L-1$ potential damage curve realizations and one intermediate technology jump realization. From the intermediate-technology, pre-damage setting there L possible jump realizations: $L-1$ potential damage curve realizations and one terminal technology jump realization. From the pre-technology, post-damage setting there is only one possible intermediate technology jump realization. From the intermediate-technology, post-damage setting there is only one possible terminal technology jump realization. The $L-1$ damage curve realizations in each of the pre-damage jump states are mutually exclusive. Below we compute the continuation values for each of the different post-jump scenarios.

E.1 Post-technology, post-damage jump HJB equations

Post-technology jump, emissions can be reduced to zero without any impact to output, and thus full abatement is optimal from the planner's perspective. As a result, temperature remains constant. The relevant state space, control set, and distortion set simplify to

$$x = \{y, \hat{k}\}, \qquad \Phi = \{i^k\}, \qquad \Gamma = \{h^k\}$$

To compute $V^{\ell,L}$ for $\ell = 1, ..., L-1$ conditioned on both a technology jump and a damage jump occurring, we simply need to solve the HJB equation

$$\begin{aligned} 0 = & \max_{i^k} \min_{h^k} \delta \log(\alpha - i^k) + \delta \hat{k} - \delta \hat{n}(y) - \delta V^{\ell,L} + \xi^k \frac{|h^k|^2}{2} \\ & + \frac{\partial V^{\ell,L}}{\partial \hat{k}} \left[\mu_k + i^k - \frac{\kappa}{2} \left(i^k \right)^2 - \frac{|\sigma_k|^2}{2} + \sigma_k h^k \right] + \frac{\partial^2 V^{\ell,L}}{\partial \hat{k} \, \partial \hat{k}'} \frac{|\sigma_k|^2}{2} \end{aligned}$$

This continuation value has a closed-form solution that we guess and verify to be of the form

$$V^{\ell,L} = \log(\alpha - i^k) - \frac{1}{\delta 2 \xi^k} |\sigma_k|^2 + \frac{1}{\delta} \left[\mu_k + i^k - \frac{\kappa}{2} \left(i^k \right)^2 - \frac{|\sigma_k|^2}{2} \right] + \hat{k} - \hat{n}(y)$$

with first order conditions given by

$$\left(h^k\right)^* = -\frac{1}{\xi^k}\sigma_k$$
$$\left(i^k\right)^* = \frac{(1+\alpha\kappa) - \sqrt{(1+\alpha\kappa)^2 - 4\kappa(\alpha-\delta)}}{2\kappa}$$

E.2 Post-technology, pre-damage jump HJB equation

The value function assuming that only a technology jump has been realized, V^L, incorporates the possibility of jumping to one of L-1 possible damage states. However, as in the post-technology, post-damage jumps case, emissions can be reduced to zero without any impact to output, and thus full abatement is optimal from the planner's perspective. As a result, temperature remains constant and the realization of a damage curve realization is inconsequential because we remain at the point where any incremental curvature has no impact on the damages. Therefore, we can ignore the intensities of the damage curve and the associated continuation values for this computation. The result is that the state variable, control set and distortion set, as well as the value function solution and first order conditions, are essentially identical to the post-technology, post-damage jumps case.

E.3 Intermediate-technology, post-damage jump HJB equation

Now we compute each of the post-damage intermediate-technology jump value functions V^ℓ where only a damage jump has been realized for $\ell = 1, ..., L-1$. Note that the damage function in this setting depends on the temperature at which the damage jump occurred, $\hat{y}$. To simplify the post damage-jump state space used for our post damage-jump value function, we define a new state variable $z = y - \hat{y} + \overline{y}$. This state variable becomes the counterpart of the temperature anomaly state variable y. The dynamics of z are the same as that of our previous temperature state variable, although z is initialized at $\overline{y}$ at the time of the jump regardless of the current temperature value. Using the new state variable, we can specify an adjusted

