

Pricing Uncertainty Induced by Climate Change

Michael Barnett

Arizona State University

William Brock

University of Wisconsin and University of Missouri

Lars Peter Hansen

University of Chicago

Geophysicists examine and document the repercussions for the earth's climate induced by alternative emission scenarios and model specifications. Using simplified approximations, they produce tractable characterizations of the associated uncertainty. Meanwhile, economists write highly stylized damage functions to speculate about how climate change alters macroeconomic and growth opportunities. How can we assess both climate and emissions impacts, as well as uncertainty in the broadest sense, in social decision-making? We provide a framework for answering this question by embracing recent decision theory and tools from asset pricing, and we apply this structure with its interacting components to a revealing quantitative illustration. (*JEL* D81, E61, G12, G18, Q51, Q54)

Received December 7, 2017; editorial decision November 19, 2019 by Editor Harrison Hong. Authors have furnished an Internet Appendix, which is available on the Oxford University Press Web site next to the link to the final published paper online.

Global efforts to mitigate climate change are guided by projections of future temperatures. But the eventual equilibrium global mean temperature associated with a given stabilization level of atmospheric greenhouse gas concentrations remains uncertain, complicating the setting of stabilization targets to avoid potentially dangerous levels of global warming.

– Allen et al. (2009)

We thank James Franke, Elisabeth Moyer, and Michael Stein of RDCEP for the help they have given us on this paper. Comments, suggestions and encouragement from Harrison Hong and Jose Scheinkman are most appreciated. We gratefully acknowledge Diana Petrova and Grace Tsiang for their assistance in preparing this manuscript, Erik Chavez for his comments, and Jiaming Wang, John Wilson, Han Xu, and especially Jieayo Wang for computational assistance. More extensive results and python scripts are available at <http://github.com/lphansen/Climate>. The Financial support for this project was provided by the Alfred P. Sloan Foundation [grant G-2018-11113] and computational support was provided by the Research Computing Center at the University of Chicago. Send correspondence to Lars Peter Hansen, University of Chicago, 1126 E. 59th Street, Chicago, IL 60637; telephone: 773-702-3908. E-mail: lhansen@uchicago.edu.

© The Authors 2020. Published by Oxford University Press.

This is an Open Access article distributed under the terms of the Creative Commons Attribution Non-Commercial License (<http://creativecommons.org/licenses/by-nc/4.0/>), which permits non-commercial re-use, distribution, and reproduction in any medium, provided the original work is properly cited. For commercial re-use, please contact journals.permissions@oup.com

doi:10.1093/rfs/hhz144

Introduction

Our ambition, like that of other researchers, is to understand better the macroeconomic consequences of climate change and conversely how the economic activity will alter the climate in the future. We see this challenge as a problem for which aggregate uncertainty is a first-order consideration and not just a second-order afterthought as it often is in quantitative macroeconomic analyses. Developing a modeling framework that could support policy discussions requires that we quantify the associated uncertainty and assess its impacts on policy design. Addressing this problem requires a structural model in the sense of Hurwicz (1966) because we will be compelled to assess possibilities that are not well represented by historical evidence. Economic dynamics necessarily play a central role. To design, say, an optimal carbon tax compels us to use measurements of the mechanism by which human activity today will affect climate in the future and an assessment of the resultant damages to human welfare. Uncertainty prevails in both the transmission mechanism and the resultant social damages. While much of the economics literature has focused on quantifying social damages, climate science investigates the transmission mechanism by which carbon emissions alter the environment. As is reflected in the Allen et al. (2009) quote, climate science quantifications embed uncertainty, both across models and within any given model. This paper pays particular attention to the interaction of the climate impacts and their economic consequences.

We build and assess *dynamic structural* economic models using:

- a. *decision theory under uncertainty*
- b. *nonlinear impulse response functions*
- c. *dynamic valuation via asset pricing*

In terms of item (a), we use a formal decision problem as a way to conduct a meaningful sensitivity analysis. While much of decision theory within economics is typically axiomatic in nature, for us the resultant recursive representations are also of vital importance for implementation. In terms of item (b), changes in emissions today alter the climate and hence economic damages in current and future time periods. Our interest in the shadow price of the human-induced externality on the climate leads us to use nonlinear counterparts to impulse response functions familiar in macroeconomics and climate science. In terms of item (c), we use asset pricing methods not only to impute market valuations but also social valuations. Our asset pricing vantage point leads us to view the shadow prices of interest as discounted expected values of the impulse responses. As we know, asset prices are “marginal” in nature. In a private market setting, they depend on the stochastic intertemporal marginal rates of substitutions of investors. Because our interest is in social valuation, the prices of interest use the marginal rates of substitution of the preferences of the fictitious planner for stochastic discounting and the pertinent relative

prices. In turn these are sensitive to the formulation of decision theory under uncertainty that we use to represent these preferences. We provide mathematical characterizations of the probability measures that adjust for *ambiguity* over how much emphasis to place on the alternative models and for the potential impact of model *misspecification*. Indeed, we use tools from items (a), (b), and (c) in ways that are intertwined. While our main focus is to apply these tools in social valuation to represent Pigouvian taxes that confront externalities in socially efficient manners, an analogous approach can be developed to study the local impacts of policy changes from socially inefficient allocations.

In this paper, we use the “social cost of carbon” as a target of measurement. Featuring this entity as a tax on an externality is an overly simplified solution to a complex policy problem, both politically and economically. Two challenges in implementing such a tax are (1) what happens to the tax revenues and (2) how do existing distortionary taxes alter an idealized choice of a carbon tax? These challenges carry with them a variety of ramifications for implementation, from determining how best to offset any undesirable distributional consequences to ensuring that proceeds are allocated in ways that are not socially wasteful.¹ Of course, there are questions about how to coordinate any such policy across a variety of political venues. These are all vital questions that are part of actual policy discourse, but not ones that we address in this particular paper. Our aim is to assess what sources of uncertainty matter the most. We use implications for the social cost of carbon to guide those discussions, although we suspect that some of the key uncertainty considerations here should also contribute to other more complex and pragmatic approaches to policy.

Our analysis targets “sensitivity” to uncertainty and potential misspecification. We approach this in two ways. First, we take a preliminary stab at exploring the uncertainty in the transmission mechanism from carbon emissions to the climate (captured by us as temperature changes). Second, we show that the “details” of the economic model can really matter, by conducting our analysis within some different economic configurations of technology and preferences.

In this paper, we feature continuous-time models and corresponding pricing methods that are familiar to financial economists. We will exploit the continuous-time recursive representations of preferences to produce revealing formulas for how alternative uncertainty components are reflected in valuation. While the continuous-time diffusion model gives some pedagogically revealing formulas, our approach has direct extensions to discrete-time models and models with jump components, although we do not develop such connections here.

1. Uncertainty and Approximation

We find it advantageous to explore three components to uncertainty:

¹ Kevin Murphy and Bob Topel have emphasized these points in direct communication.

- *risk* - uncertainty *within* a model: uncertain outcomes with known probabilities
- *ambiguity* - uncertainty *across* models: unknown weights for alternative possible models
- *misspecification* - uncertainty *about* models: unknown flaws of approximating models

The first of these components is captured in scientific discourse by introducing random shocks or impulses into models. With known distributions, this modeling approach captures *risk*. Economists often discuss risk and aversion to that risk. We frame this discussion as one in which outcomes are not known, but probabilities are. For instance economic agents “inside” rational expectation models confront risk. The literature on long-run risk assumes investors have preferences that respond to the intertemporal composition of risk using the recursive formulation originally proposed by Kreps and Porteus (1978). The long-run risk literature uses this framework in conjunction with uncertainty in macroeconomic growth rates. See, for instance, Bansal and Yaron (2004). As many previous researchers have noted, the human impact on the climate is a potentially important source of uncertainty that could play out over long horizons. See, for instance, Jensen and Traeger (2014), Cai et al. (2015), Nordhaus (2017), Hambel, Kraft, and Schwartz (2018), and, especially, Cai, Judd, and Lontzek (2017).

The second of these components, *ambiguity*, reflects the fact that there are multiple models at the disposal of decision-makers motivating the question of how much weight to assign to each of these models in terms of their credibility. This is addressed by subjective probabilities within a Bayesian framework. The robust Bayesian approach explores sensitivity to subjective inputs. Historical data alone have only limited insights in terms of how we conceptualize climate change uncertainty. Some of the potential adverse climate outcomes seem best understood by using climate models designed to help us think through the long-term consequences of human inputs into the climate system. For an example of within model ambiguity, consider the findings reported in Olson et al. (2012) for what they call the climate sensitivity parameter. Figure 3 of their paper reports Bayesian posteriors using an uninformative prior and compares this to an informative prior documenting substantial sensitivity, suggesting the importance of the subjective prior in the analysis. This is not a parameter for which “the evidence speaks for itself.” More generally, the interplay between models and evidence seems vital if we are to think through the consequences of uncertainty, broadly conceived. There are now a variety of climate models with differing implications, so how to confront cross-model uncertainty seems pertinent to an assessment of uncertainty.

In this paper, we apply an approach to model ambiguity that applies the Hansen and Miao (2018) recursive implementation of the smooth ambiguity model originally proposed and axiomatized by Klibanoff, Marinacci, and Mukerji (2005). The smooth ambiguity model provides a differential

preferential response to the uncertainty about models that is distinct from risk. Examples motivated by climate science are given in Millner, Dietz, and Heal (2013) and Lemoine and Traeger (2016), although their analyses are driven by robustness considerations. Such considerations for subjective probabilities have played an important role in Bayesian inferences. For instance, see Berger (1984). Hansen and Sargent (2007), and Hansen and Miao (2018) provide a link between the smooth ambiguity model and a recursive robust prior model.

This third component to uncertainty, potential *model misspecification*, is necessitated by the underlying complexity of the environment to be understood through the guises of insightful models. The climate environment, like the economic one, is complex. Models that we constructed of their interactions are necessarily abstractions designed to help us understand the underlying phenomenon under consideration. They are necessarily misspecified because of our desire for simplicity, and because our understanding of some of the features of the environment is limited. Other model shortcomings may be difficult to pinpoint *ex ante*. Interestingly, some well known climate models are themselves sufficiently complicated that researchers construct simplified approximations typically called emulators that capture some broad features using relatively simple time-series models. See, for instance, Li and Jarvis (2009) and Castruccio et al. (2014). Considerations like these lead us to consider potential model misspecification as an important source of uncertainty.

In summary, we formulate a social decision planner problem that includes concerns about the potential misspecification of alternative models and ambiguity over how much weight to assign to each these models. In so doing, we are following the Hansen and Miao (2018) continuous-time extension of Hansen and Sargent (2007). As we will show, this approach gives revealing continuous-time formulas for pricing uncertainty components to the SCC. Since the formulas target social valuation, not market valuation, they do not provide empirical predictions. Instead we use the SCC sensitivity to alternative sources of uncertainty as a well-posed structural setting for our quantitative investigation. This structural approach yields a probability measure encapsulating the planner's uncertainties and the corresponding aversions. Finally, we illustrate the effect of ambiguity aversion on the SCC.

2. A Model with Reserves and Climate Damages

Our model consists of an information structure and the evolution of endogenous state variables including reserves, cumulative emissions, capital, and environmental damages, along with societal preferences. Figure 1 depicts the economic model components without climate impacts and environmental damages. This model has a Brownian motion information structure and, like many in macroeconomics, is highly stylized. We use it to illustrate a framework for doing dynamic policy analysis in the presence of uncertainty in a numerically tractable setting. But we are cognizant of its limitations and hope




Figure 1 Depiction of the economic model in the absence of climate and economic damages. The model includes Brownian increment shocks, adjustment costs in capital accumulation and curvature in how investment in discovery increases the stock of new reserves.

to add some complexity in future research. The continuous-time, Brownian information structure simplifies some of the implications for social valuation, but it is not essential to the overall approach.²

2.1 Information

To assist some of our characterizations, we presume a Brownian information structure where $W \doteq \{W_t : t \geq 0\}$ is a m -dimensional standard Brownian motion and $\mathfrak{F} \doteq \{\mathfrak{F}_t : t \geq 0\}$ is the corresponding Brownian filtration with \mathfrak{F}_t generated by the Brownian motion between dates zero and t .

In what follows, we let $Z \doteq \{Z_t : t \geq 0\}$ be an exogenously specified, stochastically stable, multivariate forcing process. We write its evolution equation stochastically as

$$dZ_t = \mu_Z(Z_t)dt + \sigma_Z(Z_t)dW_t.$$

In our examples Z will be Ornstein-Uhlenbeck or Feller type processes with affine mean dynamics and either constant or linear volatility dynamics.

