

Uncertainty, Social Valuation, and Climate Change Policy**

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

June 20, 2025

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, real options theory

[†]Arizona State University

[§]University of Wisconsin and University of Missouri, Columbia

[‡]University of Chicago

*Argonne National Laboratory

**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, 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 supported in part by the University of Chicago Kenneth C. Griffin Economics Incubator and the Haddad Fund for Economics Research at the Becker Friedman Institute for Economics at the University of Chicago.

1 Introduction

Policy analyses in macroeconomics and finance are typically done by imposing a model, including a specific set of parameters, on the policy-maker and deducing the implications for welfare. In this paper, our decision-maker, a fictitious social planner, confronts concerns about uncertainty over models and their potential misspecification when exploring policy implications. Misspecification becomes germane because we envision the use of a model as form of “quantitative storytelling”, where the model provides valuable insights without necessarily being literally correct. In a world of uncertainty, unknown parameters or alternative models give rise to “multiple stories” that are potentially relevant when exploring prudent courses of action or policies. We view potential misspecification concerns as particularly relevant for the analysis of climate change policy. By considering potential misspecifications, we compute robustly optimal policies. These policies provide valuable benchmarks for the effectiveness of *ad hoc* policies implemented in market settings. The *ad hoc* policies could include ones that currently exist or others that are proposed and viewed as politically feasible.

In our application, the family of models that we consider is sufficiently rich to investigate trade-offs that emerge in the presence of uncertainty considerations. Moreover, this family allows us to explore, simultaneously and separately, alternative channels by which uncertainty impacts climate policy. The models we consider use deliberately stark simplifications along some dimensions, as do typical models in macroeconomics and finance, to reveal potentially nonlinear and durable transition mechanisms. We solve and analyze models through the requisite use of global solution methods, deliberately avoiding local approximations because the entire focus is on 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. Our focus on uncertainty leads naturally to the use of tools that come from financial economics to provide novel social valuations via nonlinear stochastic impulse responses along with uncertainty-adjusted probability measures.

We study policy assessment and marginal valuation from a social perspective of our fictitious planner rather than a market perspective. We adopt the social vantage point because climate change induced by human activity is an externality that is at the forefront of policy discussions. Market solutions do not account for the climate change externality, and so there is scope for prudent policy intervention because of the failure of the “invisible hand.”. The social valuations of interest are linked to endogenous state variables, including temperature and the stock of knowledge associated with targeted research and development (R&D).

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

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

While we do not take up the full challenge posed by uncertainties featured by Rising et al., we propose and implement methods for assessing the impact of uncertainty when the nature of this uncertainty is broadly conceived. Our methods give one formal way to incorporate “deep uncertainty” into policy analyses. Our broad approach to uncertainty is consistent with perspectives from multiple fields, including decision and control theory with contributions from statistics, engineering, and economics. These literatures formalize sagacious responses to alternative types of uncertainty, including *risk* within a model (unknown outcomes with known probabilities), *ambiguity* across models (unknown priors weighting alternative models or parameter configurations), and potential *model misspecification* (unknown flaws of fully specified probability models). It is the latter two of these constructs that gives a way to address the “deep uncertainties,” that are often present in policy discussions.

Uncertainty, as it impacts prudent policy-making, often operates through multiple channels. In our example economy, we focus on four such channels:

- i) *carbon-climate dynamics*: mapping from carbon emissions into temperature changes;
- ii) *economic damage functions*: reductions in output because of changes in atmospheric temperature;
- iii) *technological innovation*: breakthroughs for generating energy without using fossil fuels through investment in research and development (R&D);
- iv) *investment productivity*: exogenous shifts in the capital evolution that have persistent consequences.

The first three channels are explicitly linked to climate change and contribute to our analysis in specific ways. We include the fourth channel to provide a benchmark for comparison. By design, it captures the familiar “long-run risk” contributions to the macroeconomy in the absence of climate change via an investment technology.¹

As a refined approach to uncertainty quantification, we explore the impact of the specific channels of uncertainty mentioned previously within an otherwise relatively simple production-based

¹See, for example, Bansal and Yaron (2004). While Bansal and Yaron (2004) do not consider a model with aggregate investment opportunities, it is straightforward to extend it to do so. For instance, see Hansen et al. (2024).

model. Specifically, we compute solutions when the uncertainty in each channel is activated separately, along with a solution when we activate all channels simultaneously. While the composite uncertainties can be high-dimensional, this exercise allows us to ascertain which channel of uncertainty is of most concern for the decisions or policies under consideration. It identifies which knowledge gaps really matter to the decision maker or social planner, which perhaps guides us towards efforts that close the gaps.

Prudent decision making is often linked, through first-order conditions, to social valuations of endogenous state variables, including temperature and the stock of knowledge associated with targeted research and development (R&D). For instance, the marginal valuation of temperature is negative, and we refer to its negative as the social cost of global warming. It is closely linked to the commonly discussed social cost of carbon. While much is made of measurements of the so-called social cost of carbon (SCC), quantifying the social value of R&D directed at the discovery of new and economically viable green technologies has received considerably less attention.²

To elucidate some of our findings, we use asset pricing methods to investigate such marginal valuations. We extend tools from stochastic impulse response theory and from the pricing of uncertain cash flows to provide novel characterizations of the contributions to social valuations and their horizon dependence. Using an asset pricing-type representation allows us to explore the relative contributions of alternative social cash flows induced by marginal changes in initial states. They provide insights into the role of feedback mechanisms among state variables and the primary contributors to social valuation.

In our example economy, our 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. Thus, in addition to the stock of capital, 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. We think of each of these as an asset characterized by its implied intertemporal social cash flows. Importantly, the social payoffs include a direct marginal utility contribution and marginal contributions to prospects of potentially big changes in the future to either damages induced by more extreme values of climate change or to the technology opportunities for addressing climate change. The asset price representations we derive depend on a “tilted probability” distribution that reflects aversion to model ambiguity or model misspecification.

Analysis of our stylized model shows that technological innovation is the dominant uncertainty channel. For a prudent planning problem in our stylized environment, the socially and robustly efficient policy includes both immediate reductions in carbon emissions and a substantial invest-

²But even the social cost of global warming, a central input into the SCC, is altered by our inclusion of R&D investment in the toolkit of the social planner, modifying some previous analyses in a substantive way.

ment in research and development (R&D). The planner reduces carbon emissions in the short run to limit damages and to allow for more time for R&D investments to be successful. The technology innovation channel emphasizes concerns about uncertainty, which spills over to R&D investment. Within the range of uncertainty concerns we consider interesting, R&D investment becomes even more markedly increased. Our use of asset pricing tools for social valuation allows us to understand better why the fictitious social planner wants to be more pro-active in the presence of more uncertainty about the time horizon over which the technology will become a successful replacement for fossil fuels. To exposit our findings about proactive R&D investment in a simpler environment, we abstract entirely from the dynamic information revelation and uncertainty about climate damages. We use this to illustrate the forces at work with fewer moving parts. This simpler environment opens the door to an analogous result within the confines of a highly stylized real options model.

2 Example economy

We consider a simple production-based model with an AK production technology. By design, this formulation has close ties to much of the long-run risk literature with exogenously specified consumption endowments. Our production framework includes multiple (in our example two) investment opportunities. We find this to be a critical component for properly addressing investment within a model with climate change.

Our model specification will introduce two types of Poisson events, one is an informational or awareness event and the other is a technology discovery. These Poisson events are essentially “metaphors” for more complicated phenomena, but ones that could unfold relatively rapidly. Each event type allows us to capture an important contribution to climate change policy in a simple, yet revealing manner. The simplicity also allows the computations to be more tractable. The first one reveals potential damages for more extreme temperatures and allows us i) to consider a dynamic tradeoff between acting now versus waiting, and ii) to investigate the sensitivity of future policies to a more refined understanding of the impact of climate change. The second is a major technological success in the provision of clean energy. This success could come in different forms driven in part by economic viability implications. While actual advances may proceed incrementally, we find it pedagogically useful to focus on the case of a big technological success with a highly uncertain time horizon for success. The mechanisms and outcomes we elucidate will likely remain applicable in models with more complicated dynamics. Importantly, the Poisson outcomes in our setup are triggered by intensity functions that depend on endogenous state variables.