damage function as follows

$$\begin{aligned} m(z) &= \hat{n}(y) - (\lambda_1 \hat{y} + \frac{1}{2}\lambda_2 \hat{y}^2) \\ &= \lambda_1(y - \hat{y}) + \frac{1}{2}\lambda_2(y - \hat{y} + \overline{y})^2 + \frac{1}{2}\lambda_3(\ell)(y - \hat{y})^2 - \frac{1}{2}\lambda_2(\overline{y})^2 \\ &= \lambda_1(z - \overline{y}) + \frac{1}{2}\lambda_2 z^2 + \frac{1}{2}\lambda_3(\ell)(z - \overline{y})^2 - \frac{1}{2}\lambda_2(\overline{y})^2 \end{aligned}$$

Our value functions depend on a modified state vector X, the control set Φ, and the distortion set Γ:

$$x = \{\hat{k}, z, \hat{r}\}, \qquad \Phi = \{i^k, i^r, e\}, \qquad \Gamma = \{h^k, h^z, h^r, g^L\}$$

With this new state variable and adjusted damage function, we can now compute the simplified HJB equation $W^{\ell,L'}(\hat{k}, z, \hat{r})$, which has the following relationship with $v^{\ell,L'}(\hat{k}, y, \hat{r})$:

$$W^{\ell,L'}(\hat{k}, z, \hat{r}) = V^{\ell,L'}(\hat{k}, y, \hat{r}) + (\lambda_1 \hat{y} + \frac{1}{2}\lambda_2 \hat{y}^2)$$

There is only one possible jump state, which is the realization of the terminal of technological innovation with intensity $\mathcal{J}^L$. The post-technology jump continuation value is $W^{\ell,L}$, and the intermediate-technology jump continuation value is the solution to the following HJB equation:

$$\begin{aligned} 0 \;=\; & \max_{i^k, i^r, e} \min_{h, q(\theta), g^{L'}} \delta \log\left[\alpha - i^k - i^r - \alpha\phi_0' \left(1 - \frac{e}{\beta\alpha k}\right)^{\phi_1}\right] + \delta\hat{k} - \delta m(z) - \delta W^{\ell,L'}(\hat{k}, z, \hat{r}) \\ & + \frac{\partial W^{\ell,L'}}{\partial \hat{k}}(\hat{k}, z, \hat{r})\left(-\mu_k + i^k - \frac{\kappa}{2}\left(i^k\right)^2 - \frac{|\sigma_k|^2}{2} + \sigma_k h\right) + \frac{\partial^2 W^{\ell,L'}}{\partial \hat{k}\,\partial \hat{k}'}(\hat{k}, z, \hat{r})\frac{|\sigma_k|^2}{2} \\ & + \frac{\partial W^{\ell,L'}}{\partial y}(\hat{k}, z, \hat{r}) e \left(\sum_{\theta \in \Theta} \theta q(\theta)\right) + \frac{\partial^2 W^{\ell,L'}}{\partial y\,\partial y'}(\hat{k}, z, \hat{r})\frac{|\varsigma|^2}{2} e^2 \\ & + \frac{\partial W^{\ell,L'}}{\partial \hat{r}}(\hat{k}, z, \hat{r})\left(-\zeta + \psi_0 (i^r)^{\psi_1} \exp\left(-\psi_1\left(\hat{r} - \hat{k}\right)\right) - \frac{|\sigma_r|^2}{2} + \sigma_r h\right) + \frac{\partial^2 W^{\ell,L'}}{\partial \hat{r}\,\partial \hat{r}'}(\hat{k}, z, \hat{r})\frac{|\sigma_r|^2}{2} \\ & + \mathcal{J}^L g^L\left[W^{\ell,L}(\hat{k}, z, \hat{r}) - W^{\ell,L'}(\hat{k}, z, \hat{r})\right] + \xi\mathcal{J}^L\left[1 - g^L + g^L \log g^L\right] + \frac{\xi}{2} h' h + \sum_{\theta \in \Theta} q(\theta) \log q(\theta). \end{aligned}$$

where $W^{\ell,L}(\hat{k}, z, \hat{r}) = V^{\ell,L}(\hat{k}, y, \hat{r}) + (\lambda_1 \hat{y} + \frac{1}{2}\lambda_2 \hat{y}^2)$, which uses the transformation given previously.