2.2 State variable evolution

We consider an extended version of a model used by Brock and Hansen (2018). Capital K evolves as

$$dK_t = K_t \left[\zeta_K(Z_t)dt + \phi_0 \log \left(1 + \phi_1 \frac{I_t}{K_t} \right) dt + \sigma_K \cdot dW_t \right],$$

² Our continuous-time diffusion model is similar in some respects to two prior contributions. Hambel, Kraft, and Schwartz (2018) build and analyze a DICE-type model and consider damage specifications in technology and in technology growth. Our production specification is different, in particular, in our inclusion of reserves as a state variable. The structure of our model, net of climate change, bears some similarity to the ? analysis of two productive capital stock technologies with adjustment costs. Our two stocks, however, produce distinct outputs with one being the stock of reserves.

where I_t is investment and $0 < \phi_0 < 1$ and $\phi_1 > 1$. For computational purposes, we will use the evolution for $\log K$:

$$d \log K_t = \zeta_K(Z_t) dt + \phi_0 \log \left(1 + \phi_1 \frac{I_t}{K_t} \right) dt - \frac{|\sigma_K|^2}{2} dt + \sigma_K \cdot dW_t,$$

where the third dt term is the local lognormal adjustment implied by Ito's lemma.

Output is constrained by an AK model:

$$C_t + I_t + J_t = \alpha K_t,$$

where C_t is consumption, I_t is new investment in productive capital, J_t is investment in new reserves, and $\alpha > 0$ is a productivity parameter. So far, we imposed the adjustment costs in the capital evolution. Alternatively, we could posit the adjustment costs in the output constraint. This model is sufficiently streamlined so that it allows for both interpretations.³

In contrast to standard DICE models, we introduce the possibility or replenishing reserves through an investment J in exploration. We do this because the stock of known reserves does change based on new discoveries as has been captured in some models of oil reserves as we discuss below. As we will see, allowing for reserve augmentation does have important quantitative consequences for our analysis. The stock of reserves, R_t , can be at least partially replenished and evolves according to

$$dR_t = -E_t dt + \psi_0 (R_t)^{1-\psi_1} (J_t)^{\psi_1} dt + R_t \sigma_R \cdot dW_t,$$

where $\psi_0 > 0$ and $0 < \psi_1 < 1$ and E_t is the emission of carbon. For computational purposes, we use the implied evolution for $\log R$:

$$d \log R_t = - \left(\frac{E_t}{R_t} \right) dt + \psi_0 \left(\frac{J_t}{R_t} \right)^{\psi_1} dt - \frac{|\sigma_R|^2}{2} dt + \sigma_R \cdot dW_t.$$

Remark 2.1. This model of reserves has some features in common with others in the literature. The well-known Hotelling (1931) specification is a special case in which J_t is constrained to be zero and $\sigma_R = 0$. To elaborate, let

$$R_t = \int_0^{+\infty} E_{t+s} ds$$

be a total stock of reserves available from date t forward. Then $dR_t = -E_t dt$, or

$$d \log R_t = - \frac{E_t}{R_t} dt.$$

While the Hotelling constraint would give us some pedagogical simplicity and is a revealing platform for illustration, historically the stock of reserves has

³ See the Online Appendix for an elaboration.

been increasing over time because of new discoveries. This motivates why we allow for productive resources to be engaged in exploration.

Another special case is when $\psi_1 = 1$. With this specification, a nonnegativity constraint on J_t may bind for a substantial fraction of time in the solution to the planner's problem. A similar model with these features was analyzed by Casassus, Collin-Dufresne, and Routledge (2018). They treated the counterpart of J_t as an "impulse control problem" whereby J_t is optimally set to zero over time segments determined endogenously. While we view this as an interesting special case, we choose not to address it in this paper.

As a third example, Bornstein, Krusell, and Rebelo (2017) have an industry model of reserves with a counterpart to investment J_t with diminishing returns. They allow for richer dynamics by including an additional state variable they call exploration, whose evolution depends on J_t . Exploration increases the reserve stock in a proportional manner. In contrast, we conserve on state variables by having fossil fuel investment augment the reserve stock. We also allow for the current stock of reserves to alter the productivity of investment J_t in a manner that preserves a constant-returns-to-scale specification.

None of these three papers used their reserve model to explore adverse social implications of carbon admissions. While many previous researchers have imposed a Hotelling (1931)-type constraint, we are particularly interested in the impact of including investment in the new discovery of fossil fuels.

2.3 Damages

Climate literature suggests an approximation that can simplify discussions of uncertainty and its impact. Matthews et al. (2009) and others have purposefully constructed a simple "approximate" climate model:

$$T_t - T_0 \approx \beta \int_0^t E_s ds = \beta F_t, \quad (1)$$

where the F evolution pertinent to this approximation is

$$dF_t = E_t dt.$$

Within this framework, emissions today have a permanent impact on temperature in the future where β is a climate sensitivity parameter.

Of course, this is a rather stark approximation of a complex climate system, and we will entertain some alternatives. A substantial literature in climate science assesses for what purposes this is a revealing approximation, which we will discuss subsequently. There are transient components to temperature fluctuations not explicitly connected to emissions that are needed to capture a more complete characterization of temperature dynamics. These could be captured by an exogenous transient process added to βF_t in our analysis.

We focus on the component that the Matthews et al. approximation is meant to capture. Thus while actual temperature has transient departures, the

contribution to temperature change that might be most pertinent to our analysis of the economic impact of climate change could be the increment βE_t . Even with a richer specification of the climate dynamics, it could be advantageous to feature the longer-term temperature changes induced by human activity as it is not obvious why the transient components should be included when quantifying damages induced by an externality induced by carbon emissions. In this paper we use cumulative emissions, F , and not temperature, T , as the pertinent state variable.

The simplicity of the Matthews et al. approximation is sometimes used to reframe policy questions in terms of a carbon budget. Given knowledge of the parameter β , a maximal allowable change in temperature implies an intertemporal constraint on the amount of emissions and in effect could be used to justify a Hotelling-type constraint on cumulative emissions. But when there is substantial uncertainty about the climate sensitivity coefficient, β , there is corresponding uncertainty about what constraint to impose on emissions. Figure 2 depicts this uncertainty via a histogram and a smoothed density based on evidence reported by MacDougall, Swart, and Knutti (2017). They find the cross model mean value to be 1.72 degrees centegrade per one trillion tons of carbon (TtC). The .05 quantile value is 0.88, which is about half the mean value, and the .95 quantile is 2.52, showing the extensive range of parameter values. When there is substantial uncertainty about β , there is uncertainty about what constraint to impose on emissions. As an alternative, we could impose the constraint on the realized temperature change or on the admissible augmentation of carbon concentration.

Given our limited understanding of how to model damages and long-term uncertainty associated with the impact that emissions might have on the economy, some scholars have doubted the value of building so called integrated assessment models with ad hoc specifications of economic or social damages. Instead some have suggested that the social policy objectives should be framed in terms of temperature increases induced by carbon concentration targets. For recent such arguments, see Morgan et al. (2017) and Pezzey (2019). Imposing admissible temperature or concentration bounds can be represented as an extreme form of damage or penalization function with infinite damages or penalties when a threshold is exceeded. We could use this as our specification for damages, but instead we follow much of the economics-climate literature by penalizing large temperature changes through a so-called “damage function” specified exogenously. Consistent with a more general view of carbon budgeting, this damage function could be taken to be a penalty function instead of a hard constraint where the magnitude of the penalty is dictated, at least in part, by the implied climate outcomes. Recall that our aim is to assess what aspects of uncertainty have the most adverse consequences, and we see value in the modeling formalism. On the other hand, we share concerns about the literal interpretation of ours and others of the computed social costs of carbon.

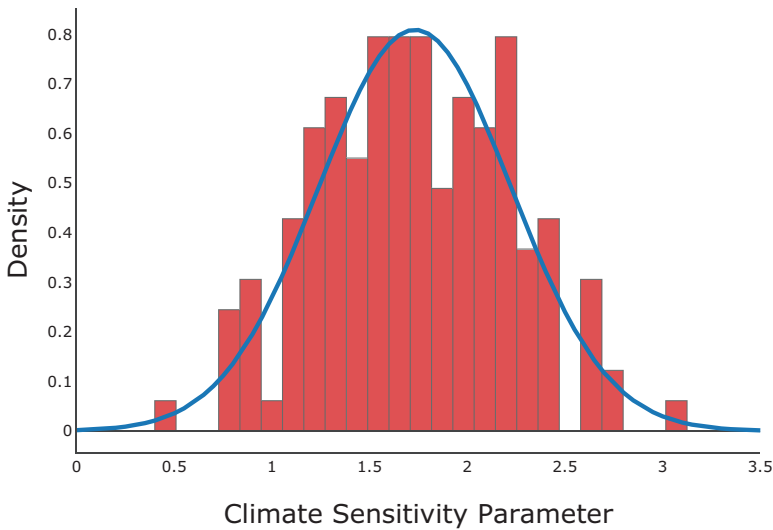


Figure 2 Climate sensitivity uncertainty. Histogram (red) and normal density approximation (blue) for the climate sensitivity parameter β across models. The climate sensitivity parameter is in units of degrees centigrade per teraton carbon. Figure based on evidence reported in Figure 3A by MacDougall, Swart, and Knutti (2017) (© American Meteorological Society, used with permission) and constructed with data provided by the authors.

In this paper, we follow much of the previous literature in economics by positing an ad hoc damage process to capture negative externalities on society imposed by carbon emissions. Just as in the case of the climate approximation, the damage specification we use is an obvious simplification. The economics literature has explored alternative damage specifications typically expressed as functions of temperature. By positing such an evolution we refrain from modeling formally any dynamics associated with adaptation including responses in advance of future temperature increases.⁴ While this model is overly simplistic, the evolution of damages captures two forms of uncertainty that interest us, one from damages that we as depict as uncertainty in the function Γ and the other from climate uncertainty parameter β .

2.4 Consumption damages

In this specification, the instantaneous contribution to the social utility function is

$$\delta(1 - \kappa)(\log C_t - \log D_t) + \delta\kappa \log E_t,$$

where $\delta > 0$ is the subjective rate of discount and $0 < \kappa < 1$ is a preference parameter that determines the relative importance of emissions in the

⁴ While the literature on modeling adaptation to climate change is limited, for a recent example focused on agriculture, see Keane and Neal (2018).

instantaneous utility function. Abstracting from damages, the instantaneous utility is the logarithm of a Cobb-Douglas composite good that depends on material consumption and an energy component that is proportional to emissions. We incorporate damages into this analysis by presuming that diminishes proportionately the material consumption component to the composite good. While in this representation, damages enter the utility function, we may equivalently think of this as a model with proportional damages to production along the lines suggested by Brock and Hansen (2018).

We model the logarithm of damages as

$$\log d = \Gamma(\beta f) + \zeta_D(z) \cdot \begin{bmatrix} f \\ 1 \end{bmatrix},$$

where ζ_D is a two-dimensional vector. With this specification, $\zeta_D(z) \cdot \begin{bmatrix} f \\ 1 \end{bmatrix}$ potentially captures two forms of uncertainty in damage/climate sensitivity by adding an exogenous shifter to the logarithm of damages. One component is deliberately proportional to the temperature anomaly. The other component could capture a distinct role for more transient changes in temperature on damages or other technological contributions that could affect damages. As we will see, this exogenous component opens the door to possible model misspecification that is at least partially disguised by the Brownian increments dW_t . The other component could capture a distinct role for more transient changes in temperature on damages or other technological contributions that could affect damages. The implied evolution for $\log D$ is

$$d \log D_t = [\nabla \Gamma](\beta F_t) \beta E_t dt + d\zeta_D(Z_t) \cdot \begin{bmatrix} F_t \\ 1 \end{bmatrix} + \zeta_D(Z_t) \cdot \begin{bmatrix} E_t \\ 0 \end{bmatrix} dt, \quad (2)$$

where $[\nabla \Gamma]$ is the first derivative of the function Γ .

In our subsequent illustration we parameterize Γ as

$$\Gamma(y) = \begin{cases} \gamma_1 y + \frac{1}{2} \gamma_2 y^2 & 0 \leq y < \bar{y} \\ \gamma_1 \bar{y} + \frac{1}{2} \gamma_2 \bar{y}^2 + \frac{1}{2} \gamma_2^+ (y - \bar{y})^2 & y \geq \bar{y}, \end{cases} \quad (3)$$

where $\gamma_2^+ \geq 0$. To illustrate the impact of damage uncertainty, we focus on the parameter γ_2^+ . For a low damage specification, we set this parameter to zero and for a high damage specification we set it to be a positive number. By setting γ_2^+ to an arbitrarily large number, we approximate a carbon budget constraint by penalizing damages in excess of \bar{y} . While the construction of \bar{y} is suggestive of a “tipping point,” previous literature has explicitly focused on tipping points with uncertain consequences. Of course, other damage functions are also of interest. Observe that the uncertainties about the economic damage function Γ , in general, or the parameter γ^+ , in particular, and the geophysics climate sensitivity parameter β are in effect multiplicative as they contribute to social




Figure 3 Economic damage uncertainty. The two curves plot D as a function of the temperature net of preindustrial levels for two alternative damage configurations. The vertical axis gives the corresponding damage percentage.

welfare. Because of this interaction, it would be misleading to simply add together the uncertainties from the two sources.