We use this model to report some quantitative illustrations. Our model specification is highly stylized, which makes calibration challenging; but we do want 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 F. While we offer some guidance for our calibration, given the stylized nature of our model specification a more formal and complete calibration is not feasible.³ As a consequence, the model outcomes we report may be thought of as illustrations of potentially important mechanisms. Since there are competing forces in play, we will find the quantitative outcomes to be a revealing and essential component to our economic analysis that follows. Perhaps this can help open the door to what Hansen and Heckman (1996) call a “symbiotic relationship” with a productive intellectual exchanges between model builders and empiricists which will advance an overall climate-economics research program.

Though our policy uncertainty analysis brings some aspects of model sensitivity “inside the decision problem” through the dynamic decision theory toolkit, we also conduct some “outside the decision problem” sensitivity analyses for some key parameters and specifications of the preferences and technology in this appendix.

2.1 Production and innovation

Prior to introducing climate change, we include three modifications. First, we consider an adjustment cost version of an AK model. This allows us to consider a baseline environment with a conventional investment option. Second, we include an energy input in a way that is mathematically similar to what is used in “DICE” models as developed by Nordhaus (2017) and others. Our rationale for this specification is a bit different leading us to different parameter settings. Third, we allow for a second investment option, which is R&D that could eventually remove the need for the energy input. In particular, our decision maker (social planner) has two forms of capital: one is our broad-based notion of a aggregate capital stock, K_t , and the other is the stock of R&D, R_t , devoted to the discovery of a new green technology. The costs of these investments by design will be the same, but the payoffs in production will be quite different in magnitude and timing. While our model is deliberately stylized, some suggest that an economically viable version of nuclear fusion could potentially eliminate the need for dirty energy in production.⁴

³As we study long-duration transitional dynamics with model nonlinearities and broad-based uncertainty, we necessarily rely on global solution methods, and not ones that approximate around stochastic steady states. Thus, computational tractability also limits the complexity of the economic framework that we can reasonably analyze. 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.

⁴See Chang (2022) and Stalard (2024) for recent discussions of progress in the development of an economically viable nuclear fusion technology.

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):

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

for $\phi_1 \geq 2$ and $0 < \phi_0 \leq 1$, where

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

and $\mathbf{1}$ 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.⁵ 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). 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 $\phi_1 = 3$.

Remark 2.1. *This type of mathematical formulation has shown up in many climate-economics papers since the work of Nordhaus. See, for instance, Nordhaus (2017). One could imagine instead a fixed proportions technology for energy and capital. In such a setting, dirty energy reduction less than the required proportion of capital must be replaced by a clean alternative subject to a convex cost. The resulting modification in the output equation would be:*

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

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

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

would then be the counterpart to what is often referred to as an abatement cost, perhaps even a cost that is external to the firm. While we find this alternative representation to be substantively interesting, in what follows we view the right-hand side of (1) as a production relation without

⁵While we could include B_t as an entry in the control vector, A_t , this is unnecessary because B_t is fully determined by emissions and capital.

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 curvature that is missed in a fixed proportion technology, but it presumes that all of the energy inputs are dirty. We have more to say about the production relation in Appendix E.⁶ Observe that when $\mathcal{E}_t = 0$ either 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 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.⁷ Investment in R&D makes this discovery more likely. Specifically, the jump to the clean technology is governed by a Poisson intensity $\mathcal{J}^L(R_t) = R_t$. We elaborate more on technology advances and the evolution of R_t in what follows.

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. 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

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

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

⁶We thank one of our referees for suggesting this representation of abatement costs and its motivation.

⁷See Chang (2022) and Ball (2023) for recent discussions of the state of nuclear fusion technologies and their promise.

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 about the future likelihood of a technological advance.⁸ The jump intensity for a new discovery is proportional to the knowledge stock R .

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 F.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

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

where $\bar{\theta}$ is the equally weighted average of the $\theta(i)$ ’s where each $\theta(i)$ 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 F.5 describes how we constructed this histogram.⁹

⁸For a recent exploration of the policy implications of R&D for a green breakthrough technology, see Jaakkola and van der Ploeg (2019).

⁹As 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

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.¹⁰




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

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 the logarithm of damages $\hat{n}$ with respect to the temperature anomaly is

$$\begin{aligned} \frac{d\hat{n}}{dy} &= \lambda_1 + \lambda_2 y & y \leq \hat{y} \\ \frac{d\hat{n}}{dy} &= \lambda_1 + \lambda_2 (y - \hat{y} + \bar{y}) + \lambda_3(\ell)(y - \hat{y}) & y > \hat{y} \end{aligned} \tag{4}$$

for $\ell = 1, \dots, L - 1$, where $\hat{y}$ denotes the temperature at which the damage jump takes place and $\bar{y}$ denotes the upper limit for when the damage jump can take place. Equation (4) has an initial condition $\hat{n}(0) = 0$. Since $\hat{y} \leq \bar{y}$, the derivative is positive. The implied damage function is the exponential of

policy. Thus, we adopt here this simpler specification to avoid having to include an additional state variable.

¹⁰One 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.

$$\hat{n}(y) = \begin{cases} \lambda_1 y + \frac{\lambda_2}{2} y^2 & 0 \leq y \leq \hat{y} \\ \lambda_1 y + \frac{\lambda_2}{2} \hat{y}^2 + \frac{\lambda_2}{2} (y - \hat{y} + \bar{y})^2 + \frac{\lambda_3(\ell)}{2} (y - \hat{y})^2 - \frac{\lambda_2}{2} (\bar{y})^2 & y \geq \hat{y} \end{cases} \quad (5)$$

Notice from specification (5) 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. The intensity function governing the jump is given by

$$\mathcal{J}^n(y) = \begin{cases} 0 & 0 \leq y \leq \underline{y} \\ r_0 \exp \left[\frac{r_1}{2} (y - \underline{y})^2 + \frac{r_2}{2} (y - \bar{y})^{-2} \right] - r_0 \exp \left[\frac{r_2}{2} (\underline{y} - \bar{y})^{-2} \right] & \underline{y} < y < \bar{y} \end{cases},$$

where $\underline{y}$ denotes the lower bound for when the damage jump can occur. The numerical values for the r 's are provided in Appendix F along with a plot of the implied intensity function. The r_2 term guarantees that the jump occurs prior to the temperature anomaly reaching $\bar{y}$.¹¹ Later we will show graphically probabilities induced by our parameter settings.

The intensity $\mathcal{J}^n$ determines the damage realization 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:

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

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 realization jump at a lower temperature ($\hat{y} = 1.75$) generates a range of damage curves that essentially includes those for a damage realization 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.

Observe that our specification of a jump in damages induces a jump in the continuation values that we determine endogenously. This jump could be thought of as 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 with more extreme concavity of the damage function, in contrast with late damage recognition. We include this informational event in order that our planner confronts a tradeoff between acting now versus waiting until more

¹¹In effect, it imposes a form of “value-matching” at $y = \bar{y}$.

information is available in the future.¹²




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 $\bar{y} = 2.5$. 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 $Y_t = \hat{y} = 2.25$.

Remark 2.2. *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). Our specification of damage uncertainty focuses on information revelation about the severity of climate damage function curvature in the future rather than stark “falling off the cliff” specifications used in some other settings. We intend this as an initial illustrative platform for addressing the layers of important uncertainties in climate damages and prudent society responses. It opens the door to an investigation of acting now versus waiting until more information is revealed.*

The 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 by Waidelelch et al. (2024).

¹²We 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 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 continuation value under risk for each calendar date, t as follows:

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

where $\mathfrak{F}_t$ reflects the date t information available to the planner. The second equality expresses a backward recursion linking future continuation values to the current one. The following differential equation gives the local representation:

$$0 = \delta (\log C_t - \log N_t - V_t) + \lim_{\tau \downarrow 0} \frac{1}{\tau} [\mathbb{E} (V_{t+\tau} - V_t \mid \mathfrak{F}_t)]. \quad (6)$$

This equation becomes a Hamilton-Jacobi-Bellman (HJB) equation once we substitute a candidate value function expressed as a function of the state vector. With a minor abuse of notation, $V_t = V(X_t)$.

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 (6):

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

with a robust counterpart.

Remark 2.3. *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 the Online Appendix. In Appendix C we describe the extension allowing for other values of this substitution parameter.*

3 State and control variables

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

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

where

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

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. We compute a representation for the continuation values expressed as a value function $V(X_t) = V_t$.

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:

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

Consequently, the consumption-capital ratio inclusive of damages satisfies:

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

Remark 3.1. *An alternative and similar model that is straightforward to explore has output diminished by damages. In this case,*

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

where C_t^d is consumption inclusive of damages. While we report results with proportional damages to consumption, we conjecture that similar results could be obtained with proportional damages to output.