E.4 Pre-technology, post-damage jump HJB equation

Now we compute each of the post-damage, pre-technology jump values functions V^ℓ where only a damage jump has been realized for $\ell = 1, ..., L-1$. This set-up is nearly identical to the intermediate-technology, post-damage setting except that the only jump realization is to the intermediate technology state with $\phi_0' = \frac{1}{2}\phi_0$. Note that the damage function in this setting depends on the temperature at which the damage jump occurred, $\hat{y}$. To simplify the post damage-jump state space used for our post damage-jump value function, we define a new state variable $z = y - \hat{y} + \overline{y}$. This state variables becomes the counterpart to the temperature anomaly state variable y. The dynamics of z are the same as that of our previous temperature state variable, although z is initialized at $\overline{y}$ at the time of the jump regardless of the current temperature value. Using the new state variable, we can specify an adjusted damage function as follows

$$\begin{aligned} m(z) &= \hat{n}(y) - (\lambda_1 \hat{y} + \frac{1}{2}\lambda_2 \hat{y}^2) \\ &= \lambda_1(y - \hat{y}) + \frac{1}{2}\lambda_2(y - \hat{y} + \overline{y})^2 + \frac{1}{2}\lambda_3(\ell)(y - \hat{y})^2 - \frac{1}{2}\lambda_2(\overline{y})^2 \\ &= \lambda_1(z - \overline{y}) + \frac{1}{2}\lambda_2 z^2 + \frac{1}{2}\lambda_3(\ell)(z - \overline{y})^2 - \frac{1}{2}\lambda_2(\overline{y})^2 \end{aligned}$$

Our value functions depend on a modified state vector X, the control set Φ, and the distortion set Γ:

$$x = \{\hat{k}, z, \hat{r}\}, \qquad \Phi = \{i^k, i^r, e\}, \qquad \Gamma = \{h^k, h^z, h^r, g^{L'}\}$$

With this new state variable and adjusted damage function, we can now compute the simplified HJB equation $W^\ell(\hat{k}, z, \hat{r})$, which has the following relationship with $v^\ell(\hat{k}, y, \hat{r})$:

$$W^\ell(\hat{k}, z, \hat{r}) = V^\ell(\hat{k}, y, \hat{r}) + (\lambda_1 \hat{y} + \frac{1}{2}\lambda_2 \hat{y}^2)$$

There is only one possible jump state, which is the realization of the technological innovation with intensity $\mathcal{J}^L$. The intermediate-technology jump continuation value is $W^{\ell,L'}$, and the pre-technology jump

continuation value is the solution to the following HJB equation:

$$
\begin{aligned}
0 \;=\; & \max_{i^k, i^r, e} \min_{h, q(\theta), g^L} \delta \log\left[\alpha - i^k - i^r - \alpha\phi_0\left(1 - \frac{e}{\beta\alpha k}\right)^{\phi_1}\right] + \delta\hat{k} - \delta m(z) - \delta W^\ell(\hat{k}, z, \hat{r}) \\
& + \frac{\partial W^\ell}{\partial \hat{k}}(\hat{k}, z, \hat{r})\left(-\mu_k + i^k - \frac{\kappa}{2}\left(i^k\right)^2 - \frac{|\sigma_k|^2}{2} + \sigma_k h\right) + \frac{\partial^2 W^\ell}{\partial \hat{k}\,\partial \hat{k}'}(\hat{k}, z, \hat{r})\frac{|\sigma_k|^2}{2} \\
& + \frac{\partial W^\ell}{\partial y}(\hat{k}, z, \hat{r}) e\left(\sum_{\theta \in \Theta} \theta q(\theta)\right) + \frac{\partial^2 W^\ell}{\partial y\,\partial y'}(\hat{k}, z, \hat{r})\frac{|\varsigma|^2}{2} e^2 \\
& + \frac{\partial W^\ell}{\partial \hat{r}}(\hat{k}, z, \hat{r})\left(-\zeta + \psi_0 (i^r)^{\psi_1} \exp\left(-\psi_1\left(\hat{r} - \hat{k}\right)\right) - \frac{|\sigma_r|^2}{2} + \sigma_r h\right) + \frac{\partial^2 W^\ell}{\partial \hat{r}\,\partial \hat{r}'}(\hat{k}, z, \hat{r})\frac{|\sigma_r|^2}{2} \\
& + \mathcal{J}^L g^{L'}\left[W^{\ell, L'}(\hat{k}, z, \hat{r}) - W^\ell(\hat{k}, z, \hat{r})\right] + \xi \mathcal{J}^L\left[1 - g^{L'} + g^{L'} \log g^{L'}\right] + \frac{\xi}{2} h' h + \sum_{\theta \in \Theta} q(\theta) \log q(\theta).
\end{aligned}
$$

where $W^{\ell, L'}(\hat{k}, z, \hat{r}) = V^{\ell, L'}(\hat{k}, y, \hat{r}) + (\lambda_1 \hat{y} + \frac{1}{2}\lambda_2 \hat{y}^2)$, which uses the transformation given previously.

E.5 Pre-technology, pre-damage jump HJB equation

The detailed, full-form pre-technology, pre-damage jump HJB equation is given by:

$$
\begin{aligned}
0 \;=\; & \max_{i^k, i^r, e} \min_{h, q(\theta), g^\ell} \delta \log\left[\alpha - i^k - i^r - \alpha\phi_0\left(1 - \frac{e}{\beta\alpha k}\right)^{\phi_1}\right] + \delta\hat{k} - \delta \hat{n}(y) - \delta V(\hat{k}, y, \hat{r}) \\
& + \frac{\partial V}{\partial \hat{k}}(\hat{k}, y, \hat{r})\left[-\mu_k + i^k - \frac{\kappa}{2}\left(i^k\right)^2 - \frac{|\sigma_k|^2}{2} + \sigma_k h\right] + \frac{\partial^2 V}{\partial \hat{k}\,\partial \hat{k}'}(\hat{k}, y, \hat{r})\frac{|\sigma_k|^2}{2} \\
& + \frac{\partial V}{\partial y}(\hat{k}, y, \hat{r}) e\left(\sum_{\theta \in \Theta} \theta q(\theta)\right) + \frac{\partial^2 V}{\partial y\,\partial y'}(\hat{k}, y, \hat{r})\frac{|\varsigma|^2}{2} e^2 \\
& + \frac{\partial V}{\partial \hat{r}}(\hat{k}, y, \hat{r})\left(-\zeta + \psi_0 (i^r)^{\psi_1} \exp(-\psi_1 \hat{r}) - \frac{|\sigma_r|^2}{2} + \sigma_r h\right) + \frac{\partial^2 V}{\partial \hat{r}\,\partial \hat{r}'}(\hat{k}, y, \hat{r})\frac{|\sigma_r|^2}{2} \\
& + \sum_{\ell=1}^{L} \mathcal{J}^\ell g^\ell\left[V^\ell(\hat{k}, y, \hat{r}) - V(\hat{k}, y, \hat{r})\right] + \xi \sum_{\ell=1}^{L} \mathcal{J}^\ell\left[1 - g^\ell + g^\ell \log g^\ell\right] + \frac{\xi}{2} h' h + \sum_{\theta \in \Theta} q(\theta) \log q(\theta).
\end{aligned}
$$

This HJB equation depends on a value function difference. In forming this difference, we must use $V^\ell(\hat{k}, y, \hat{r}) = W^\ell(\hat{k}, z, \hat{r}) - (\lambda_1 \hat{y} + \frac{1}{2}\lambda_2 \hat{y}^2)$, subtracted from the pre-technology, pre-damage jump value function. The remaining details for the pre-technology, pre-damage jump HJB equation are as specified in the main text.