In our computational example, we use the two damage functions depicted in Figure 3. The low damage specification is implemented by setting $\gamma_2^+ = 0$. In terms of the previous environmental economics literature, we imagine the case in which $\gamma_2^+ = 0$ as an approximation to Nordhaus (2018). One can see from this figure that our 3° C percentage loss is approximately the same as that of Nordhaus and Moffatt (2017), who say, “...the estimated impact is $-2.04(+2.21)\%$ of income at 3° C warming... We also considered the likelihood of thresholds or sharp convexities in the damage function and found no evidence from the damage estimates of a sharp discontinuity or high convexity...” Weitzman (2012) argues for a steeper degradation in the damages and motivates his construction of an alternative damage function on the basis of uncertainty considerations. Rather than simply impose an approximation to Weitzman’s damage function we illustrate an uncertainty adjustment by positing an alternative even steeper function over some of the temperature increment region and consider the impact of weighting the two possibilities. This allows us to characterize the uncertainty contribution explicitly.

There are two interconnected forms of uncertainty in the evolution of damages that we will capture in conjunction with Equation (2), one from the specification of the damage function Γ and the other from climate uncertainty parameter β .

2.4.1 Damages to macroeconomic growth. Alternatively, suppose that damages diminish growth in the capital evolution:⁵

$$d \log K_t = \zeta_K(Z_t) dt - \Gamma(\beta F_t) dt - \zeta_D(Z_t) \cdot \begin{bmatrix} F_t \\ 1 \end{bmatrix} dt + \phi_0 \log \left(1 + \phi_1 \frac{I_t}{K_t} \right) dt - \frac{|\sigma_K|^2}{2} dt + \sigma_K \cdot dW_t.$$

Not surprisingly, and as discussed in previous literature (see, for instance, the recent discussion in Diaz and Moore 2017), this difference can have an important impact on computations of the social cost of carbon.⁶ Examples of empirical analyses that seek to bear on this issue are Dell, Jones, and Olken (2012) and Burke, Hsiang, and Miguel (2015), who have different perspectives on the importance of heterogeneity and nonlinearity based on reduced-form panel data evidence. From our perspective, this reinforces the notion of damage rate uncertainty.

Several researchers, including Dell, Jones, and Olken (2012), Burke, Hsiang, and Miguel (2015), Burke, Davis, and Diffenbaugh (2018), and Colacito, Hoffmann, and Phan (2019), have looked empirically at the relation between macro growth and temperature. Dell, Jones, and Olken (2012) explore cross-country evidence including lagged effects. They document the largest impacts of temperature on macroeconomic growth occur for low income countries. While they find evidence for a long-term impact, the quantitative magnitude of the impact is much reduced. The climate-economic system potentially has feedbacks in both directions and a single equation approach may be a flawed way empirically to deduce the long-term impacts. The heterogeneity in the impacts across economies at different stages of economic development does seem to be both empirically and substantively important. Unfortunately our simplified analysis in this paper is not designed to confront this heterogeneity, although the consequences of uncertainty will remain for a more refined analysis.

Figure 4 uses reported evidence from Burke, Davis, and Diffenbaugh (2018) exploiting cross-country variation in development and temperature exposure. They report cross-country evidence with temperature and its square regressors (in addition to fixed effects).⁷ Their featured econometric specification has a

⁵ Bansal, Kiku, and Ochoa (2017) and Hambel, Kraft, and Schwartz (2018) give alternative stochastic models of damages to macroeconomic growth. Both use a recursive utility specification for preferences with a risk-based approach where the decision-maker knows the probabilities.

⁶ The material in section 9 of Diaz and Moore's (2017) supplementary online material directly speaks to this point. See Moyer et al. (2014), who provide an initial illustration to show that modifying a DICE-type model to include damages to the growth rate of productivity could have a big impact on the SCC.

⁷ Relatedly, Burke, Hsiang, and Miguel (2015) show how a quadratic specification for the temperature impact on growth can capture the heterogenous temperature responses previously documented by Dell, Jones, and Olken (2012) and others.

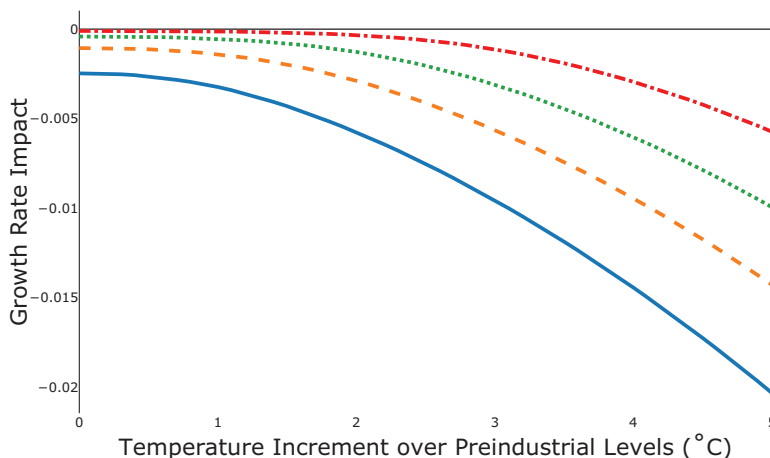


Figure 4

Macroeconomic growth rate damages. The reported quintiles are constructed using estimates from Burke, Davis, and Diffenbaugh (2018) provided by the authors. The blue solid line represents the probability .2 quintile, and the red dot-dashed line represents the .8 quintile. The intermediate curves are the .4 and .6 quintiles.

homogeneous growth response to temperature and abstracts from more lagged impacts that might emerge through adaptation.

Our growth damage function is constructed from the estimated coefficients from Burke, Davis, and Diffenbaugh (2018). Our γ_1 and γ_2 roughly correspond to the linear and quadratic temperature effects, respectively, on economic growth in their global effect regression.⁸ There are nontrivial issues in converting this evidence to a single region, say world, model, leading us to make some ad hoc choices in how we report and subsequently use their evidence.⁹

As we will see this quadratic specification of temperature on economic damages will have rather dramatic implications for the policy implications of our climate-economic model, and we include this in large part to illustrate the impact of damage uncertainty. We have some skepticism as to how far one can go in using developing country responses to quantify more generally global responses to temperature changes by extrapolating from lower income countries in locations with higher temperature.¹⁰ Moreover, given historical evidence alone it is likely to be challenging to extrapolate climate impacts on a world scale to ranges in which many economies have yet to experience.

⁸ See Figure 1A and the estimated coefficients β_1 and β_2 from equation 1 in their methods section.

⁹ The preindustrial temperature level corresponds to a value of approximately 13°C in temperature levels as measured by historical records. We use 13°C as the baseline for the construction of the temperature anomaly values that arise in our model. This value is in line with the median no damage temperature value estimated in Burke, Davis, and Diffenbaugh (2018). We thank Marshall Burke for answering our questions about their work and directing us to the GitHub repository for the relevant inputs need for our computations. Neither he nor his coauthors bear responsibility for how we used their very interesting evidence.

¹⁰ These studies do include fixed country and time effects.

Both richer dynamics and alternative nonlinearities may well be essential features of the damages that we experience in the future due to global warming. Burke, Davis, and Diffenbaugh (2018) give a thoughtful treatment of the impact of parameter uncertainty that we exploited when constructing Figure 4 and that we draw on in our computations that follow.¹¹

3. Implications of Hamilton-Jacobi-Bellman Equations

We start by deducing the relatively standard optimization implications of our model in the absence of ambiguity and model misspecification concerns. The following notation will be used in setting up social planner Hamilton-Jacobi-Bellman (HJB) equations. Let the state vector X_t include $\log K_t, \log R_t, \log D_t, F_t, Z_t$, and let the action vector A_t include $\frac{I_t}{K_t}, \frac{J_t}{K_t}$ and $\frac{E_t}{R_t}$. Write the composite state equation as

$$dX_t = \mu_X(X_t, A_t)dt + \sigma_X(X_t)dW_t,$$

where $\sigma_X(x)' \sigma_X(x)$ is nonsingular m by m matrix. Let n denote the number of states. In what follows we use lower-case letters to denote potential realized values. For instance, d is a possible realization of $\log D_t$, k is a possible realization of $\log K_t$ and r is a potential realized value of $\log R_t$. In terms of the actions, i and j are possible realizations of the investment ratios $\frac{I_t}{K_t}$ and $\frac{J_t}{K_t}$ and e is a possible realization of emissions $\frac{E_t}{R_t}$. We denote the value function by $V(x)$. For our alternative model specifications, some of the state variables enter into the value function in ways that we can exploit for computational simplicity.

3.1 Consumption damages

The HJB equation for this setup abstracting from robustness is

$$0 = \max_{a \in \mathbb{A}} -\delta V(x) + \delta(1 - \kappa) [\log(\alpha - i - j) + k - d] + \delta \kappa (\log e + r) + \frac{\partial V}{\partial x}(x) \cdot \mu_X(x, a) + \frac{1}{2} \text{trace} \left[\sigma_X(x)' \frac{\partial^2 V}{\partial x \partial x'}(x) \sigma_X(x) \right], \quad (4)$$

where \mathbb{A} is a constraint set for the realized action or decision a . As part of a guess and verify approach, the implied value function coefficient for the logarithm of damages is $\kappa - 1$. The pertinent terms for the first-order conditions for the actions or controls are:

$$\delta(1 - \kappa) [\log(\alpha - i - j)] + \delta \kappa \log e + (\kappa - 1) \left([\nabla \Gamma](\beta f) \beta + \zeta_D(z) \cdot \begin{bmatrix} 1 \\ 0 \end{bmatrix} \right) e \exp(r) + V_f(x) e \exp(r) + V_k(x) \phi_0 \log(1 + \phi_1 i) + V_r(x) (-e + \psi_0 \exp[\psi_1(k - r)] j^{\psi_1}).$$

¹¹ While cross-country differences in the long-term impact of temperature on growth is likely to be pronounced, interestingly Colacito, Hoffmann, and Phan (2019) also find that seasonal differences are important in an advanced economy like that of the United States.

The first-order conditions for i , j , and e are¹²

$$-\frac{\delta(1-\kappa)}{\alpha-i-j} + \frac{\phi_0\phi_1 V_k(x)}{1+\phi_1 i} = 0, \quad (5)$$

$$-\frac{\delta(1-\kappa)}{\alpha-i-j} + V_r(x)(\psi_0\psi_1)j^{\psi_1-1} \exp[\psi_1(k-r)] = 0, \quad (6)$$

$$\frac{\delta\kappa}{e} + V_f(x)\exp(r) - V_r(x) + (\kappa-1) \left([\nabla\Gamma](\beta f)\beta + \zeta_D(z) \cdot \begin{bmatrix} 1 \\ 0 \end{bmatrix} \right) \exp(r) = 0. \quad (7)$$

We denote the solution for the investment-capital ratio as $i^*(x)$ and for the exploration-capital ratio as $j^*(x)$. The first-order conditions for the two investments can be solved separately from the first-order condition for emissions. Moreover, there is a further simplification as the first-order condition for investment in capital implies the affine relationship (conditioned on state variables)

$$\phi_0\phi_1 V_k(x)(\alpha - i^* - j^*) = \delta(1-\kappa)(1 + \psi_1 i^*),$$

which can be exploited in computation.

3.1.1 Relative prices of capital and reserves. As is typical in the investment literature, we define the relative price π , sometimes referred to as Tobin's q , as the marginal rate of substitution between capital and consumption:

$$\pi(x) = V_k(x) \left[\frac{\alpha - i^*(x) - j^*(x)}{\delta(1-\kappa)} \right] = \frac{1 + \phi_1 i^*(x)}{\phi_0\phi_1}, \quad (8)$$

where the second relation follows from the first-order conditions (5) for investment in new capital. While the first-order conditions are for the investment-capital ratio, the value function argument is the logarithm of capital. These two adjustments net out in our construction of π .

Analogously, we define the relative price ρ as the marginal rate of substitution between the reserve stock and consumption:

$$\rho(x) = V_r(x) \left[\frac{\alpha - i^*(x) - j^*(x)}{\delta(1-\kappa)} \right] = \frac{j^*(x)^{1-\psi_1} \exp[\psi_1(r-k)]}{\psi_0\psi_1},$$

where the second equality is implied by the first-order conditions (6) for investment in new reserves.

In the construction of these prices, we use the marginal utility of consumption. Depending on the interpretation of the model, we could use either C_t or the damaged counterpart C_t/D_t as the numeraire good. Use of the latter replaces the marginal utility contribution $\frac{\alpha - i^*(x) - j^*(x)}{\delta(1-\kappa)}$ with $\frac{\alpha - i^*(x) - j^*(x)}{\delta(1-\kappa)\exp(d)}$ in the price constructions. Thus, in both cases, the formulas would include an additional multiplication by $\exp(d)$ under the second choice of numeraire good.

¹² In imposing first-order condition (5), we allow for "disinvestment," that is, we permit $i < 0$. This outcome is not prevalent in our model solution, however.