4 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 in this section focuses on the minimizing player and

the implications for terms that enter the HJB equation involving the evolution of the value function. We temporarily omit the other terms, remembering that these are also important for deriving the maximizing control law and solving the HJB equation.

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

4.1 Brownian motion misspecification

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

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

Suppose that the state vector process X has a local mean increment $\mu(X_t, A_t)dt$ and stochastic increment $\sigma(X_t, A_t)dW_t$, where A_t is a decision or action taken at time t . 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 $\mathcal{X}$. In the absence of potential misspecification, the drift or local mean of the value function $V(X)$ is given by¹⁴

$$\frac{\partial V}{\partial x'}(X_t)\mu(X_t, A_t) + \frac{1}{2}\text{trace} \left[\sigma(X_t, A_t)' \frac{\partial^2 V}{\partial x \partial x'}(X_t) \sigma(X_t, A_t) \right]. \quad (7)$$

In this equation, we omit time dependence and think of V as the value function for an infinite-horizon discounted problem. Formula (7) captures the time increment to risk confronted by the decision-maker.

When the decision-maker (in our case the planner) entertains possible misspecification we replace (7) with the solution to

Problem 4.1.

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

¹³See Cerreia-Vioglio et al. (2025) for an axiomatic formulation of misspecification aversion.

¹⁴We 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' .

for a penalty parameter ξ .

The minimization captures a form of uncertainty aversion, analogous to risk aversion, since the minimizing objective will be less than or equal to (7). The penalty parameter ξ restrains the concern for robustness to model misspecification. The quadratic penalty in h is a local measure of “relative entropy” or Kulback–Leibler divergence.¹⁵ A limiting choice of $\xi \approx \infty$ implies a minimizing choice of $h = 0$ with an implied contribution given by (7). The minimizer is

$$h^* = -\frac{1}{\xi} \sigma(x, a)' \frac{\partial V}{\partial x}(x) \quad (8)$$

with a minimized objective:

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

Notice that the minimizing drift, h^* , is potentially state dependent. When σ depends on the action a , the drift of interest for valuation and interpretation depends on the maximizing action a expressed as a function of the state. The drift vector, h^* , has a relatively larger contribution when the value function is more adversely exposed to the Brownian increments. The parameter ξ governs the magnitude of the distortions. A smaller value of ξ results in drift adjustments with a larger magnitude.¹⁶

4.2 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 $\ell = 1, 2, \dots, L$. Let $\mathcal{J}^\ell$ denote a state-dependent intensity for 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 $V(X)$:

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

¹⁵For instance, see Anderson et al. (2003) for an elaboration.

¹⁶While this looks obvious from formula (8), it is a bit more subtle because the value function implicitly depends on ξ .

capturing the jump-risk contribution to the decision problem.

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

$$\xi \sum_{\ell=1}^L \mathcal{J}^\ell(x) \left[1 - g^\ell(x) + g^\ell(x) \log g^\ell(x) \right].$$

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

Problem 4.2.

$$\begin{aligned} \min_{g^\ell \geq 0} \quad & \sum_{\ell=1}^L \mathcal{J}^\ell(x) g^\ell(x) \left[V^\ell(x) - V(x) \right] \\ & + \xi \sum_{\ell=1}^L \mathcal{J}^\ell(x) \left[1 - g^\ell(x) + g^\ell(x) \log g^\ell(x) \right]. \end{aligned}$$

Minimization problem 4.2 has a quasi-analytical solution:

$$g^{\ell*}(x) = \exp \left(-\frac{1}{\xi} \left[V^\ell(x) - V(x) \right] \right), \quad (10)$$

with a minimized objective:

$$\xi \sum_{\ell}^L \mathcal{J}^\ell(x) \left[1 - \exp \left(-\frac{1}{\xi} \left[V^\ell(x) - V(x) \right] \right) \right] \leq \sum_{\ell=1}^L \mathcal{J}^\ell(x) \left[V^\ell(x) - V(x) \right], \quad (11)$$

which we use in place of (9). Notice from (10), that the distorted intensity is diminished when the jump in the continuation value function is larger, and conversely when the jump in the continuation value is smaller.

Remark 4.3. *In our example economy, we compute V^ℓ 's by solving value functions conditioned on the damage jumps ($\ell = 1, 2, \dots, L-1$) and the technology jump ($\ell = L$) using the same baseline intensities for the remaining jump possibilities.*

Remark 4.4. *Since*

$$\frac{\mathcal{J}^\ell(x)}{\sum_{\ell=1}^L \mathcal{J}^\ell(x)}$$

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

gives the baseline jump distribution conditioned and damage realization jump happening, the uncertainty-adjusted counterpart is¹⁸

$$\frac{g^\ell(x)\mathcal{J}^\ell(x)}{\sum_{\ell=1}^L g^\ell(x)\mathcal{J}^\ell(x)}.$$

In light of (10), this is constructed using an exponential tilting formula.

4.3 Planner HJB equation

To form the planner HJB equation, we incorporate misspecification concerns for both the diffusive and jump contributions from solving Problems 4.1 and 4.2, respectively:

Problem 4.5.

$$\begin{aligned} 0 = & \max_a -\delta V + \delta \log c - \delta \hat{n}(y) + \\ & + \frac{\partial V}{\partial x'}(x)\mu(x, a) + \frac{1}{2}\text{trace} \left[\sigma(x, a)' \frac{\partial^2 V}{\partial x \partial x'}(x) \sigma(x, a) \right] - \frac{1}{2\xi} \frac{\partial V}{\partial x'}(x) \sigma(x, a) \sigma(x, a)' \frac{\partial V}{\partial x}(x) \\ & + \xi \sum_{\ell}^L \mathcal{J}^\ell(x) \left[1 - \exp \left(-\frac{1}{\xi} \left[V^\ell(x) - V(x) \right] \right) \right] \end{aligned}$$

When finding a solution to this problem, we iterate between solving for the maximizing and minimizing controls and solving a high-dimensional (conditionally) linear system approximating the second-order differential equation for the value function. See Appendix I for details.

The planner Problem 4.5 takes as inputs the L continuation value functions for each of the jump possibilities. These have to be computed in advance. There are $L - 1$ potential damage curve realizations and one technology jump realization. The $L - 1$ damage curve realizations are mutually exclusive. We compute the continuation value functions allowing for one jump, either a technology jump or a damage curve jump, using the intensity $\mathcal{J}^\ell$. If a technology jump occurs first, temperature remains constant, but there is a possibility of a damage-curve realization at some time in the future. This damage curve realization is inconsequential because the temperature anomaly remains constant at the value where any incremental curvature has no impact on the damages. Thus, we drop the damage curve intensities from the computation and the post-technology jump continuation value functions are, therefore, all the same, which simplifies our computations. Additional details on the post-jump value functions for the three possible scenarios are given in Appendix A.

Finally, we note that this specification has a direct link to the long-run risk and recursive preference settings that have been used extensively in the macro-asset pricing literature.

¹⁸The exponential tilting formula turns out to be analogous to a particular smooth ambiguity (Klibanoff et al. (2005)) specification that we describe in Appendix D.

Remark 4.6. *Model misspecification aversion is mathematically equivalent to risk aversion in a recursive utility formulation obtained by setting the risk version parameter, γ , as*

$$\gamma = \frac{1 + \xi}{\xi}$$

This insight holds for discrete-time specifications within the class studied by Kreps and Porteus (1978) and Epstein and Zin (1989). It also carries over to continuous-time recursive utility as in Duffie and Epstein (1992), extended to include jump risk.¹⁹

The robustness concerns provide a substantially different perspective to calibrate uncertainty aversion than the pure risk aversion counterpart, a perspective that we find attractive. This is especially true for climate economic analyses, where it is clearly challenging to assign probabilities with full confidence to important model ingredients. Moreover, the mathematical equivalence will not carry over for the uncertainty decompositions, except in a very mechanical way.

4.4 First-order conditions for the investments and emissions

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

The first-order conditions for investment in new capital are

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

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

$$Q_t = \frac{C_t \left[\frac{\partial V}{\partial \hat{k}}(X_t) \right]}{\delta K_t} \tag{12}$$

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

is the “Q” from the theory of investment, abstracting from damages. In addition to the partial derivative of the value function, this term includes an adjustment for the marginal utility of consumption. The division by K_t occurs because of our choice of $\log K_t$ as a state variable. Consumption damages scale Q_t in (12) by $1/N_t$. The term $\frac{\partial V}{\partial \hat{k}}$ reflects the social valuation, which may be distinct from a marginal valuation.