F Damage function intensity construction

F.1 From intensities to densities over jump times

Consider the following simplified process for the temperature anomaly:

$$\tilde{y}(t) = y_o + \alpha t. \tag{24}$$

We use $\alpha = 0.017$ in the computations that follow. The integral of interest is:

$$\int_0^{\frac{y-y_o}{\alpha}} \mathcal{J}[\tilde{y}(t)]dt$$

With the change of variables formula, we rewrite this integral as:

$$\frac{1}{\alpha}\int_{y_o}^{y} \mathcal{J}(\tilde{y})d\tilde{y},$$

implying a cumulative distribution function:

$$F(y) = 1 - \exp\left[-\frac{1}{\alpha}\int_{y_o}^{y} \mathcal{J}(\tilde{y})d\tilde{y}\right]$$

where we vary y over the interval of interest. The corresponding density is:

$$f(y) = \frac{1}{\alpha}\exp\left[-\frac{1}{\alpha}\int_{y_o}^{y} \mathcal{J}(\tilde{y})d\tilde{y}\right]\mathcal{J}(y).$$

F.2 Intensity calibration

We now discuss our calibration of the damage intensity function obtained by inverting the mode of the density. Recall that the intensity for damages is

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

$\overline{y}$	y_m	$y_{.1}$	$y_{.9}$	r_1
2.0	1.75	1.63	1.87	2.90
2.5	2.00	1.77	2.25	0.36
3.0	2.25	1.90	2.62	0.11

Table 28: Parameter inputs and into the intensities and implications. y_m is the mode as imposed, $y_{.1}$ is the implied probability .1 and quantile and $y_{.9}$ is the implied probability .9 quantile. In all cases $\underline{y} = 1.5$, $\mathsf{r}_0 = 1.5$, $\mathsf{r}_2 = 0.0001$. The temperature input used for setting r_1 is given by the simplified model (24).

For simplicity set $\mathsf{r}_2 = 0$. Find the mode, y_m, of the density setting the derivative of the density to zero:

$$\frac{\mathsf{r}_1}{\alpha}(y_m - \underline{y})\left[\mathcal{J}^n(y_m) + \mathsf{r}_0\right] - \left(\frac{1}{\alpha}\right)^2 \mathcal{J}^n(y_m)^2 = 0.$$

We specify r_1, so that $y_m = \frac{1}{2}\left(\overline{y} + \underline{y}\right)$. Note that r_1 enters into the expression for $\mathcal{J}^n$.

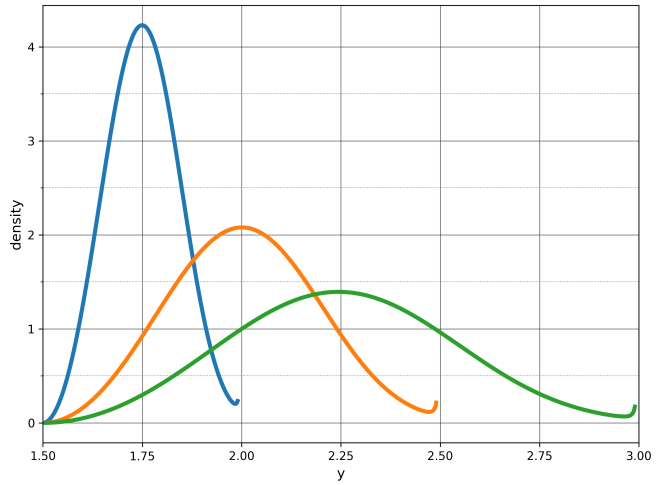

Figure 14: Densities for different choices of $\overline{y}$. Blue density is for $\bar{y} = 2$, the orange density is for $\bar{y} = 2.5$ and the green density is for $\bar{y} = 3$.