3.1.2 Social cost of carbon. The social marginal rate of substitution between emissions and consumption is commonly referred to as the social cost of carbon (SCC). Thus it is a shadow price of the resource allocation problem for a hypothetical planner. It could be implemented via a Pigouvian tax that would correct the private shadow price for the externality, although we use this way to assess the impact of uncertainty, when conceived broadly. Following previous literature, we start by representing this social cost in terms of partial derivatives of the value function of the social planner. We then apply an asset pricing perspective to interpret components to this social cost. This follows in part discussions in Golosov et al. (2014). Cai, Judd, and Lontzek (2017) have a more ambitious exploration of the risk consequences for the social cost of carbon. We also embrace an asset pricing interpretation, but we will show how to extend the analysis to include forms of uncertainty other than risk. Our purpose in making this asset pricing link goes beyond the particular example economy that we posited. This same perspective also allows researchers to understand better the components to the social cost applicable in more general settings.

The marginal utility of emissions as a function of the state vector is given by

$$\frac{\delta\kappa}{e^* \exp(r)} = \frac{V_r(x)}{\exp(r)} - V_f(x) + (1-\kappa) \left([\nabla\Gamma](\beta f)\beta + \zeta_D(z) \cdot \begin{bmatrix} 1 \\ 0 \end{bmatrix} \right),$$

which follows from the first-order conditions (7). Dividing by the marginal utility of consumption gives

$$scc(x) = \left[\frac{V_r(x)}{\exp(r)} - V_f(x) + (1-\kappa) \left([\nabla\Gamma](\beta f)\beta + \zeta_D(z) \cdot \begin{bmatrix} 1 \\ 0 \end{bmatrix} \right) \right] \left[\frac{\alpha - i^*(x) - j^*(x)}{\delta(1-\kappa)} \right].$$

As with the constructions of q^* and r^* , the scaling by capital nets out when forming the marginal rate of substitution used in the social cost of carbon construction.

The social cost induced by the externality is captured by the two terms:

$$ecc(x) = -V_f(x) + (1-\kappa) \left([\nabla\Gamma](\beta f)\beta + \zeta_D(z) \cdot \begin{bmatrix} 1 \\ 0 \end{bmatrix} \right), \tag{9}$$

scaled by the current period marginal utility for consumption. Both of these can in turn be expressed as expected discounted values of future social damages. To motivate this representation, consider impulse response functions for the logarithm of damages in the future induced by a marginal change in emissions today. This is necessarily a nonlinear impulse response and hence will be state dependent. The marginal emissions change induces an impact on $\log D_{t+u}$ given by¹³

$$\left([\nabla\Gamma](\beta F_t)\beta + \zeta_D(Z_t) \cdot \begin{bmatrix} 1 \\ 0 \end{bmatrix} \right) + \int_0^u [\nabla^2\Gamma](\beta F_{t+\tau})\beta^2 E_{t+\tau} d\tau. \tag{10}$$

The first contribution in (10) occurs on impact and is independent of u because emissions at t have an (approximately) permanent impact on the logarithm

¹³ Following our earlier notational convention, $[\nabla^2\Gamma]$ denotes the second derivative of Γ .

of damages. The second term (10) reflects the nonlinear dependence of the logarithm of damages on state variable f . It includes an integral because of the accumulative impact of emissions on this state variable. Because these are expressed as marginal utilities, we discount using the subjective rate δ . Using a simple integration-by-parts argument, we write

$$ecc(x) = (1 - \kappa) \mathbb{E} \left[\int_0^\infty \exp(-\delta\tau) [\nabla^2 \Gamma](\beta F_{t+\tau}) \beta^2 E_{t+\tau} d\tau \mid X_t = x \right] + [\nabla \Gamma](\beta F_t) \beta + \zeta_D(Z_t) \cdot \begin{bmatrix} 1 \\ 0 \end{bmatrix} \quad (11)$$

divided by the current period marginal utility of consumption. In formula (11), we use the notation \mathbb{E} to denote the expectation operator.¹⁴ In the Online Appendix we show that formulas (9) and (11) coincide.

Thus the ecc is an expected discounted impulse response of marginal damages induced by current period emissions divided by the current period marginal utility of consumption. The discounting here is with respect to the subjective rate of discount because we are working with marginal utilities. This overall approach of representing the ecc as a discounted expected value extends to more complex models of climate dynamics. But so far, we have presumed knowledge of the climate dynamics when constructing this cost. We will have much more to say about uncertainty adjustments in the next section.

3.2 Damages to macroeconomic growth

We briefly describe the corresponding set of calculations of the model in which there are damages to capital evolution. In this specification, we no longer make reference to an explicit damage state variable. The pertinent terms from the HJB equation for optimization are given by

$$\delta(1 - \kappa) \log(\alpha - i - j) + \delta \kappa \log e + V_k(x) \phi_0 \log(1 + \phi_1 i) + V_r(x) [-e + \psi_0 \exp[\psi_1(k - r)] j^{\psi_1}] + V_f(x) e \exp(r).$$

Even with the modifications, the first-order conditions for i and j remain the same. The value function and its derivatives are different, however, as is the first-order condition for e :

$$\frac{\delta \kappa}{e} + V_f(x) \exp(r) - V_r(x) = 0.$$

¹⁴ The second term in (11) also can be written as a discounted expectation of

$$\delta \left([\nabla \Gamma](\beta F_t) \beta + \zeta_D(Z_t) \cdot \begin{bmatrix} 1 \\ 0 \end{bmatrix} \right),$$

which is the same for all $\tau \geq 0$.

Thus the implied marginal utility for emissions satisfies

$$\frac{\delta\kappa}{e^*\exp(r)} = \frac{V_r(x)}{\exp(r)} - V_f(x).$$

We now think of $-V_f$ divided by the marginal utility of consumption to be the external contribution to the social cost of carbon. The instantaneous utility cost induced by a marginal change in e is given by $-V_k(x)[\nabla\Gamma(\beta f)\beta]$, and as we show in the Online Appendix:

$$ecc(x) = \mathbb{E} \left[\int_0^\infty \exp(-\delta\tau) V_k(X_{t+\tau}) [\nabla\Gamma(\beta F_{t+\tau})\beta] d\tau \mid X_t = x \right].$$

Changing the numeraires at each date from utils to consumption entails replacing V_k by the relative price π^* as given by formula (8) so that the social costs being discounted weight marginal damages by π^* .

4. Incorporating Additional Uncertainty Components

As formulated so far, the planner’s problem only features risk and not other components of uncertainty. We now explore multiple ways to capture a broader notion of uncertainty, beyond just risk, that exploit some simplifications that emerge from our continuous-time formulation. In what follows, we capture ambiguity and model misspecification concerns conveniently with two parameters (ξ_p, ξ_m) following an approach suggested by Hansen and Sargent (2007) and extended to continuous time by Hansen and Miao (2018). From a computational/mathematical perspective, they act as penalization parameters that restrain the sensitivity analysis of alternative models (ξ_p) , and the exploration of the potential misspecification of those models (ξ_m) . An outcome of the computation will be an alternative probability measure that reflects aversions to model ambiguity and to the potential misspecification of each of the models under consideration by the social planner. In constructing such a measure, we borrow convenient mathematical tools used extensively for pricing derivative claims. The measure emerges as part of our solution to an HJB equation for the planner who designs policies that are aimed to be sensibly robust in the presence of this uncertainty. In effect, this probability is an uncertainty-based pricing measure. In this section, we derive this adjusted probability measure under various settings of uncertainty and its implications for social valuation, and Section 5 illustrates its impact in a quantitative example.

4.1 Discounting, uncertainty, and pricing

Our analysis shows how an asset pricing perspective adds new twists to the environmental economics literature. Discussions of the questions “what should the discount rate be for social valuation?” have been extensive in the environmental economics literature to date. This discourse sometimes alludes

to ad hoc uncertainty adjustments. A detailed version of such an exploration is provided in Gollier (2013), including references to ambiguity aversion as a motivation for wanting to alter discount rates. The discussion of discount rates often includes both a subjective discount rate contribution, δ in our model, and a growth rate adjustment. While our formulas for the SCC only include the former, this is because we expressed the costs to discounting in utility units. Had we used instead a consumption numeraire, a consumption growth adjustment would have been present in our analysis as well. But even here, the theory of asset pricing typically uses a *stochastic* discount factor process when there are shocks to the macroeconomy. Differential exposure to these shocks should be discounted in different ways as encoded conveniently in stochastic discount factors. It is perhaps more germane to ask “what should the social *stochastic* discount factor be for social valuation?” Producing interest rate counterparts over alternative horizons depends on both the price of uncertainty and the exposure to that uncertainty, but these adjustments are a feature of the joint properties of the stochastic discounting and the uncertain social costs to be discounted. Consistent with Gollier’s reference to forward rates, the compounding of stochastic discount factors over multiple periods of time can have substantively important valuation consequences giving rise to a potentially important term structure for risk prices.

We next provide an overview of how we incorporate a broad notion of uncertainty into valuation. In a nutshell, our uncertainty measures adds an important dimension to stochastic discounting and the remainder of this section shows how to construct this measure.

4.2 An overview

We purposely limit our exploration of alternative probability measures to those that are “disguised” from the planner and not trivially revealed through observations.¹⁵ Roughly speaking, consider alternative probabilities that can be represented as likelihood ratios. Because we focus on models with Brownian information structures, it is most convenient to use changes of measures familiar in mathematical finance justified mathematically by the Girsanov theorem. As is well known from the theorem, the implied change of probability measure includes a possibly history-dependent drift distortion within the Brownian increment. That is, dW_t under the alternative probability measure can be expressed as

$$dW_t = H_t dt + dW_t^H, \quad (12)$$

where dW_t^H a Brownian increment under the change of measure and $H = \{H_t : t \geq 0\}$ is a history-dependent drift distortion process. The drift distortion

¹⁵ We accomplish this formally by considering only alternative probability measures that are absolutely continuous over finite intervals of time.

allows for considerable flexibility, but this formulation is not “without loss of generality.”¹⁶ It is a restriction enforced by the likelihood ratio formulation.

To implement concerns about misspecification, we necessarily penalize or constrain the corresponding drift distortions. For our alternative ways to depict ambiguity aversion and model misspecification, we show the corresponding adjustments to the Hamilton-Jacobi-Bellman (HJB) equation of the robust social planner. These adjustments introduce a minimization problem to the HJB equation formulation so that the planner solves a max-min, or equivalently a two-player, zero-sum game specified recursively rather than only a maximization problem. The minimization is over alternative probabilities represented conveniently as drift distortions. We then use the minimization problem to construct a specific probability measure that gives the valuation adjustment that we are looking for. For adding specificity, we start by describing more formally the resultant preferences.

4.3 Continuation values

We use continuation values to define the preferences recursively. Continuation values are prospective and computed by solving a forward stochastic differential equation. As in dynamic programming, a terminal value along with a forward-looking evolution equation imply continuation value processes for each hypothetical decision or action process. Looking forward, for Markov decision problems of the type we consider for a social planner, the equation for the continuation value evolution alters the HJB equations previously described.

Let $U = \{U_t : t \geq 0\}$ denote the continuation value process posed in continuous time. Write

$$dU_t = \mu_{U,t} dt + \sigma_{U,t} \cdot dW_t,$$

where a recursive representation of the value function implies the restriction:

$$0 = \mu_{U,t} + v_t - \delta U_t. \tag{13}$$

This representation of preferences translates into an HJB equation once we use the Markov structure and the Ito formula to depict the drift $\mu_{U,t}$ in terms of value function derivatives and the local evolution of the Markov state. For an action or decision process A and value function V , the local dynamic coefficients for the continuation value process are:

$$\begin{aligned} \mu_{U,t} &= \frac{\partial V}{\partial x}(X_t) \cdot \mu_X(X_t, A_t) + \frac{1}{2} \text{trace} \left[\sigma_X(X_t)' \frac{\partial^2 V}{\partial x \partial x'}(X_t) \sigma_X(X_t) \right] \\ \sigma_{U,t} &= \left[\frac{\partial V}{\partial x}(X_t) \right]' \sigma_X(X_t). \end{aligned}$$

The instantaneous utility v_t depends on the action as a function of the state. Optimization leads us to include the maximization as in (4).

¹⁶ Although there are ways to further generalize some of the formulations which follow, these are beyond the scope of this paper.

Under the (local) change of measure captured by (12), this is modified to be:

$$0 = \mu_{U,t} + v_t + \sigma_{U,t} \cdot H_t - \delta U_t. \tag{14}$$

Alternative specifications of aversions to uncertainty will lead us to restrain the drift distortion processes H in different ways.

4.4 Model misspecification

Initially, we explore model misspecification for a single model. Allowing for arbitrary misspecification leads to a degenerate outcome. Instead we consider ways of penalizing distortions using a well-studied construct in the applied probability literature called “relative entropy.” The approach has been used previously in the literature on robust control theory. For instance, see James (1992) for a continuous-time formulation. We use the adaptation and extension by Hansen and Sargent (2001), and Hansen et al. (2006). Anderson, Brock, and Sanstad (2018) used a discrete-time formulation of this approach to study an alternative energy climate model with concerns for model misspecification.