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

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

Thus

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

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 capital adjusted for uncertainty via the term ςH_t :

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

This equation continues to hold if we divide both terms by N_t . The implied social cost of carbon is

$$\frac{\left[-\frac{\partial \hat{V}}{\partial y}(X_t) \right] (\bar{\theta} + \varsigma H_t) - \mathcal{E}_t \left[\frac{\partial^2 V(X_t)}{\partial y^2} \right] |\varsigma|^2}{\delta \left(\frac{N_t}{C_t} \right)}, \quad (13)$$

and the social benefit is

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

both expressed in terms of damage-adjusted consumption as the numeraire. Notice that the drift distortion 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. This is in contrast to model-based measures that abstract from uncertainty.

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 for temperature.

5 Marginal valuations as asset prices

We next describe a type of local sensitivity analysis that help us interpret some of our main findings. This approach draws on and extends ideas from asset pricing. In terms of an asset pricing analogy, we are very much taking an intertemporal cash flow rather than instantaneous returns 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 important Poisson events that will likely be realized in multiple decades. The reported economic values will be in units of (damaged) consumption.

While there is an extensive applied mathematics literature that supports local sensitivity in diffusion environments, we draw on a particular formulation described in Hansen and Souganidis (2025) because it is well designed for our application.²⁰ It entails using asset pricing-type representations with probabilistic uncertainty adjustments for robustness and marginal contributions to infrequent Poisson events of the type we consider here.

Our planner has three controls: emissions, investment in new capital, and investment in R&D. Associated with these three sets of first-order conditions for the controls are three marginal social valuations: the value of capital $\left(\frac{\partial V}{\partial k}\right)$, the value for 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 particularly interested in the social value of R&D and how uncertainty aversion impacts this valuation.

The asset-pricing type representations of the marginal valuations, expressed in terms of discounted social cash flows, allows us to inspect and deconstruct the contributions to social valuation. The representation of the social values that we use is expressed in terms of an associated expected discounted social cash flow. The two central ingredients that go into these formulas are stochastic marginal impulse responses and uncertainty-adjusted probability measures.

We construct our marginal value representations prior to the first jump. This jump could be a damage curve realization 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.

²⁰See Gobet and Munos (2005) for a rigorous mathematical development of parameter sensitivity for diffusion models and the many literature citations given there. A focal point of this literature is efficient computation for high-dimensional problems, whereas our interest is on using the resulting representations to interpret and deconstruct model implications. In contrast to our analysis and Hansen and Souganidis (2025), this prior literature abstracts from implications of preferences for robustness. It does, however, justify many of the core mathematical components to our analysis.

Given this perspective, we adjust the discounting to incorporate jump probabilities. The resulting discount rate we use for time τ is:

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

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

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

The $g^{*\ell}$'s in these formulas are the uncertainty adjustment factors for the intensities.

We next describe a type of local sensitivity analysis that allows us to deconstruct the implications of the model. Although there is a vast applied mathematics literature that supports local sensitivity in diffusion environments, we draw on the particular formulation described in Hansen and Souganidis (2025) that is well designed for our application. It involves using asset pricing-type representations with probabilistic uncertainty adjustments for robustness and marginal contributions of infrequent Poisson events of the type we consider here.

A key input in this and other related forms of sensitivity is a marginal impulse response processes, Λ , using what are called variational processes. 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. By construction, the process Λ has the same dimension 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.

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

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

where the expectation $\tilde{\mathbb{E}}$ reflects the diffusion dynamics incorporating the minimizing drift distortions, h^* implied by robustness. Uncertainty-adjusted jump intensities are encoded on the state-dependent discount factor process $\{Dis_t : t \geq 0\}$.

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

$$\Lambda_t \cdot Scf_t.$$

In this calculation, Scf_t is a three-dimensional vector with each entry measuring a marginal contri-

bution of a date t state variable to date t cash flow. The dot product with Λ_t reflects the stochastic response of the marginal date t to an initial change in one of the state variables. 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.

5.1 Marginal-utility flow

Flow i) measures the impact of marginal change in an initial change in a state variable on the date t contribution to utility (adjusting for damages):

$$\text{flow i} \stackrel{\text{def}}{=} \delta \left(\frac{\partial \log c}{\partial x} - \frac{\partial \hat{n}}{\partial x} \right) (X_t), \quad (14)$$

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

Again using the chain rule, we write the term (14) 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.

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.

5.2 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:

$$\text{flow ii} \stackrel{\text{def}}{=} \xi \sum_{\ell=1}^L \frac{\partial \mathcal{J}^\ell}{\partial x} \left(1 - \exp \left[-\frac{1}{\xi} (V^\ell - V) \right] \right). \quad (15)$$

The following inequality captures an adjustment for uncertainty:

$$\xi \sum_{\ell=1}^L \frac{\partial \mathcal{J}^\ell}{\partial x} \left(1 - \exp \left[-\frac{1}{\xi} (V^\ell - V) \right] \right) \leq \sum_{\ell=1}^L \frac{\partial \mathcal{J}^\ell}{\partial x} (V^\ell - V). \quad (16)$$

This contribution reflects the impact of state dependence on jump intensity, a contribution that is absent in most “rare event models” in asset pricing. As we shall see, this contribution is important to capture in order to understand the impacts of uncertainty on investment.

5.3 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:

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

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.

5.4 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 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. The same adjustment was prominent in the first-order conditions that we derived.

6 Worst-case distributions

We consider the implied uncertainty-adjusted probabilities that emerge from our analysis for the four sources of uncertainty: climate, damages, productivity, and technology. These are solutions to the minimization problem, evaluated at the solutions to the maximization problem. These should not be viewed as “best-guess” distributions. Rather, they are the “worst-case” distributions subject to penalization that are a vehicle by which the planner constructs robustly optimal courses of action. Although we feature results for two specifications of the penalty parameter: $\xi = .05$ (more aversion) and $\xi = .1$ (less aversion), we also report some results for a larger range of values.²¹ Since the ξ parameter setting is a feature of planner preferences, as external researchers we see little reason to commit to a single “calibrated” value of ξ . It is straightforward to run our solution code for other values of ξ .

While it is hard to interpret directly the magnitudes of ξ , 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.²² 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 refined informational inputs would be most helpful in the design of better policies. In effect, the parameter, ξ , provides a one-dimensional way to isolate the uncertainty inputs from the various sources that are most consequential for the decision problem.

We start by reporting the drift distortions for two capital stock processes. These are given in Table 1. 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

²¹Recall the mathematical equivalence between robustness and recursive utility risk adjustments noted in Remark 4.6. The corresponding risk aversion parameter settings from recursive utility for less aversion ($\xi = .1$) and more aversion ($\xi = .05$) 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.

²²For instance, see Good (1952) and Chamberlain (2000).

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.²³ 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 the reported drift distortions.

	capital	knowledge stock
more aversion ($\xi = .05$)	-.184	-.008
less aversion ($\xi = .1$)	-.096	-.003

Table 1: Drift distortions for capital stock evolution and the knowledge stock evolution at the initial time period

Figure 3 reports the distorted counterparts of the smoothed density for the unknown θ , that characterizes the temperature response to carbon emissions. To compute these altered densities, we used the robust prior-smooth ambiguity counterpart to the model misspecification adjustments as described in Appendix D. The smooth ambiguity parameter is set to induce the same conditional mean shift as implied by the misspecification parameter ξ for the calendar date of interest. As displayed in Figure 3, the densities show a modest shift to the right based on uncertainty concerns.




Figure 3: Distorted density plots for the climate models using the density reported in Figure 1 as the baseline subjective probability specification.

²³In Appendix F.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.

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 a damage jump or a technology jump. By conditioning first on the state variable history, we find that the probability that a jump has occurred in the time period t to be:

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

which increases in t . The implied density obtained by differentiating with respect to t is:

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

We express f as:

$$f_t = \sum_{\ell=1}^L f^\ell(t) \tag{18}$$

where

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

While formula (18) gives an additive decomposition, each f^ℓ is not itself a density because it does not integrate to unity unless $\ell = 1$. The computations so far condition on the path of the state variables. To obtain the densities of interest, we take expectations over the future states variables.

Figure 4 gives the densities for the first jump time. The “neutrality” density is the baseline ($\xi = \infty$) density and the other two include adjustments for specification uncertainty. The modal jump date occurs around 38 years and remains about the same with or without ambiguity aversion. This computation gives a quantitative sense that the potential jumps are a source of *long-term* uncertainties. Under a “business as usual” vantage point, the counterpart density would have a mode that occurred at a much earlier date with a much greater contribution from damages. The uncertainty-adjusted densities shift the probability from the left tail to the right tail of the distribution.