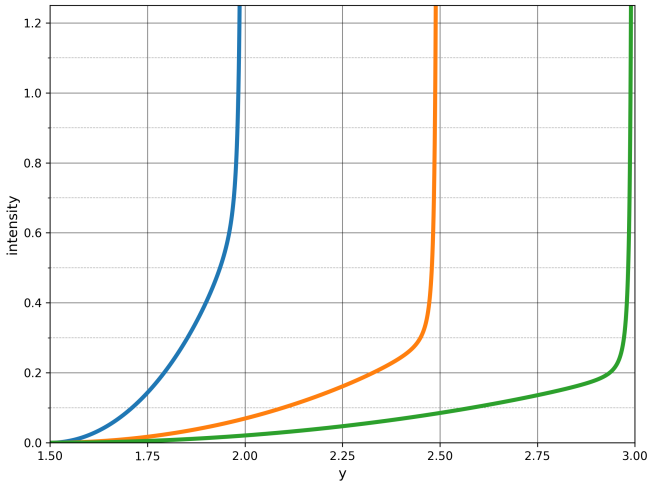

Figure 15: Intensities for different choices of ($\overline{y}$. Blue density is for $\bar{y} = 2$, the orange intensity is for $\bar{y} = 2.5$ and the green intensity is for $\bar{y} = 3$.

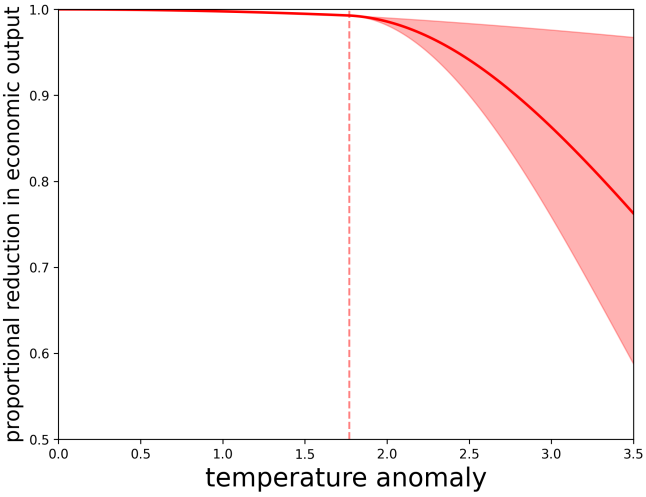

Figure 16: Damage Red

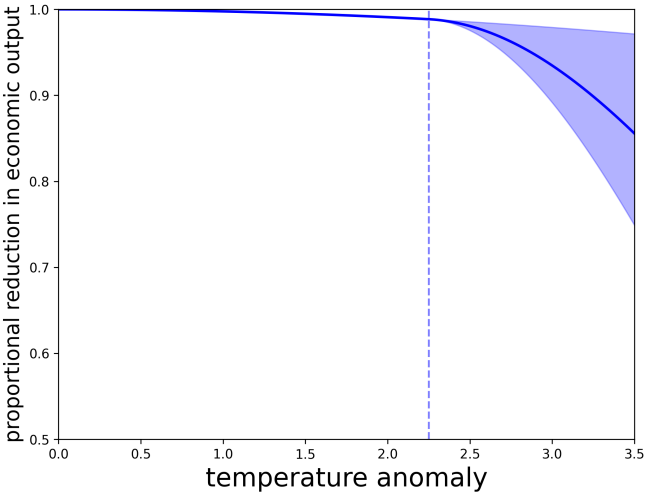

Figure 17: Damage Blue

We use the orange intensity for our baseline calculations. In this appendix we report results for the blue intensity and the green intensity using the r values reported in Table 28 for $\bar{y} = 2.5$.