As shown by Hansen and Sargent (2019b), this formulation can be viewed as a special case of the recursive variational decision theory axiomatized by Maccheroni, Marinacci, and Rustichini (2006). This approach introduces a quadratic penalty in (14)

$$0 = \min_{h \in \mathbb{R}^m} \mu_{U,t} + v_t + \sigma_{U,t} \cdot h - \delta U_t + \frac{\xi_m}{2} h \cdot h = \mu_{U,t} + v_t - \delta U_t - \frac{1}{2\xi_m} \sigma_{U,t} \cdot \sigma_{U,t}, \tag{15}$$

where the minimized value is:

$$H_t^* = -\frac{1}{\xi_m} \sigma_{U,t}.$$

Here, ξ_m determines how much the planner is concerned about misspecification. Large values of ξ_m capture low concern about misspecification, while for small values of ξ_m this concern is much more pronounced.

Next, we describe a more structured approach to parameter uncertainty.

4.5 Parameter ambiguity

Dynamic models typically have unknown parameters for which theory and data are only partially informative. Recall from Figure 2, that there is substantial uncertainty in the climate sensitivity parameter β used in the Matthews et al. approximation. Similarly, Figures 3 and 4 illustrate uncertainty in the specification of damages. There may be very little reason to commit to a specific measure of central tendency in the case of Figures 2 and 4 or an arbitrary weighting of the high and low damage specifications in Figure 4 when solving the model. We could perform calculations based on imposing alternative values on the fictitious social planner and check for sensitivity of the analysis. Here,

we suggest an alternative strategy whereby the planner confronts parameter ambiguity and model specification with caution.

Let θ denote a possible parameter configuration unknown to the planner in a set Θ . For each possible parameter realization θ , there is dynamic evolution given by:

$$dX_t = \mu_X(X_t, A_t | \theta)dt + \sigma_X(X_t)dW_t.$$

For a value function V and a decision process $\{A_t : t \geq 0\}$

$$\mu_{U,t}(\theta) = \frac{\partial V}{\partial x}(X_t) \cdot \mu_X(X_t, A_t | \theta).$$

Let $P_t(d\theta)$ be a date t reference prior/posterior over a set of possible values of Θ conditioned on date t information. In a dynamic setting, the distinction between a prior and posterior becomes blurred as “yesterday’s posterior” is “today’s prior”. The values of θ can index unknown parameters or a discrete set of models or both. Rather than fully embrace this posterior, the planner explores deviations. Let $Q_t(\theta)$ be a relative density that satisfies:

$$\int_{\Theta} Q_t(\theta)P_t(d\theta) = 1,$$

which is used to alter the posterior distribution. Let $G_t(\theta)$ be a drift distortion that can depend on the unknown parameter. Then the drift distortion that interests us is an H_t that satisfies

$$\begin{aligned} \sigma_X(X_t)H_t = & \left(\int_{\Theta} [\mu_X(X_t, A_t | \theta) + \sigma_X(X_t)G_t(\theta)] Q_t(\theta)P_t(d\theta) \right) \\ & - \int_{\Theta} \mu_X(X_t, A_t | \theta)P_t(d\theta), \end{aligned} \tag{16}$$

as a possible drift distortion for the Brownian motion. Notice that if Q_t is identically one, then $H_t = \int G_t(\theta)P_t(d\theta)$ solves this equation. Before proceeding, there is one technical restriction that we must impose on how the drift depends on the unknown parameter vector.

Remark 4.1. Recall that we allow for σ_X to be singular (e.g., $m < n$). Instead, we restrict the m by m matrix $(\sigma_X)' \sigma_X$ to be nonsingular. Allowing σ_X to have more rows than columns requires some explanation because there may not exist a solution H_t to the equation. We rule this problem out by presuming that the parameter vector to be fully disguised by the local dynamics. Suppose there is some (potentially conditional) linear combination of the n -dimensional state vector that has locally predictable dynamics for which the Brownian exposure is zero. We restrict the implied drift for this linear combination to be independent of θ . For example, in our model there is no diffusion component to the state dynamics for F . These same dynamics do not depend on an unknown parameter.

To accommodate this structured uncertainty, in restricting the local mean of the continuation value, we now alter minimization problem (15) along the lines suggested in the Hansen and Miao (2018):

$$\begin{aligned}
 0 = & \min_{q, \int q(\theta)P_t(d\theta)=1} \min_{g(\theta) \in \mathbb{R}^m} -\delta U_t + v_t \\
 & + \int_{\Theta} \left[\mu_{U,t}(\theta) + \sigma_{U,t} \cdot g(\theta) + \frac{\xi_m}{2} g(\theta) \cdot g(\theta) \right] q(\theta) P_t(d\theta) \\
 & + \xi_p \int_{\Theta} [\log q(\theta)] q(\theta) P_t(d\theta), \tag{17}
 \end{aligned}$$

where we have penalized the choice of density distortion q with a scaled version of the relative entropy divergence:

$$\int_{\Theta} [\log q(\theta)] q(\theta) P_t(d\theta),$$

which has been used extensively in the applied probability and statistics literature. Letting q be one makes this divergence zero, and letting the parameter ξ_p become arbitrarily large restricts the posterior distortion q to be arbitrarily close to unity.

This minimization has a very tractable quasi-analytical solution, which is important for numerical implementation. The minimizing $g(\theta)$ does not depend on θ and has a solution analogous to that for minimizing h for the model misspecification problem:

$$G_t^*(\theta) = -\frac{1}{\xi_m} \sigma_{U,t}.$$

The minimizing density distortion

$$Q_t^*(\theta) = \frac{\exp\left[-\frac{1}{\xi_p} \mu_{U,t}(\theta)\right]}{\int_{\Theta} \exp\left[-\frac{1}{\xi_p} \mu_{U,t}(\theta)\right] P_t(d\theta)},$$

which tilts the resultant posterior toward θ 's for which the value function drift is relatively low. Substituting these solutions in to the objective in (17) gives:

$$-\delta U_t + v_t - \xi_p \log \int_{\Theta} \exp\left[-\frac{1}{\xi_p} \mu_{U,t}(\theta)\right] P_t(d\theta) - \frac{\xi_m}{2} \sigma_{U,t} \cdot \sigma_{U,t}. \tag{18}$$

Remark 4.2. This approach, absent model misspecification, can be viewed as a continuous-time version of a “smooth ambiguity” model. Klibanoff, Marinacci, and Mukerji (2005) represent uncertainty as a two-stage lottery whereby one stage is used to capture risk conditioned on a model θ , which for us is depicted as a Brownian increment, and another stage to depict ambiguity

over models (indexed by θ). They suppose that there are distinct preference representations of aversions associated with this two-stage lottery. In this paper, we follow Hansen and Miao (2018) in our use of a continuous-time formulation along with the robustness interpretation. To connect our formulation to that of Klibanoff, Marinacci, and Mukerji (2005), notice that the outcome of the minimization problem depicted in (18) includes a term given on the left-hand side of the inequality

$$-\xi_p \log \int_{\Theta} \exp \left[-\frac{1}{\xi_p} \mu_{U,t}(\theta) \right] P_t(d\theta) \leq \int_{\Theta} \mu_{u,t}(\theta) P_t(d\theta).$$

The term on the left is recognizable as the exponential certainty equivalent and less than the posterior mean $\int_{\Theta} \mu_{u,t}(\theta) P_t(d\theta)$. Hansen and Miao (2018) derive this as a continuous-time limit of recursive smooth ambiguity preferences.

Remark 4.3. As an alternative ambiguity adjustment in a continuous-time Brownian setting, Chen and Epstein (2002) propose an instant-by-instant restriction on the potential subjective probabilities $Q_t(\theta)P_t(d\theta)$ assigned to the alternative models. The decision-maker is uncertain about Q_t but instead restricts it to be in the convex set that can be state-dependent. The Chen and Epstein (2002) preference specification is a recursive implementation of the max-min utility formulation axiomatized by Gilboa and Schmeidler (1989). Hansen and Sargent (2019b) motivate state dependence in the date-by-date constraint set as a form of time variation in parameters and show how to construct such an ambiguity set using a refinement of relative entropy. The formulation in Hansen and Sargent (2019b) combines this approach with concerns that each of the models in the ambiguity set might be misspecified. This amalgam is very much analogous to the extension of the smooth ambiguity formulation we proposed here. The asset pricing methods that we describe in what follows are also applicable to the uncertainty averse preferences proposed in Hansen and Sargent (2019b).

4.6 Parameter learning

Learning adds state variables to the analysis. For sufficiently simple examples, there could be sufficient statistics that make learning recursions straightforward and tractable to implement recursively. These sufficient statistics would need to be included among the set of state variables and the drift distortions to the underlying Brownian motion would alter their evolution. Also, depending on what coefficients are uncertain, the choice of action could affect the learning and the social planner problem as we have posed it here, as the social planner might have incentives to “experiment.” To the extent such a channel exists, designing a policy with this incentive in mind would add controversy to the analysis, as it does in macroeconomic policy in other settings.¹⁷ For some key climate

¹⁷ For example, see Cogley et al. (2008) for a discussion of robustness and experimentation in a monetary policy setting with learning.

parameters, learning can happen at best very slowly. In our computations we will omit the learning channel altogether. Although this will substantially simplify our calculations, there are also convincing climate science-related reasons to embrace this approximation. For instance, Roe and Baker (2007) write, “The envelope of uncertainty in climate projections has not narrowed appreciably over the past 30 years, despite tremendous increases in computing power, in observations, and in the number of scientists studying the problem... foreseeable improvements in the understanding of physical processes, and in the estimation of their effects from observations, will not yield large reductions in the envelope of climate sensitivity.” This perspective is consistent with the Bayesian computations of Olson et al. (2012) for what they call the climate sensitivity parameter that we mentioned earlier.

4.7 HJB equation and implications

We now propose a modified HJB equation for the social planner that includes concerns about model misspecification and ambiguity. In light of this evidence of very slow learning, we use a time invariant probability P in place of P_t as an approximation. The value function dynamics given in Equation (17) imply a counterpart HJB Equation to (4) with damages entering preferences (or equivalently scaling consumption):

$$\begin{aligned}
 0 = \max_{a \in \mathbb{A}} \min_{q > 0, \int q P(d\theta) = 1} \min_{g \in \mathbb{R}^m} & -\delta V(x) + \delta(1 - \kappa) [\log(\alpha - i - j) + k - d] + \delta \kappa (\log e + r) \\
 & + \frac{\partial V}{\partial x}(x) \cdot \left[\int_{\Theta} \mu_X(x, a | \theta) q(\theta) P(d\theta) + \sigma_X(x) g \right] \\
 & + \frac{1}{2} \text{trace} \left[\sigma_X(x)' \frac{\partial^2 V}{\partial x \partial x'}(x) \sigma_X(x) \right] \\
 & + \frac{\xi_m}{2} g' g + \xi_p \int_{\Theta} [\log q(\theta)] q(\theta) P(d\theta). \tag{19}
 \end{aligned}$$

See the Online Appendix for more details on our numerical implementation.¹⁸ This max-min problem provides a state-dependent action a^* as well as state-dependent density q^* and a drift distortion g^* . We now show how to use these latter two objects to construct an uncertainty adjusted probability by constructing a corresponding drift for the state dynamics. The ambiguity-adjusted probability over the parameter space Θ is $q^*(\theta | x)P(d\theta)$ and the drift as a function of the Markov state is given by

$$\mu^*(x) = \int_{\Theta} \mu_X[x, a^*(x) | \theta] q^*(\theta | x) P(d\theta) + \sigma_X(x) g^*(x). \tag{20}$$

In Section 3, we represented the external contribution to the social cost of carbon as expected discounted future marginal damages induced by a marginal

¹⁸ We also provide a Jupyter Notebook on <https://github.com/lphansen/Climate> with access to the code for the project and a user interface with more details on the implementation and the resultant accuracy.

change in emissions for all future time periods where the time $t + \tau$ contribution is

$$(1 - \kappa)[\nabla^2 \Gamma](\beta F_{t+\tau})\beta^2 E_{t+\tau} + \delta(1 - \kappa) \left([\nabla \Gamma](\beta F_t)\beta + \zeta_D(Z_t) \cdot \begin{bmatrix} 1 \\ 0 \end{bmatrix} \right)$$

scaled by the marginal utility of consumption. This same logic extends once we incorporate the alternative uncertainty sources, but with qualification. The expectation is now computed using the conditional *ambiguity-adjusted probability measure*. Instead of computing this expectation directly, we may infer it from our ambiguity-adjusted HJB solution to the planner’s problem as

$$ecc^* = \frac{\delta \kappa}{e^* \exp(r)} - \frac{V_r(x)}{\exp(r)},$$

where the right-hand side is the marginal utility emissions minus the private contribution from the value function. As in Section 3.1.2, this follows from the first-order condition for emissions from the planner’s HJB equation. See the Online Appendix for an elaboration.

As an alternative to evaluating the discounted value using the ambiguity-adjusted probability, suppose we use the original unadjusted probabilities to evaluate the expected discounted value of the future marginal social costs. Call this $\overline{ecc}(x)$. We take the difference between the two discounted expected values

$$ucc^*(x) = [ecc^*(x) - \overline{ecc}(x)]$$

divided by the marginal utility of consumption or its damaged counterpart to be the uncertainty component to the SCC of carbon, inclusive of both model ambiguity and model misspecification adjustments.