Figure 4: Density for the time of the first jump under the alternative uncertainty aversions.



(a) Contribution of technology jump.



(b) Contribution of damage jumps.

Figure 5: Additive decomposition of the jump-time densities. For the baseline probability specification, the technology discovery contribution is 61%, for less aversion it is 48%, and for more aversion it is 24%.

Next, we decompose this density into two components, one for the damage jumps and the other for the technology jump using the additive decomposition based on formula (18) to depict the relative contributions of the two jump types. The left panel of Figure 5 shows how uncertainty aversion is reflected in the worst-case density of technology, making a delayed success more likely. The delayed prospects for success makes *R&D* investment *less attractive*. Consistent with this finding, the right panel in Figure 5 shows how uncertainty aversion makes it more likely that the first jump will be a damage jump under the worst-case density. Although uncertainty aversion has little impact on the modal date for a damage jump, the distribution across damage curves is sensitive to this 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 ($\xi = \infty$) distribution

for λ_3 is uniform $[0, .33]$ ²⁴ As we see in Figure 6 the uncertainty adjust probabilities are tilted to the right, with a quite extreme shift when $\xi = .05$. The $\xi = .1$ adjustment looks quite modest, although perhaps more prominent than the counterpart reported in Figure 3. In the next section, we explore the various uncertainty channels in a more formal way.




Figure 6: Uncertainty-adjusted probabilities across the alternative damage curves $\lambda_3(\ell)$.

7 A quantitative assessment of the impacts of uncertainty on prudent policy-making

We consider the impact of uncertainty on the social value of R&D and the social cost of global warming (SCGW), measured using partial derivatives of the two capital stocks scaled by the inverse marginal utility of consumption. The social cost of carbon (SCC) implied by the SCWG follows from formula (13).²⁵ In addition we report R&D investment as a fraction of output and emissions. All of the quantities were computed for the initial time period,

7.1 Uncertainty deconstruction

We initially present findings that assess 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

²⁴We used a twenty equally spaced discrete support points in our computations.

²⁵Abstracting from the second derivative and H contributions, both of which turn out to be quite small ($< 5\%$ of the total value), the SCC (in units of \$ per ton of carbon) is $SCWG \times 1000 \times \bar{\theta}$, where $1000 \times \bar{\theta} \approx 1.86$.

minimization problem to the corresponding actions and outcomes when all channels are activated simultaneously.

Table 2 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 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 $\xi = .05$, the percent reductions are larger across the board.

Uncertainty channel	SVRD		R&D investment / output	
	$\xi = .1$	$\xi = .05$	$\xi = .1$	$\xi = .05$
baseline	41.1 (66%)	41.1 (43%)	.0063	.0062
climate uncertainty	42.0 (68%)	43.0 (45%)	.0066	.0068
damage uncertainty	42.4 (68%)	43.7 (46%)	.0067	.0071
productivity uncertainty	40.3 (65%)	39.6 (42%)	.0061	.0058
technology uncertainty	57.6 (93%)	77.5 (82%)	.0123	.0223
all channels	62.2	95.0	.0144	.0335

Table 2: Social value of research and development and the R&D investment-output ratios when different uncertainty channels are activated. The numbers in the parentheses for the SVRD the percent reductions 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.²⁶ 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 8. When we activate only the uncertainty concerns about the productivity channel, R&D investment decreases slightly when going from $\xi = .1$ to $\xi = .05$. In this case, the planner increases capital investment due to uncertainty in productivity growth.

Table 3 presents the results of the uncertainty decomposition for SCGW and emissions. The patterns are very similar to those reported in Table 2. The technology channel is again the dominate contributor to the uncertainty responses. Emissions respond proportionately much less than R&D investment in the face of uncertainty concerns.

²⁶The contributions are not constructed to be additive.

uncertainty channel	SCGW		emissions	
	$\xi = .1$	$\xi = .05$	$\xi = .1$	$\xi = .05$
baseline	54.4 (58%)	54.4 (29%)	9.32	9.32
climate uncertainty	55.7 (59%)	56.9 (30%)	9.30	9.29
damage uncertainty	56.8 (60%)	59.2 (32%)	9.29	9.27
productivity uncertainty	54.0 (57%)	53.4 (29%)	9.32	9.32
technology uncertainty	83.8 (89%)	133.0 (71%)	9.09	8.79
all-channels	94.2	186.8	9.01	8.47

Table 3: Social cost of global warming and emissions when different uncertainty channels are activated. The numbers in the parentheses for the SCGW are the percent reductions when not activating all channels simultaneously.

In Appendix G, we provide additional results about the sensitivity of the social valuations and robustly optimal policy choices derived from the model for alternative choices of ξ . The outcomes for this broader range of ξ values are qualitatively consistent with our two featured specifications. (See Tables 15 and 16.)

7.2 Stochastic simulations

Figure 7 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 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 revelation 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 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 γ_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 γ_3), and there is a prominent increase in R&D investment, presumably because of the forthcoming pronounced damages.

With the exception of the path 2 technological breakthrough, the emissions responses depicted in Panel B of Figure 7 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.




Figure 7: Illustrative stochastic pathways computed with baseline probabilities and controls given by the $\xi = .1$ specification of uncertainty aversion. Path 1: no jumps. Path 2: technology jump (at year 17). Path 3: damage jump (at year 38, $\gamma_3 = .052$). Path 4: damage jump (at year 37, $\gamma_3 = .246$).

7.3 Responses to damage-curve information

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. There will also be an intertemporal aspect of “wait and see,” which is sometimes alluded to policy discussions. 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 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 8 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. This response is relative to the already cautious initial response by the planner to climate change and model uncertainty leading to a nearly 15% – 20% reduction in emissions. The R&D responses become much more heterogeneous with a more complete understanding of damages, as the investment decisions become tailored to potential

damages induced by more extreme global warming. While the increased heterogeneity in the response dominates the comparison, the median R&D investment to output ratio is notably higher (.036) after the damage curve realization than before (.029). Thus, there is some initial overall pullback on R&D investment prior to having a more complete picture of the prospective damages due to global warming. These results emphasize the intertemporal implications of uncertainty in our framework: initial robustly optimal choices supported by the tilted “worst-case” probability distribution; “wait and see” considerations in anticipation of the arrival of further information; and post-jump responses that often call for strengthened emissions reductions and substantially heterogeneous R&D investment conditional on a better understanding of the damages to economic opportunities from more extreme global warming.




Figure 8: Densities for actions just before and just after one of the damage jumps. The left panel reports the densities for R&D investment and the right panel depicts the densities for emission. The median investment level is 2.9 before and 3.6 after damage curvature is revealed. The median emission response is 17.2 before and 15.3 after the damage curvature is revealed.

8 A simplified model with only a technology jump

As a simplified illustration of why the social planner’s R&D investment is more proactive, we consider a model without a damage informational event. Instead we assume there is a single known damage curve as depicted in Figure 9. This constructed damage curve assumes the additional curvature from the λ_3 term begins at a temperature anomaly of $2^\circ C$, and the value of λ_3 is the average across the set of possible $\lambda_3(\ell)$ considered in our main specification. Given we now only have a single potential damage curve, in this example there is only a single jump, a technology discovery. Also, for simplicity we close down the misspecification aversion from the diffusion contribution and focus solely on the prospective technology discovery. For illustration, we consider two values of

aversion parameter: $\xi = .1$, which we refer to as “less aversion,” and $\xi = .05$, which we refer to as “more aversion.”




Figure 9: The single damage function constructed by imposing $\hat{y} = 2$ and the arithmetic average of the $\lambda_3(\ell)'s$.

Figure 10 gives the baseline and two worst-case densities for the technology jump. This figure shows the delay in the prospective success under the uncertainty-adjusted probability measure. For example, the median in the absence of robustness concerns is about 33 years. The enhanced median for the less averse specification is about 43 years, and about 64 years for the more averse specification.

This uncertainty adjustment in isolation makes the R&D investment less attractive. An analogous post-jump impact is reflected in inequality (16) which also reduces the attractiveness of the investment *holding fixed* the continuation value difference. As we shall see, the continuation value differences push in the opposite direction.