We compute \overline{ecc} and hence ucc^* as follows:

i) integrate:

$$(1 - \kappa) \int_{\Theta} \left([\nabla \Gamma](\beta f)\beta + \zeta_D(z) \cdot \begin{bmatrix} 1 \\ 0 \end{bmatrix} \right) P(d\theta);$$

ii) integrate:

$$(1 - \kappa) \int_{\Theta} [\nabla^2 \Gamma](\beta f)\beta^2 e^* \exp(r) P(d\theta);$$

iii) solve a Feynman Kac equation to compute the discounted expected value of the future damage flow given in (ii) using the baseline probability measure;

iv) add the solution from part (i) to the solution from part (iii) to form \overline{ecc} .

We apply the analogous approach for the model in which damages alter economic growth. This basic construct is much more generally applicable including to models with richer climate dynamics.

The altered probability is not meant to represent the beliefs of the social planner. This constructed probability gives the planner a way to confront more general forms of uncertainty other than risk. Conveniently, the outcome of our robustness analysis to alternative probabilities can be captured and computed by specifying two parameters that serve as preference parameters for the decision-maker, ξ_p and ξ_m . Although we do not dictate what these should be, we find it revealing to look at the implied ambiguity-adjusted probabilities and the corresponding relative entropies to assess what probabilities are of most concern to the decision-maker.¹⁹

Remark 4.4. Since the writing of Good (1952), robust Bayesians have suggested that an implied “worst-case probability” under which the decision-maker optimizes is worthy of careful inspection. The ambiguity-adjusted probability measure that emerges from the HJB equation is arguably difficult to interpret in this light, because it depends on endogenous state variables. To construct this worst-case probability, we appeal to a result from two-player, zero-sum differential games. Just like in dynamic programming, there is a date zero static game that the HJB equation provides a solution for. Provided that a so-called “Bellman-Isaacs condition” is satisfied, the orders of maximization and minimization can be exchanged as of date zero without altering the implied value to the game. See Fleming and Souganidis (1989) for a formal discussion. To compute the worst-case probability, exchange orders in the static game by first maximizing conditioned on the probability and then minimizing over probabilities subject to penalization. The outcome of this static minimization with the order of extremization reversed gives the worst-case probability from a robust Bayesian perspective. For further discussion, see Hansen et al. (2006).²⁰

Remark 4.5. The term “social cost of carbon” can have different meanings depending on the context. While we featured the Pigouvian taxation interpretation, there is another construct that may be more pertinent to current usage by governments, say as is reflected in the Green Book prepared by HM Treasury (2018). Consider a marginal change in emissions from an existing equilibrium that may not be socially efficient. To formalize this with a similar perspective, we would impose the stochastic evolution of the pertinent economic state variables specified exogenously in our HJB equation formulation. For instance, we could solve for a competitive equilibrium

¹⁹ See Anderson, Hansen, and Sargent (2003), and Anderson, Brock, and Sanstad (2018) for alternative ways to link the parameter ξ_m to entropy measures and to so-called “detection error probabilities” used to assess how statistically close the ambiguity-adjusted probability measure is to the reference or baseline probability.

²⁰ The material in appendix D of their paper is particularly relevant to this topic.

abstracting from climate impacts and then impose the resultant actions on the planner's problem. Instead of computing the action "a" as in HJB Equation (19), we would dispense with the maximization and impose the solution for the action from the competitive problem. We would continue to solve the minimization problem to produce an ambiguity adjusted probability to use for social valuation. With this approach, we would still compute the social marginal rate of substitution of emissions and consumption as an alternative measure of the social cost of carbon. This cost also can be represented as the valuation of a social cash flow for the implied economic damages using the ambiguity adjusted probability measure from the altered HJB equation.

5. An Illustration

In this section, we illustrate our analysis. To provide a basic understanding of the economic model, we start by investigating a steady-state version of our model without climate impacts. Given the homogeneity imposed, this version of the model possesses a steady state in the appropriate ratios of variables. This was by design. We use these relations to gain an initial understanding of our baseline parameter configuration and to set the stage for assessing how the efficient allocation is altered by incorporating the climate externality. We then we introduce a climate/damage externality and show how uncertainty alters emissions and the social cost of carbon. As we will illustrate, the damage specification acts similarly to a Hotelling-like constraint on emissions.

5.1 Steady state without climate impacts

To illustrate "how the model works" we start with a deterministic version of the model without damages and investigate the steady-state implications.

Table 1 lists the technology and preference parameters, and Table 2 gives the steady-state values associated with our parameters. The economic model at this level of abstraction is difficult to calibrate in a fully convincing way. Thus, this table is not the outcome of a formal moment matching approach sometimes used in the macro calibration literature. In addition to its simplicity, the notion of capital in our setup should be broad based in including human capital and forms of intangible capital in addition to physical capital. Similarly, the reserves in our models could include both oil and coal.²¹ See the Online Appendix for more details.

The emissions trajectory implicit in this fixed point ignores the climate externality in perpetuity, so the outcome essentially will be to "fry the planet." Absent climate impacts, by design our model has sufficient homogeneity

²¹ We formally imposed two steady-state targets in our parameter settings, one on the reserves to capital ratio and the other on the growth rate of capital. Had we not included the possibility of investment for the discovery of new reserves, we would have been led to a rather different "calibration strategy," including some speculation about a substantially larger stock of "potential reserves."

Table 1
Technology (top) and preference (bottom) configurations.

Parameter	Value
α	.115
ϕ_1	16.7
ϕ_0	.060
$\bar{\mu}_K$	-.035
ψ_0	.113
ψ_1	.143
δ	.010
κ	.032

Table 2
Steady states for the model specification without climate impacts. The values with a superscript ^a were imposed when setting the parameters.

Variable	Value
Investment/capital ^a : i	.090
Growth rate of capital ^a : η	.020
Marginal value of capital ^a : π	2.50
Emissions/reserves ^a : e	.015
Reserves/capital ^a : $\exp(r - k)$.980
Exploration/capital: j	2.72×10^{-4}
Consumption/capital: c	.0247
Marginal value of reserves: ρ	.0545

whereby there is steady growth implying a fixed point in ratios. Under the Matthews et al. (2009) approximation, temperature will grow without bound. In the competitive steady state associated with our parameter settings, emissions grow at 2% while the subjective discount factor is 1%. This implies that log damages will grow at roughly 4% given our quadratic specification of log damages. This means that the discounted future social costs will be infinite at the deterministic steady state. The solution to the social planner’s problem will avoid this extreme outcome as it will be desirable to limit the growth of emissions and keep the damage integral finite.

5.2 Consequences of climate and damage uncertainty

Our first set of results are computed in a stochastic version of the model²² using the smooth ambiguity specification of preferences applied to both climate sensitivity and to the damage uncertainty depicted in Figure 3. In particular, we make the following modeling simplifications:

- i) $\xi_m = \infty$,
- ii) $\zeta_D(Z_t) \cdot \begin{bmatrix} F_t \\ 1 \end{bmatrix} = \zeta_{D,2}(Z_t)$.

²² See the Online Appendix for more details on the volatility parameters.

In regards to item (i), we do not mean to diminish the importance of model misspecification and plan to do comparative analysis of the distinct consequences of both uncertainty components in future research. We impose the restriction in item (ii), to simplify computation, though it also removes a potentially interesting source of variation for emissions. Moreover, as we discussed in Section 4.4, activating both would open an interesting additional channel for model misspecification concerns to affect prudent climate/economics policy.

As we discussed previously, associated with this ambiguity adjustment are altered probabilities assigned to the alternative damage specifications and altered densities for the climate sensitivity parameter β . As we see no easy way to give a “primitive interpretation” for the magnitude of the smooth ambiguity parameter ξ_p , we instead look at the distributional consequences of this parameter setting. With this in mind, we begin by looking at the implied densities and probabilities.

We start by assigning baseline probabilities of one half to each of the damage specifications. Once we introduce damages, there is no even approximate stochastic steady state of interest. As a result, this induces state dependence in the worst-case or adjusted probabilities that is prominently reflected in the dynamic evolution of state variables. The dependence on the state variable f that measures cumulative emissions turns out to have a particularly pronounced impact on the worst-case densities. The altered probabilities become greater as the emissions trajectories push towards relatively higher damages towards the region where the two damage specifications depicted in Figure 3 diverge. This pattern is evident in the second column of Table 3, where we report entropies for a deterministic path simulated from the state initialization that matches the steady states from the competitive model without climate impacts. The entropies only start to have notable distortions on this path 50 years out. Prior to this date, altering probabilities has little impact on the decision problem because the two damage specifications agree. The simulated path for the state variables is from the solution to the planner’s problem in which emissions are relatively modest. Exposure to large environmental degradation is delayed until well into the future under this trajectory.

Figure 5 depicts the distorted climate sensitivity densities that condition on each of the damage function specifications. This figure gives three densities for the climate sensitivity parameter β . One reproduces the normal approximation from Figure 2 and the other two are the ambiguity adjusted densities conditioned on each of the two damage specifications. These are shifted to the right to capture the caution implicit in the ambiguity adjusted probabilities. The distortions are notably larger conditioned on the high-damage specification, which is to be expected. The high damage specification is of most concern to the planner while the adjusted weights reported in Table 3 even up to 100 years are modest. Conditioned on the high damage specification the adjusted density for β loads up probability in the right-tail with the second mode of the density becoming more prominent.

Table 3
 Entropies relative to the baseline normal density with a mean of 1.73 and a standard deviation of .493. For the “weighted damage” specification, the baseline probabilities are one half for each damage specification in Figure 3. The implied worst-case probabilities for the low damage specification are given in parentheses. For the “low damage” specification, probability one is placed on the low damage specification. The worst-case means and standard deviations are reported in parentheses. For the “high damage” specification, probability one is placed on the high damage specification. The value used for ξ_p is $\frac{1}{4000}$.

Year	Weighted damage (low damage prob)	Low damage (mean, SD)	High damage
0	.005 (.50)	.010 (1.80, .502)	.004
25	.010 (.50)	.032 (1.86, .510)	.008
50	.026 (.50)	.054 (1.89, .515)	.018
75	.112 (.46)	.071 (1.91, .518)	.087
100	.197 (.42)	.084 (1.93, .520)	.162




Figure 5
 Probability densities for the climate sensitivity parameter. The blue solid curves represent the baseline probability density, the red dot-dashed curves represent the ambiguity-adjusted density conditioned on the low damage model, and the green dashed curves represent the ambiguity-adjusted densities conditioned on the high damage model.

Figure 6 plots the implied social cost of carbon over a 100-year time horizon. This figure also includes a contribution that quantifies the impact of the uncertainty-adjusted probability measure. The private contribution to this cost is relatively speaking, very small and can safely be ignored. In contrast, the uncertainty component is substantial and accounts for roughly half of the social cost of carbon for this example. Not surprisingly, given our depiction of the adjusted densities in Figure 5, the relative importance of the uncertainty adjustment (as well as the cost itself) becomes more prominent at say 100 years out than at zero. The units are 2010 U.S. dollars per ton of carbon.

Figure 7 gives two emissions trajectories, one computed when we abstract from ambiguity aversion and the other from the same social planner’s problem as was used in the Table 3 and Figure 5. Both trajectories decay much like in a Hotelling exhaustible resource allocation problem. However, this outcome is not induced by the potential exhaustion of the resource because our model allows for investment in new reserves. Instead, the potential for severe damages




Figure 6
 Social cost of carbon decomposition. The units are 2010 U.S. dollars per ton of carbon. The costs are computed at the socially efficient allocation. The blue solid curve represents the total social cost of carbon. The private contribution is negligible relative to the other components and is not plotted. The red dashed curve represents the uncertainty contribution.




Figure 7
 Emissions comparison. The figure reports emissions paths under ambiguity aversion (blue solid line) and ambiguity neutrality (red dashed line). In each case, the socially efficient allocations are used under the respective ambiguity preferences.

restrain the emissions for the fictitious planner because of the presence of the climate externality.²³ While the curves in Figure 6 hold fixed the emissions and

²³ Note that the initial value of emissions is actually higher here than in our steady-state setting that ignores climate impacts. This finding emerges because the initial decrease in the marginal social value of holding reserves

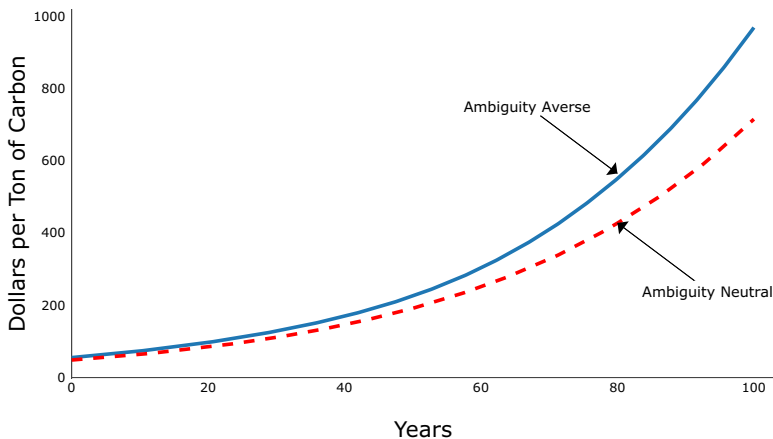


Figure 8

Social cost of carbon trajectories computed under ambiguity aversion (blue solid line) and under ambiguity neutrality (red dashed line). In each case the socially efficient allocations are used under the respective ambiguity preferences. The units are 2010 U.S. dollars per ton of carbon.

other allocations implied by the model, in Figure 8, we report the total social cost of carbon with and without the ambiguity averse preferences. Both trajectories grow like the resource price in a Hotelling model, but not surprisingly, the social cost of carbon is higher when the planner is averse to ambiguity.

We next report results from a “sensitivity to the prior” type analysis familiar in robust Bayesian methods. We change rather substantially the ex ante weights to the two damage specifications by focusing on two extremes. In the first one, we simply embrace the “low damage” specification by assigning probability one to this specification while continuing to focus on climate sensitivity. In the second one, we feature the “high damage” specification by assigning all of the weight on this specification.

In making these comparisons, we hold fixed the parameter ξ_p . Alternatively we might hold fixed relative entropies at perhaps some date and adjust the ξ_p parameter accordingly. This becomes an issue because for the fixed ξ_p the relative entropies differ across damage function specifications as is evident in Table 3. Consistent with the computation we reported earlier, Figure 9 shows that for the “high damage” configuration, the distortions become quite pronounced with a fatter right-hand tail for the climate sensitivity for longer time periods in the future.

Figures 10 and 11 depict the conditioning outcomes for emissions and the social cost of carbon, respectively. The emissions and social cost of carbon trajectories when the ex ante equal weights are used are quite similar to those

increases emissions over that in the steady-state economy. While at the outset this impact offsets the additional climate-induced social costs, it is only a transient phenomenon.

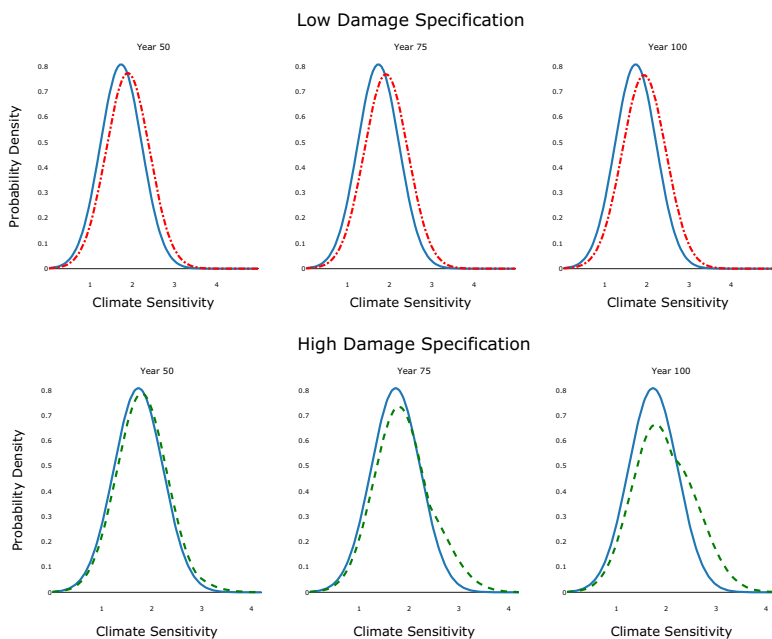


Figure 9 Conditional probability densities for the climate sensitivity parameter. The top panel presumes the low damage specification occurs with probability one, and the bottom panel presumes the high damage specification occurs with probability one. The blue solid curves represent the baseline probability density; the red dot-dashed curves represent the ambiguity-adjusted densities for the low damage specification; and the green dashed curves represent the ambiguity-adjusted densities for the high damage specification.

that emerge when we feature only the high damage specification. In contrast, the emissions trajectory is higher and the social cost of carbon lower when entertaining only the low damage specification. This finding is explicitly tied to our parameter ξ_p . A larger relative entropy penalty pushes the one-half/one-half outcomes closer to an intermediate location. Figure 8 illustrates this for the limiting case in which the ambiguity/robustness parameter is infinite.

To understand the plotted outcomes it is revealing to compare the adjusted probability densities. Of particular interest are the green densities reported in Figure 5 and the corresponding ones reported in the bottom of panel of Figure 9. For instance, consider what happens at year 100. In Figure 5, the density for the climate sensitivity parameter conditioned on the high damage specification is even more substantial than the corresponding curve in the lower panel of Figure 9, where only the high damage specification is entertained by the planner. But in the ex ante one-half/one-half case, the marginal density for the climate sensitivity parameter averages over the two damage specifications and adjustments conditioned on the low damage configuration are much smaller than those that condition on the high damage specification. About 40% of the ambiguity-adjusted weight goes to the low damage specification, making it




Figure 10
Emissions comparison. The values are computed at the socially efficient allocation simulated along a deterministic path. The blue solid curve repeats the trajectory give in Figure 7. The green dashed curve conditions on the high damage specification, and the red dot-dashed curve conditions on the low damage specification.

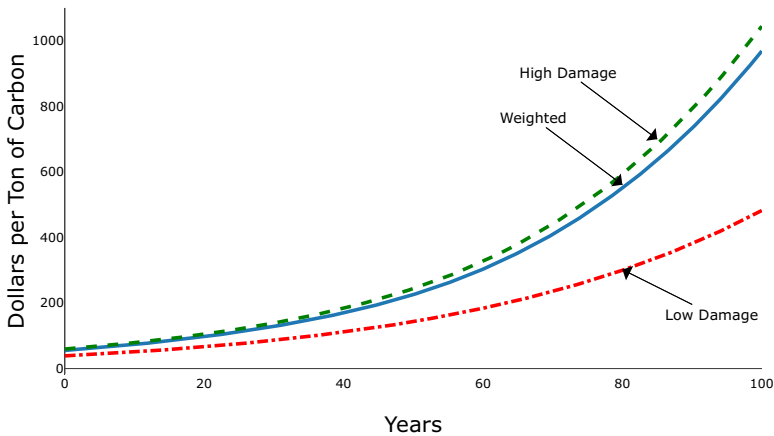


Figure 11
Social cost of carbon comparison. The values are computed at the socially efficient allocations simulated along deterministic paths. The units are 2010 U.S. dollars per ton of carbon. The blue solid curve repeats the trajectory give in Figure 8. The green dashed curve conditions on the high damage specification, and the red dot-dashed curve conditions on the low damage specification.

important in the low damage contribution in the marginal density for the climate sensitivity parameter. More generally, the marginal densities are similar for the different time periods even though the densities conditioned on the high damage specification differ in ways that are quantitatively important. Consistent with the similarities in the ambiguity-adjusted densities, there is an overall similarity in trajectories for both the emissions and the social cost of carbon, as reported in Figures 10 and 11.

Table 4

Emissions and social cost of carbon external and uncertainty contributions. The values are computed at the socially efficient allocations for deterministic pathways. The top panel gives the values at 0, 50, and 100 years for the ambiguity-neutral setting of the growth damages model. The bottom panel gives the values at 0, 50, and 100 years for the ambiguity-averse setting of the growth damages model.

Ambiguity neutral: $\xi_p = \infty$				
Year	Emissions	SCC - total	SCC - uncertainty	Entropy
0	2.4	240	0	0
50	2.0	708	0	0
100	1.8	1,996	0	0
Ambiguity averse: $\xi_p = \frac{1}{175}$				
0	1.4	411	209	.15
50	1.2	1,168	590	.17
100	1.1	3,244	1,638	.19

5.3 Climate change and growth damages

For the macroeconomic growth damage specification, we incorporate estimates of Burke, Davis, and Diffenbaugh (2018) used as in the construction of Figure 4. The results from this growth specification of damages are much more extreme than those displayed in the previous figures. What follows are the impacts observed in emissions and the external and uncertainty contributions to the social cost of carbon.

Table 4 provides the implications for emissions and the social cost of carbon along a simulated deterministic path for 100 years. As before, the initial states for this path match the steady states from the competitive model without climate impacts. For these comparisons, we hold fixed relative entropies at time 100 to be close to those in the consumption damage ambiguity averse setting. Given the specification differences, this compels us to adjust the ξ_p parameter.

The socially efficient emissions are remarkably small and the social cost of carbon remarkably high even under ambiguity neutrality. The uncertainty adjustment is substantial, making the numbers all the more extreme.

As we noted earlier, using growth damages from tropical, underdeveloped regions may well overstate damages to growth for other economies for reasons many economists have discussed (see, e.g., Sachs 2001). We conjecture that, to use this evidence in a more revealing way, it requires explicit regional heterogeneity coupled with a more complete accounting the economic differences in the regions. Distinguishing long-run from short-run growth responses could also change the nature of the evidence as suggested in the earlier work of Dell, Jones, and Olken (2012).²⁴ Hence, we view our growth analysis as a call for more serious probes into the sources and consequences of economic damages.

²⁴ Dell, Jones, and Olken (2012) consider only linear specification for temperature on macroeconomic growth rates. Nonlinearity could well alter their short-run/long-run decomposition.

5.4 Discussion and extensions

We have shown how uncertainty can potentially matter for the social cost of carbon. Our model is very stylized, and our calculations are no doubt sensitive to some of the modeling details. Whenever one engages, like we have, in quantitative storytelling, the outcome is in part about the model and in part about the social problem that it addresses. We explicitly constructed the framework to include multiple “stories.” In what follows, we conjecture about potential extensions of our analysis.

Our social costs of carbon, and in particular, the uncertainty components, are sensitive to the parameter ξ_p . Our particular choice of ξ_p is made for sake of illustration, but by conveniently using relative entropy, we have reduced the ambiguity aversion representation to a single parameter. Instead of being committed to a single parameter value, we may think of our framework as providing a disciplined way to perform a prior/posterior sensitivity analysis for uncertain damage and climate sensitivity parameters indexed by the choice of ξ_p .

The discount rate choice δ will matter as it does in other discussions of climate policy. Changing the subjective discount rate will certainly alter our emissions and cost numbers. Moreover, stochastic discounting in social valuation depends on both the subjective rate of discount in preferences and the ambiguity-adjusted probability measure that we characterized. Along a similar vein, we find it revealing and advantageous to focus on distinct contributions to valuation as well as quantifying their overall impact. While our example economy is special, the decomposition we propose has much more general applicability.

One familiar observation about Hotelling-type models is that as the price rises, backstop technologies become viable, which can give an upper bound on the price. The analogous observation applies in our setting with the potential for green energies to become profitable in the future. While such a technology is absent in our model, extensions that incorporated this will also place a new source of uncertainty and a new channel by which uncertainty affects the economic performance in future time periods. While the model would have to change and the computations would be altered, we suspect that uncertainty, broadly conceived, would continue to play an important role in a quantitative investigation. Relatedly, as carbon presents more of a challenge for society in the future and as technology advances, carbon sequestration may become an attractive form of mitigation. The potential for this and other forms of mitigation to become socially productive would certainly alter our quantitative findings, but they would also open the door to new sources of uncertainty.

While the computations in this section focused on model ambiguity, as we argued earlier in the paper, potential model misspecification is also a concern. This misspecification may be disguised by the Brownian increments making it difficult for the planner to detect model deviations. In future work, we hope to investigate misspecification concerns as a third component to the uncertainty pertinent to climate change.

In this paper, we abstracted from active learning and its impact on the uncertainty adjustments. While learning about carbon sensitivity may be modest in the current environment, if we experience more rapid climate change in the future, learning also could be more pronounced. This is absent from our model, but it could be an important consideration. This form of learning, however, occurs in times of potentially high economic damages making it costly for society to defer action while waiting to learn more. This said, we believe learning to be an interesting extension of our analysis.

6. Impulse-Response Approximation for Climate Dynamics

Recall that a central component to the social cost of carbon is the response function or trajectory for damages to an emissions impulse. A variety of papers in the climate science literature have used transfer function and impulse response methods to approximate the much more complex output that emerges from climate models. This approach aims to provide useful summaries of model implications or syntheses to support tractable emulation and facilitate model comparison. Some examples include Li and Jarvis (2009), Joos et al. (2013), and Castruccio et al. (2014). The Matthews et al. (2009) approximation is a particularly simple version of such a linearized response function. In what follows, we describe some more recent model comparisons that we find to be particularly revealing. These findings suggest further important research should be done that incorporates model uncertainty from climate science and expose further modeling challenges to be faced in embracing this evidence.

Carbon-climate dynamics are often represented in two component parts, the dynamic response of CO_2 concentration to a change in emissions and the dynamic response of temperature to a change in CO_2 concentration via radiative forcing. Combining the two, as in the Matthews et al. approximation, entails a convolution of these response trajectories. Nonlinearity plays a role connecting the two components as it is typically the logarithm of ratio of current concentration to the preindustrial counterpart that determines radiative forcing that is used as an input into the dynamic mapping from CO_2 concentration to temperature. See, for instance, Pierrehumbert (2014).