Figure 10: Jump time densities for the technology jump only model. The uncertainty parameters are $\xi = .05$ for “more aversion” and $\xi = .1$ for “less aversion.”

	flow i	flow ii	flow iii	sum	R&D investment	capital investment
more aversion	15.0 (16%)	59.9 (64%)	18.7 (20%)	93.6	.0284	.750
less aversion	9.3 (14%)	41.0 (63%)	14.9 (23%)	65.3	.0151	.764
neutrality	5.6 (12%)	31.6 (68%)	9.3 (20%)	46.5	.0075	.773

Table 4: Three flow contributions to the social value of R&D for the technology jump only model given by formulas: (14), (15), and (17), respectively. For “more aversion,” $\xi = .05$; and for “less aversion,” $\xi = .1$. Each flow contribution has been divided by the marginal utility of (damaged) consumption. Both investments are expressed as a fraction of output.

Recall the contributions of social cash flow, flow i given by (14), flow ii given by (15), and flow iii given by (17), to the valuation of R&D stock. Table 4 shows their quantitative importance as they contribute to the overall valuation. From this table, we see that flow ii, the marginal impact of a jump, is the dominant contributor to the overall marginal valuation, accounting for well over sixty percent of the total valuation. Notice that all three terms become larger when we increase the uncertainty aversion. This is true even though the success prospects are delayed more under the uncertainty-adjusted probability distributions. Flow i, the direct marginal utility contribution, is a notably small contributor, but this reflects only the pre-jump impacts. Flow iii, the marginal utility when you jump, captures an analogous contribution after the technology jump has occurred. The last two columns of Table 4 report the values of these ratios in the initial period of time. As we expect from the R&D investment first-order conditions, the larger marginal valuation under enhanced uncertainty aversion is accompanied by larger R&D investment-to-output ratios. This increase is essentially offset by a decline in the capital investment-to-output ratio.

Since flow ii is the dominant contributor to valuation, we further investigate its construction. Observe from formula (15) that the difference between the post-jump continuation value and pre-jump counterpart is a prominent contributor to flow ii. The technology discovery resolves an important contributor to the uncertainty confronting the fictitious social planner. The post-jump continuation values are not very sensitive to uncertainty aversion. In contrast, uncertainty concerns lower the pre-jump continuation values. It is the difference between the post-jump and pre-jump value functions that contributes to flow ii, and the pre-jump increases induced by misspecification aversion make flow term ii larger. This, in turn, induces more investment in R&D.

To complete our discussion, Figure 11 depicts the separate flow contributions of the capital stock and R&D stock of knowledge state variables to the marginal valuation of the R&D stock of knowledge. This allows us to see the dynamic impact of the state interactions on valuation. In this

simplified model, the only nonzero entry of the gradient vector, $\frac{\partial \mathcal{J}^\ell}{\partial x}$, is the partial with respect to the knowledge stock of R&D for this single ℓ model. As a consequence, only the marginal change in this state channel contributes to the flow term ii. In contrast, the only nonzero entry for the gradient vector, $\frac{\partial V^\ell}{\partial x}$, is the partial with respect to capital since $R\&D$ investments are no longer desirable. Thus, only marginal changes in capital contribute to term iii.




Figure 11: Horizon plot of decomposition of the social value of R&D by state. The aversion is set at $\xi = .05$ for “more aversion,” at $\xi = .1$ for “less aversion,” and at $\xi = \infty$ for “baseline”.

As we see from Figure 11, the impact of the R&D state variable is monotone decreasing with the exception of a short initial segment for more aversion. The broad-based capital contribution, however, starts at zero, but it builds gradually over longer horizons with delayed peak effects. Consistent with our Table 4, the overall magnitudes of both contributions are larger under enhanced aversion. Recall that these decompositions by horizon depend on stochastic responses to an initial change in the R&D knowledge stock under the uncertainty adjusted probabilities.²⁷ Since the uncertainty-adjusted probabilities delay the R&D success, the capital contributions by horizon are pushed to the right in this figure.

We described 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. While the second force dominates in our calculations, this finding is sensitive to the range of aversions that are considered. Table 5 shows that for very high misspecification aversions (very low values of ξ), the first force can dominate leading to a reduction in the R&D investment.

²⁷The temperature impacts are negligible and omitted from the plots in Figure 11.

	knowledge stock valuation	R&D investment / output
$\xi = \infty$	44.6	.0075
$\xi = .10$	63.4	.0151
$\xi = .05$	87.9	.0284
$\xi = .01$	88.3	.0300
$\xi = .009$	82.6	.0256
$\xi = .008$	75.0	.0212
$\xi = .007$	66.4	.0166
$\xi = .006$	56.2	.0119
$\xi = .005$	44.4	.0072

Table 5: Social value of R&D (technology jump only) as a function of the robustness parameter, ξ .

9 Comparison to real options theory

There are similarities, but also differences, in our R&D investment problem and a real options problem. For instance, our decision maker can influence the timing of the R&D payoff through investment, but there is not a specific choice of when to exercise the option. In our environment, there is an implicit investment cost because of our output constraint. Arguably, this plays a somewhat similar role to the cost of exercising an option, but the latter cost is only incurred at the date the option is exercised. We end our paper by exploring the ways in which our findings have counterparts to insights from real options theory.

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, as shown in Appendix H, 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 make the two payoffs comparable, 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.

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.

Real options theory imposes a value matching condition and a smooth pasting condition used to deduce the exercise threshold. Write the value function as $\delta[w + G^d(y)]$ for the discrete-payoff model and $\delta[w + G^f(y)]$ for the flow-payoff model where δ is the subjective rate of discount assumed to be equal to the riskless rate of return, and y is a realization of the stochastic income process. In both cases, we show in the appendix that smooth pasting implies that exercise threshold $\bar{y}$ satisfies: $\delta G_y(\bar{y}) = 1$.

As 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 . Larger values of the penalty parameter imply less aversion. 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: 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 $\xi \rightarrow \infty$.²⁸

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 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

²⁸These results have counterparts to risk aversion results reported in Miao and Wang (2007).




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

differences make our uncertainty aversion impact not monotone as displayed in Table 5 and in contrast to the real options findings reported in Figure 12.

10 Conclusions

We propose and implement methods for assessing the consequences of uncertainty on social values and prudent courses of action to address climate change. We use the policy problem to assess the relative importance of the alternative channels by which uncertainty concerns alter robustly optimal decisions and outcomes. Asset-pricing representations for valuations of endogenous state variables allow us to quantify the sources of the stochastic cash flows that contribute to marginal valuations and to assess the importance of state interdependence.

Our substantive findings emphasize the importance of including endogenous R&D investment into quantitative assessments of socially prudent courses of action along with current and near-term future carbon reductions. Our calculations expose the limitations of commonly proposed policy solutions that involve a gradual decrease in emissions to a net zero target without explicitly featuring a role for R&D investments and speculating on their potential successes.

We find an important interaction between R&D investments that have uncertain long-term payoff successes. More uncertainty in the payoff to this investment can result in an increase in the R&D investment relative to output designed to accelerate the possible discovery of a new economically viable green technology. This provides a substantively important example when *more aversion* to potential model misspecification leads to *bolder actions* on the part of the decision

maker.²⁹ As a stylized dynamic response, we find that the R&D responses in the future can be highly sensitive to information about damages to economic opportunities in the future. While emissions respond at the outset to potential damages to economic opportunity that might be realized with more extreme global warming, there is also an aspect of “wait and see.” The response density, when conditioned on more knowledge about damages, is shifted to the left implying an overall reduction in emissions.

Our models are stark along some dimensions to illustrate why uncertainty considerations are important and how they matter. A more realistic policy environment requires additional modeling complexity including heterogeneous geographical exposure to uncertainty, multiple policy-makers, and alternative forms of adaptation investments in the development of new technologies exposed to different forms of uncertainty. 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 will remain a serious possibility.

We conclude with speculative reflections on pragmatic policy challenges. We speculate that there will be an important role for government investment in developing new green technologies, at least until economic viability becomes more self-evident. A common concern to taking strong immediate action by skeptics is that the costs of using inefficient government approaches undermine attractiveness of public solutions, as opposed to market solutions. A blunt way of putting this objection is, “This problem won’t 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, broadly conceived, may be large enough to push for efforts to overcome political distortions and embark on investment directed to the discovery of economically viable clean alternatives, at least at the early stages of development. What is missing in our analysis so far is a more serious probe into the political economy of large-scale public investment projects. Nevertheless, our formulation and robustness analysis feature 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.

²⁹A potential for more broadly-based uncertainty aversion to induce a more proactive policy was demonstrated in a substantially different monetary setting by Sargent (1999) and Cogley et al. (2008).