Impulse response and transfer functions, while pedagogically and computationally convenient, are inherently linear tools of analysis. As discussed in Joos et al. (2013), there is a nontrivial issue over what range of inputs might serve as a good approximation. The impulse response functions that contribute to the social cost of carbon can accommodate nonlinearity by allowing for explicit state dependence in the responses and by calculating local approximations evaluated at the stochastic outcome of the planner's problem. Indeed, a small change in emissions in a nonlinear stochastic system with uncertain random consequences in the future can be pertinent to the social valuation. Given a nonlinear stochastic diffusion evolution, these responses could be computed recursively using what is called the first variation of the

process. Such computations, while they have conceptual appeal, would seem to be tractable only for small scale nonlinear stochastic systems. Perhaps nonlinear emulation methods also would be valuable inputs into studies like ours.

7. Conclusion

We have shown how to apply continuous-time decision theory and asset pricing tools to confront multiple components of uncertainty for the purposes of social valuation. The framework we developed incorporates both concerns about model uncertainty and model misspecification. The resultant methods allow for these broader notions of uncertainty to be integrated formally into decision-making. We apply these tools to study the economic impacts of climate change through the lens of the social cost of carbon.

While the methods are more generally applicable, our example illustrates the impact of the interacting uncertainty components coming from climate and economic modeling. In effect, the impact of these uncertainties is multiplicative: and when both are large, together their impact can be truly substantial. As a result, the social cost of carbon shows notable increases when both sources of uncertainty are acknowledged. This aspect of the analysis is particularly pertinent when the decision-maker is averse to ambiguity over models and to potential model misspecification. Just as risk aversion is theory of “caution,” so too are preference-based concerns about ambiguity and misspecification.²⁵ We believe these components to be particularly relevant for assessing the economic impacts of climate change, and we expect them to be pertinent for social valuation applied in other settings.

We are sympathetic to concerns that readers might have of our seemingly simplistic use of the social cost of carbon. Yet, for the purposes of this paper, the social cost of carbon serves as a metric to guide our assessment of what components of uncertainty are most impactful. The development of richer models of the underlying economy that include research aimed at mitigation or for the development of viable green technologies are appealing extensions of our analysis.

For quantifying the consequences of uncertainty in revealing ways, we suspect that we have scratched the surface so to speak. For purposes of illustration, we have imposed overly simplified specifications of climate and economic dynamics. Moreover, the approximate climate models we consider potentially understate the importance of nonlinearities in the climate dynamics. Within the confines of risk analyses, important research on climate tipping points has been done by Lenton et al. (2008), Cai et al. (2015), Cai, Lenton, and Lontzek (2016), and Cai, Judd, and Lontzek (2017). We suspect that adopting

²⁵ Even for financial markets, what is called risk aversion may be better conceived as investor concerns about these other components to uncertainty. For example, see Hansen and Sargent (2019a).

a broader perspective on uncertainty could contribute productively to this line of research.

References

- Allen, M. R., D. J. Frame, C. Huntingford, C. D., Jason J.A. Lowe, M. Meinshausen, and N. Meinshausen. 2009. Warming caused by cumulative carbon emissions towards the trillionth tonne. *Nature* 458:1163–66.
- Anderson, E. W., L. P. Hansen, and T. J. Sargent. 2003. A quartet of semigroups for model specification, robustness, prices of risk, and model detection. *Journal of the European Economic Association* 1:68–123.
- Anderson, E. W., W. A. Brock, and A. Sanstad. 2018. Robust consumption and energy decisions. In *Climate change economics: The role of uncertainty and risk*, eds. V. V. Chari and R. Litterman. Hoboken, NJ: Wiley.
- Bansal, R., and A. Yaron. 2004. Risks for the long run: A potential resolution of asset pricing puzzles. *Journal of Finance* 59:1481–509.
- Bansal, R., D. Kiku, and M. Ochoa. 2017. Price of long-run temperature shifts in capital markets. Working Paper.
- Berger, J. O. 1984. The robust Bayesian viewpoint. In *Robustness of Bayesian analysis*, vol. 4, ed. J. B. Kadane, 63–124. North-Holland, Amsterdam.
- Bornstein, G., P. Krusell, and S. Rebelo. 2017. Lags, costs, and shocks: An equilibrium model of the oil industry. Working Paper.
- Brock, W. A., and L. P. Hansen. 2018. Wrestling with uncertainty in climate economic models. In *Climate change economics: The role of uncertainty and risk*, ed S. V. V. Chari and R. Litterman. Hoboken, NJ: Wiley.
- Burke, M., W. M. Davis, and N. S. Diffenbaugh. 2018. Large potential reduction in economic damages under UN mitigation targets. *Nature* 557:549–53.
- Burke, M. B., S. M. Hsiang, and E. M. Miguel. 2015. Global non-linear effect of temperature on economic production. *Nature* 527:235–39.
- Cai, Y., K. L. Judd, and T. S. Lontzek. 2017. The social cost of carbon with climate risk. Technical Report, Hoover Institution, Stanford, CA.
- Cai, Y., K. L. Judd, T. M. Lenton, T. S. Lontzek, and D. Narita. 2015. Environmental tipping points significantly affect the cost-benefit assessment of climate policies. *Proceedings of the National Academy of Sciences* 112:4606–611.
- Cai, Y., T. M. Lenton, and T. S. Lontzek. 2016. Risk of multiple interacting tipping points should encourage rapid CO₂ emission reduction. *Nature Climate Change* 6:520–25.
- Casassus, J., P. Collin-Dufresne, and B. R. Routledge. 2018. Equilibrium commodity prices with irreversible investment and non-linear technologies. *Journal of Banking and Finance* 95:128–47.
- Castruccio, S., D. J. McInerney, M. L. Stein, F. L. Crouch, R. L. Jacob, and E. J. Moyer. 2014. Statistical emulation of climate model projections based on precomputed GCM runs. *Journal of Climate* 27:1829–44.
- Chen, Z. and L. Epstein. 2002. Ambiguity, risk, and asset returns in continuous time. *Econometrica* 70:1403–43.
- Cogley, T., R. Colacito, L. P. Hansen, and T. J. Sargent. 2008. Robustness and U.S. monetary policy experimentation. *Journal of Money, Credit and Banking* 40:1599–623.
- Colacito, R., B. Hoffmann, and T. Phan. 2019. Temperature and growth: A panel analysis of the United States. *Journal of Money, Credit and Banking* 51:313–68.
- Dell, M., B. F. Jones, and B. A. Olken. 2012. Temperature shocks and economic Growth: Evidence from the last half century. *American Economic Journal: Macroeconomics* 4:66–95.
- Diaz, D., and F. Moore. 2017. Quantifying the economic risks of climate change. *Nature Climate Change* 7:774–82.

- Eberly, J. C., and N. Wang. 2009. Reallocating and pricing illiquid capital: Two productive trees. AFA 2010 Atlanta Meetings Paper.
- Fleming, W. H., and P. E. Souganidis. 1989. On the existence of value functions of two-player, zero-sum stochastic differential games. *Indiana University Mathematics Journal* 38:293–314.
- Gilboa, I., and D. Schmeidler. 1989. Maxmin expected utility with non-unique prior. *Journal of Mathematical Economics* 18:141–53.
- Gollier, C. 2013. *Pricing the planet's future: The economics of discounting in an uncertain world*. Princeton, NJ: Princeton University Press.
- Golosov, M., J. Hassler, P. Krusell, and A. Tsyvinski. 2014. Optimal taxes on fossil fuel in general equilibrium. *Econometrica* 82:41–88.
- Good, I. J. 1952. Rational decisions. *Journal of the Royal Statistical Society, Series B (Methodological)* 14:107–14.
- Hambel, C., H. Kraft, and E. Schwartz. 2018. Optimal carbon abatement in a stochastic equilibrium model with climate change. Working Paper.
- Hansen, L. P., and J. Miao. 2018. Aversion to ambiguity and model misspecification in dynamic stochastic environments. *Proceedings of the National Academy of Sciences* 115:9163–68.
- Hansen, L. P., and T. J. Sargent. 2001. Robust control and model uncertainty. *American Economic Review* 91:60–66.
- . 2007. Recursive robust estimation and control without commitment. *Journal of Economic Theory* 136:1–27.
- . 2019a. Macroeconomic uncertainty prices when beliefs are tenuous. Working Paper.
- . 2019b. Structured uncertainty and model misspecification. Working Paper.
- Hansen, L. P., T. J. Sargent, G. Turmuhambetova, and N. Williams. 2006. Robust control and model misspecification. *Journal of Economic Theory* 128:45–90.
- HM Treasury. 2018. *The green book: Central government guidance on appraisal and evaluation*. London: National Archives.
- Hotelling, H. 1931. The economics of exhaustible resources. *Journal of Political Economy* 39:137–75.
- Hurwicz, L. 1966. On the structural form of interdependent systems. In *Logic, methodology and philosophy of science proceeding of the 1960 International Congress*, vol. 44, eds. P. S. E. Nagel and A. Tarski, 232–39. New York: Elsevier.
- James, M. R. 1992. Asymptotic analysis of nonlinear stochastic risk-sensitive control and differential games. *Mathematics of Control, Signals and Systems* 5:401–17.
- Jensen, S., and C. Traeger. 2014. Optimal climate change mitigation under long-term growth uncertainty: Stochastic integrated assessment and analytic findings. *European Economic Review* 69:104–25.
- Joos, F., R. Roth, J. S. Fuglested, 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* 13:2793–825.
- Keane, M. P., and T. Neal. 2018. The impact of climate change on U.S. agriculture: The roles of adaptation techniques and emissions reductions. Working Paper, University of New South Wales.
- Klibanoff, P., M. Marinacci, and S. Mukerji. 2005. A smooth model of decision making under uncertainty. *Econometrica* 73:1849–92.

- Kreps, David M. and Evan L. Porteus. 1978. Temporal resolution of uncertainty and dynamic choice. *Econometrica* 46:185–200.
- Lemoine, D., and C. P. Traeger. 2016. Ambiguous tipping points. *Journal of Economic Behavior & Organization* 132:5–18.
- Lenton, T. M., H. Held, E. Kriegler, J. W. Hall, W. Lucht, S. Rahmstorf, and H. J. Schellnhuber. 2008. Tipping elements in the Earth's climate system. *Proceedings of the National Academy of Sciences* 105:1786–93.
- Li, S., and A. Jarvis. 2009. Long run surface temperature dynamics of an A-OGCM: The HadCM3 4 x CO₂ forcing experiment revisited. *Climate Dynamics* 33:817–25.
- Maccheroni, F., M. Marinacci, and A. Rustichini. 2006. Dynamic variational preferences. *Journal of Economic Theory* 128:4–44.
- MacDougall, A. H., N. C. Swart, and R. Knutti. 2017. The Uncertainty in the transient climate response to cumulative CO₂ emissions arising from the uncertainty in physical climate parameters. *Journal of Climate* 30:813–27.
- Matthews, H. Damon, N. P. Gillett, P. A. Scott, and K. Zickfeld. 2009. The proportionality of global warming to cumulative carbon emissions. *Nature* 459:829–32.
- Millner, A., S. Dietz, and G. Heal. 2013. Scientific ambiguity and climate policy. *Environmental and Resource Economics* 55:21–46.
- Morgan, M. Granger, P. Vaishnav, H. Dowlatabadi, and I. L. Azevedo. 2017. Rethinking the social cost of carbon dioxide. *Issues in Science and Technology* 33.
- Moyer, E. J., M. D. Woolley, N. J. Matteson, M. J. Glotter, and D. A. Weisbach. 2014. Climate impacts on economic growth as drivers of uncertainty in the social cost of carbon. *Journal of Legal Studies* 43:401–25.
- Nordhaus, W. 2018. Projections and Uncertainties about climate change in an era of minimal climate policies. *American Economic Journal: Economic Policy* 10:333–60.
- Nordhaus, W. D. 2017. Revisiting the social cost of carbon. *Proceedings of the National Academy of Sciences* 114:1518–23.
- Nordhaus, W. D., and A. Moffatt. 2017. A survey of global impacts of climate change: Replication, survey methods, and a statistical analysis. Technical Report.
- Olson, R., R. Sriver, M. Goes, N. M. Urban, H. D. Matthews, M. Haran, and K. Keller. 2012. A climate sensitivity estimate using Bayesian fusion of instrumental observations and an Earth system model. *Journal of Geophysical Research Atmospheres* 117:1–11.
- Pezzey, J. C. V. 2019. Why the social cost of carbon will always be disputed. *Wiley Interdisciplinary Reviews: Climate Change* 10:e558.
- Pierrehumbert, R. T. 2014. Short-lived climate pollution. *Annual Review of Earth and Planetary Science* 42:341–79.
- Roe, G. H., and M. B. Baker. 2007. Why is climate sensitivity so unpredictable? *Science* 318:629–32.
- Sachs, J. D. 2001. Tropical underdevelopment. Technical Report.
- Weitzman, M. L. 2012. GHG targets as insurance against catastrophic climate damages. *Journal of Public Economic Theory* 14: 221–44.