Appendices

A Post-jump HJB equations

As noted previously, there are L possible jump realizations from the pre-jump setting: $L-1$ potential damage curve realizations and one technology jump realization. From the post-technology, pre-damage setting there $L-1$ possible damage curve jump realizations, and from the pre-technology, post-damage setting there is only one possible 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.

A.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

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

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

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

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

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

with first order conditions given by

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

A.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 damage curve intensities 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.

A.3 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 $\ell = 1, \dots, L-1$. Note that the damage function in this setting depends upon 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 $z = y - \hat{y} + \bar{y}$. This state variable becomes the counterpart to the temperature anomaly state variable y . The dynamics of z are the same as our previous temperature state variable, though z is initialized at $\bar{y}$ at the time of the jump regardless of what the current temperature value is. Using the new state variable, we can specify an adjusted damage function as follows

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

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

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

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 technological innovation realization with

intensity $\mathcal{J}^L$. The post-technology jump continuation value is $W^{\ell,L}$, and the pre-technology jump continuation value is the solution to the following HJB equation:

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

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.

B Pre-technology, pre-damage jump HJB equation

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

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

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.

C Social planner preferences: IES not equal to one

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

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

where $\exp(V)$ gives an ordinaly equivalent representation of preferences since $\exp(\cdot)$ is an increasing function. These preferences are dynamically consistent with a recursive representation.³⁰

The following differential equation gives the local representation:

$$\lim_{\epsilon \downarrow 0} \frac{1}{\epsilon} (V_{t+\epsilon} - V_t) = -\frac{\delta}{1-\rho} \left[\left(\frac{(C_t)^{1-\rho}}{\exp[(1-\rho)V_t]} \right)^{1-\rho} - 1 \right] \quad (19)$$

which is a backward recursion linking future continuation values to the current one. Introducing stochasticity under a unitary risk specification and rearranging terms gives:

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

As in case the when $\rho = 1$, we use the adjustments derived in Section 4 to replace the local mean of the continuation value with the robust counterpart.

D Incorporating ambiguity aversion

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

Problem D.1.

$$\begin{aligned} \min_{q, \int_{\Theta} q(\theta) dP_t(\theta) = 1} & \quad \frac{\partial V}{\partial x'}(x) \int_{\Theta} \mu(x, a | \theta) q(\theta) dP_t(\theta) \\ & + \chi \int_{\Theta} q(\theta) \log q(\theta) dP_t(\theta), \end{aligned}$$

³⁰This 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 make the transformation increasing.

where χ is a penalty parameter.

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

$$q_t^*(\tilde{\theta}) = \frac{\exp\left(-\frac{1}{\chi} \frac{\partial V}{\partial x'}(X_t) \mu(X_t, A_t \mid \tilde{\theta})\right)}{\int_{\Theta} \exp\left[-\frac{1}{\chi} \frac{\partial V}{\partial x'}(X_t) \mu(X_t, A_t \mid \theta)\right] dP_t(\theta)}.$$

The minimized objective is

$$-\chi \log \int_{\Theta} \exp\left[-\frac{1}{\chi} \frac{\partial V}{\partial x'}(x) \mu(x, a \mid \theta)\right] dP_t(\theta). \quad (20)$$

Representation (20) implies an exponential adjustment for model ambiguity concerns.³¹ We allow the baseline probability to be time dependent to allow for recursive learning, although we will abstract from this learning in our application.

E Production function interpretation of abatement

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

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

for

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

as a production function.³² Notice that the candidate production function is homogeneous of degree one in (e, k) .

³¹As noted by Hansen and Miao (2018), this exponential adjustment can be viewed as a continuous-time version of a smooth ambiguity adjustment of the type advocated by Klibanoff et al. (2005).

³²For notational simplicity, we drop the dependence of ϕ_0 on z .

First, we 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 $\alpha(1 - \phi_0) \geq 0$ when $k \rightarrow \infty$ and $\alpha > 0$ when $\beta \alpha k \leq e$. 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

$$\begin{aligned}
\begin{bmatrix} r_1 & r_2 \end{bmatrix} \begin{bmatrix} \frac{\partial^2}{\partial e^2} \text{output} & \frac{\partial^2}{\partial e \partial k} \text{output} \\ \frac{\partial^2}{\partial e \partial k} \text{output} & \frac{\partial^2}{\partial k^2} \text{output} \end{bmatrix} \begin{bmatrix} r_1 \\ r_2 \end{bmatrix} &= \frac{\partial^2}{\partial e^2} \text{output} \begin{bmatrix} r_1 & r_2 \end{bmatrix} \begin{bmatrix} 1 & -\frac{e}{k} \\ -\frac{e}{k} & \left(\frac{e}{k} \right)^2 \end{bmatrix} \begin{bmatrix} r_1 \\ r_2 \end{bmatrix} \\
&= \frac{\partial^2}{\partial e^2} \text{output} \left(r_1 - r_2 \frac{e}{k} \right)^2 \\
&\leq 0.
\end{aligned}$$

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: $\frac{de}{dp}$.

F 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 macroeconomics moments, and (iii) indirect calibration comparing outcomes from the simplified model settings to observable empirical macroeconomic and asset pricing moments. We validate these calibrated parameters choices by examining how the calibrated model fits additional macroeconomic and asset pricing empirical moments not included in the calibration. In addition, we conduct sensitivity analysis for key parameters and specifications related to preferences, the abatement technology, and climate damages as additional validation of our quantitative results. We provide the details for the calibration, validation, and sensitivity analysis in what follows.

The model parameters chosen for preferences and the dynamics of the capital stock, knowledge stock, climate, and climate damages are given in Tables 6 – 10. Steady-state values fit to empirical moments for the capital dynamics calibration are given in Table 11. The initial values and ranges for state variables are provided in Tables 12 – 13. Descriptions for these choices are given below.

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

Table 6: Capital dynamics

Parameter	Value
ζ	0
ψ_0	.1
ψ_1	.5
σ_r	[0, 0.0078]
ϱ	.0011

Table 7: 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 8: Climate dynamics

Parameter	Value
δ	.01
ρ	1
ξ	{.05, .1}

Table 9: Preferences

Parameter	Value
λ_1	.00017675
λ_2	$2 \times .0022$
$\{\lambda_3(\ell)\}_{\ell=1,\dots,L-S}$	$\{\frac{1}{3} \frac{\ell-1}{L-S-1}\}_{\ell=1,\dots,L-S}$
r_1	1.5
r_2	2.5
$\underline{y}$	1.5
$\bar{y}$	2

Table 10: 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 11: 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 12: State Variable Initial Values

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

Table 13: State variable ranges

F.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 the standard log utility case. 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.

F.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 conditions for investment, solve for κ given (i^k, π) :

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

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

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

- 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.

F.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 $\varrho = .0011$, 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 knowledge stock volatility is set to $\sigma_r = [0, .0078]$, matching annual percent changes in the time series of U.S. R&D capital stock from the BLS database.

F.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 estimated parameter values from Cai and Lontzek (2019) and Barrage and Nordhaus (2023). While the value of ϕ_0 is highly uncertain, we choose $\phi_0 = .5$ as a reasonable benchmark for the fraction of lost output in order to achieve zero emissions. We also consider $\phi_0 = .1$, consistent with Barrage and Nordhaus (2023), for a sensitivity analysis comparison. The value for the emissions intensity of output $\beta = .12$ 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 match 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 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 by Waidelech et al. (2024). We choose baseline values for our intensity function of $\underline{y} = 1.5$ and $\bar{y} = 2.5$, motivated by the literature on climate thresholds and tipping points which we elaborate on in Remark F.1. We set $r_0 = 1.5$ and $r_2 = 0.0001$, and then impose $r_1 = 0.36$, 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.³³

Remark F.1. *The choices of $\underline{y}$ and $\bar{y}$, 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°C global warming. A significant likelihood of passing multiple climate tipping points exists above $\sim 1.5^\circ\text{C}$, particularly in major ice sheets. Tipping point likelihood increases further in the Paris range of 1.5 to $< 2^\circ\text{C}$ warming.” Rogelj et al. (2018) and Rogelj et al. (2019) suggest a 1.5°C target for limiting damages from such events, while noting 2.0°C as a potentially more plausible target that also includes risks of more severe damage consequences.*

Consensus on the plausibility and timeline for such thresholds is far from decided however. 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 “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 significant realization of climate damages. 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.

³³For sensitivity analysis, we also consider cases with $\bar{y} = 2.0$ and $\bar{y} = 3.0$. 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.

(2021), rather than a stark “falling off the cliff” specifications used in some other settings.

F.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.

F.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 $R_0 = 11.2 \times \varrho$, 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).

F.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

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

where $\tilde{\mathbb{E}}[\cdot \mid \mathfrak{F}_t]$ denotes the expectation under the uncertainty-adjusted probability distribution, $N_t = \exp[\hat{n}(Y_t)]$, and $\epsilon = 1/12$ is the monthly specification. The results are provided in Table 14. 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 14: 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.

F.8 Model sensitivity

We briefly discuss sensitivity analysis exercises that provide further insight into the economic mechanisms underlying our quantitative implications: changing the abatement technology scaling parameter (ϕ_0), changing the subjective rate of discount (δ), and changing the intertemporal elasticity of substitution (IES) (ρ). More detailed analysis and computational results are provided in the Online Appendix.

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 $\phi_0 = .1$. Incorporating the Nordhaus calibration in our setting 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.

The prior environmental economics literature has explored sensitivity of the SCC to changes in the “discount rate.” Often these analyses feature the implied discount rate used in computing present values abstracting from stochastic discounting. In our setting with misspecification, stochastic discounting and endogenously determined changes in the probability measure are central ingredients in valuation, as is expected from our understanding of asset pricing. Our results show that increasing δ leads to drops in the valuations, confirming a sensitivity often noted in the environmental economics literature.

We also consider alternative values of the IES using the preferences defined in Appendix C. Much of the asset pricing literature has studied the consequences of changing the IES on asset valuation within the setting of an endowment economy. Our economy is a production economy, however, and changing the IES has a big impact on production outcomes. As expected from growth models, the investment 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.³⁴

G Extended Analysis of the SVRD and SCGW

The next two tables present the uncertainty decompositions for an even more modest and more extreme specification of ξ .

Uncertainty channel	SVRD		R&D investment / output	
	$\xi = .2$	$\xi = .025$	$\xi = .2$	$\xi = .025$
baseline	41.1 (83%)	41.1 (32%)	0.0063	0.0063
climate uncertainty	41.6 (84%)	44.7 (35%)	0.0064	0.0074
damage uncertainty	41.8 (84%)	46.0 (36%)	0.0065	0.0079
productivity uncertainty	40.7 (82%)	38.0 (29%)	0.0062	0.0054
technology uncertainty	48.4 (97%)	96.4 (74%)	0.0087	0.0345
all channels	49.8	129.5	0.0092	0.0622

Table 15: Social value of research and development and the R&D investment-output ratios when different uncertainty channels are activated. The numbers in the parentheses for the SVRD the percent reductions when not activating all channels simultaneously.

³⁴When 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.

uncertainty channel	SCGW		emissions	
	$\xi = .2$	$\xi = .025$	$\xi = .2$	$\xi = .025$
baseline	54.4 (78%)	54.4 (14%)	9.32	9.32
climate uncertainty	55.1 (79%)	59.3 (15%)	9.31	9.26
damage uncertainty	55.6 (80%)	63.7 (16%)	9.30	9.24
productivity uncertainty	54.2 (78%)	52.4 (13%)	9.32	9.33
technology uncertainty	66.6 (96%)	212.2 (54%)	9.22	8.41
all-channels	69.4	390.6	9.19	7.55

Table 16: Social cost of global warming and emissions when different uncertainty channels are activated. The numbers in parentheses for the SCGW are the ratios relative to activating all channels simultaneously.

Many climate economics papers do not consider technology uncertainty. Table 17 shows implications for SCGW when we abstract from this uncertainty. To illustrate the impacts, we explore even lower values of ξ . Both the climate and damage channels are now prominent, with the damage channel being the more notable contributor.

Uncertainty channel	SCGW			
	$\xi = .1$	$\xi = .05$	$\xi = .025$	$\xi = .0125$
baseline	54.4 (94%)	54.4 (89%)	54.4 (80%)	54.4 (67%)
climate uncertainty	55.7 (96%)	56.9 (93%)	59.3 (87%)	63.6 (78%)
damage uncertainty	56.9 (98%)	59.2 (97%)	63.7 (94%)	71.8 (88%)
productivity uncertainty	53.9 (93%)	53.4 (87%)	52.4 (77%)	50.1 (62%)
climate+damage+productivity uncertainty	57.7	61.1	67.9	81.2

Table 17: SCGW under different uncertainty-channel specifications. The technology uncertainty channel is switched off.

H Real options application

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 $\sigma = .3$.³⁵ 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

$$0 = -\delta G + G_y\alpha + \frac{\sigma^2}{2}G_{yy} - \frac{\sigma^2\delta}{2\xi}(G_y)^2 \quad (21)$$

where the minimizing h solves:

$$h^* = -\frac{\sigma}{\xi}V_y = -\frac{\sigma\delta}{\xi}G_y. \quad (22)$$

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

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

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

³⁵Miao and Wang (2007) use these numbers in the illustration they report in their Figure 4.

equation:

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

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

$$h^* = -\frac{\sigma}{\xi} V_y^o = -\frac{\delta \sigma}{\xi} F_y.$$

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 (21) with $V(w, y) = \delta[w + G(y)]$. The minimizing h^* is given by formula (22). 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 $\xi \rightarrow \infty$.

Figure 13 plots the derivatives for each model specification and for alternative values of the robustness parameter, ξ . 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 13: Plots of δG_y for different values of ξ .

We plot the pre-jump drift distortions h^* in Figure 14. 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^* are expressed 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.



(A) discrete-payoff model



(B) flow-payoff model

Figure 14: Plots of h^* for different values of ξ .

I 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.

I.1 Policy iteration

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

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

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

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

I.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 its 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.

I.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 systems³⁶. 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

³⁶We 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.

obtain the final expression

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

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

I.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

- 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* .
- Bansal, Ravi and Amir Yaron. 2004. Risks for the Long Run: A Potential Resolution of Asset Pricing Puzzles. *Journal of Finance* 59 (4):1481–1509.
- Barnett, Michael, William A. Brock, and Lars Peter Hansen. 2020. Pricing Uncertainty Induced by Climate Change. *Review of Financial Studies* 33 (3):1024–1066.
- 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.

- Brock, W. and A. Xepapadeas. 2017. Climate change policy under polar amplification. *European Economic Review* 99:263–282.
- Brook, Barry W., Erle C. Ellis, Michael P. Perring, Anson W. Mackay, and Linus Blomqvist. 2013. Does the terrestrial biosphere have planetary tipping points? *Trends in Ecology and Evolution* 28:396–401.
- Cai, Yongyang and Thomas S Lontzek. 2019. The social cost of carbon with economic and climate risks. *Journal of Political Economy* 127 (6):2684–2734.
- 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.
- 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.
- Duffie, D and L G Epstein. 1992. Stochastic Differential Utility. *Econometrica* 60:353–394.
- Epstein, Larry G. and Stanley E. Zin. 1989. Substitution, Risk Aversion and the Temporal Behavior of Consumption and Asset Returns: A Theoretical Framework. *Econometrica* 57 (4):937–969.
- Figueres, Christiana, Corinne Le Quéré, Anand Mahindra, Oliver Bäte, Gail Whiteman, Glen Peters, and Dabo Guan. 2018. Emissions Are Still Rising: Ramp Up the Cuts. *Nature* 564:27–30.
- 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.

- Gobet, Emmanuel and Remi Munos. 2005. Sensitivity Analysis Using Itô–Malliavin Calculus and Martingales, and Application to Stochastic Optimal Control. *SIAM Journal on Control and Optimization* 43 (5):1676–1713.
- Good, Irving J. 1952. Rational Decisions. *Journal of the Royal Statistical Society. Series B (Methodological)* 14 (1).
- 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, Qinqian 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 Thomas J. Sargent. 1995. Discounted Linear Exponential Gaussian Control. *IEEE Transactions on Automatic Control* 40:968–971.
- . 2001. Robust Control and Model Uncertainty. *The American Economic Review* 91 (2):60–66.
- Hansen, Lars Peter and Panagiotis Souganidis. 2025. Stochastic Responses and Marginal Valuation. Manuscript, University of Chicago.
- 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.
- 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. Fuglestedt, 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.
- Levitán, 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.
- Maenhout, P. J. 2004. Robust Portfolio Rules and Asset Pricing. *Review of Financial Studies* 17:951–983.
- 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.
- 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 Streffer, 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.
- Sargent, Thomas. 1999. *Policy Rules for Open Economies (Discussion of Laurence Ball)*, chap. 3, 141–154. University of Chicago Press.
- Stalard, Esme. 2024. Nuclear Fusion: New Record Brings Dream of Clean Energy Closer. *BBC, February 8* .
- 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.
- 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.