

Wrestling with Uncertainty in Climate Economic Models[‡]

William Brock^{*}

Lars Peter Hansen^{**}

July 24, 2017

Abstract

This paper uses insights from decision theory under uncertainty to explore research challenges in climate economics. We embrace a broad perspective of uncertainty with three components: risk (probabilities assigned by a given model), ambiguity (level of confidence in alternative models), and misspecification (potential shortfalls in existing models). We survey recent climate science research that exposes the uncertainty in climate dynamics that is pertinent in economic analyses and relevant for using models to provide policy guidance. The uncertainty components and their implications for decision theory help us frame this evidence and expose the modeling and evidential challenges.

[‡]We thank Mike Barnett, Amy Boonstra, Diana Petrova and Bob Litterman for helpful comments on an earlier draft.

^{*}University of Wisconsin and University of Missouri, Columbia

^{**}University of Chicago

1 Introduction

Many economic policies aim to improve economic outcomes by confronting externalities missed by market allocations. The human impact on the climate is one such externality. In such circumstances, economists have long advocated cost-benefit analysis to help in the design of prudent policy making. This language is now embraced by governmental agencies and underlies measures of what is referred to as the social cost of carbon.

There are at least two ways to define and measure the social cost of carbon. One way is to measure the long-term damages avoided by small decreases in carbon dioxide (CO_2) emissions in a given year. Given the marginal nature of the conceived change, it is possible to use observed market prices associated with existing and perhaps socially inefficient resource allocations as inputs into this measurement. Another way is to compute Pigouvian taxes as part of the implementation of an allocation that is socially efficient. Pigouvian tax rates are also often represented in terms of shadow prices, but these prices may be distinct from observed prices since the shadow prices that are pertinent to Pigouvian taxation are associated with the social optimum. Since both approaches rely on marginal analysis, their representations are very similar even if the measurements of the pricing inputs can be very different.¹

Conceptually, we find it instructive to represent the social cost of carbon under both approaches as having three ingredients:²

- i) impulse responses from changes in emissions to adverse climate outcomes such as temperature changes in future dates;
- ii) shadow prices of the climate outcome relative to a consumption numeraire;
- iii) stochastic discount factor process that assigns (shadow) values to the numeraire consumption good at different dates in the future and different realized states.

In terms of input i), economic dynamics has long embraced impulse response function analysis at least since the fundamental paper by Frisch (1933). These impulse responses trace out implications of a change today on economically relevant variables in the future. For instance, the change today could be human induced CO_2 emissions and the economically

¹The failure to use price inputs that are the marginal rates of substitution evaluated at the efficient allocation sometimes has small consequences as reported in the first two rows of Table 1 in Nordhaus (2017).

²See, for instance, Golosov et al. (2014) equation 9.

relevant variables in the future could be economic damages. To apply impulse response analysis appropriately to such calculations requires the measurement of the stochastic nature of the responses and explicit inputs from climate science. The relative price in ii) expresses damages in a chosen numeraire. Since the responses are dynamic, we must assign value weights using market or shadow prices depending upon the horizon. Given the intertemporal nature of the valuation, it is *stochastic discounting* featured in iii that matters for this valuation. Stochastic discount factors both discount the future and adjust for risk. The stochastic contribution to valuation, while well understood from the literature on asset pricing, is typically abstracted from or treated with some naïveté in governmental computations.

Local analysis gains simplicity in part because prices or shadow prices are presumed invariant to small changes, but they are hardly the complete answer to relevant policy questions. For instance, changes in human-induced CO_2 emissions may respond dynamically to local policy changes aimed at climate change mitigation. For a more complete assessment of climate change mitigation, it is important to assess the aggregate global consequences. One might hope in a sufficiently linear world that the local analysis computed at the suboptimal allocation is scalable. But the rationale for holding prices, including stochastic discount factors fixed, breaks down. As is well appreciated in the literature on dynamic stochastic general equilibrium models, global analysis of alternative macro policies must account, at least implicitly, for the impact on prices or shadow prices used to make comparisons across alternative time horizons. Alternatively, we might aim to use fully specified and calibrated models to quantify the social benefits of alternative (other than fully implement Pigouvian tax) policies using local or global calculations. Indeed this is a central rationale for the computational approach pioneered by Cai et al. (2015) for climate economic models. Their approach to dynamic stochastic equilibrium analysis is by design well suited for global assessment of more general climate mitigation policies.

To undertake local or global cost-benefit analyses of the types we describe requires that we quantify the prospective costs and benefits potentially in a probabilistic manner. This is not just a challenge that we can simply hand over to economic empiricists skilled at measuring some of the potential economic damages as a result of temperature changes. Decades ago, Koopmans (1947) in his famous critique of the Burns and Mitchell (1946) empirical characterization of business cycles, articulated the limited nature of conclusions that follow measurement without theory. Koopmans was particularly interested in implications for economic dynamics, as are we. Addressing policy challenges with intertemporal

implications, including in climate science and economics, requires structural models. It is not sufficient to rely on historical and cross-sectional evidence. For instance, historical evidence alone on temperature dynamics will miss some of the potentially adverse consequences in the future of the human impact on the climate. In particular, climate and economics models are essential components to our understanding the dynamic mapping between CO_2 changes induced by humans, the resulting changes in temperature and subsequent damages. Reliance on explicit models becomes essential with inputs from both economics and climate science.

Our focus in this essay is on the uncertainty in dynamics implied by economic and climate systems and its potential consequences for valuation. While we are not committed to the social cost of carbon, as a single target of measurement, however conceived, as a single target of measurement, we do find the categorization of the modeling ingredients i), ii), and iii) to be valuable as a device for organizing our discussion. We recognize important steps by a variety of researchers in building dynamic stochastic equilibrium models, but these models require inputs from climate science. Much of our essay explores these latter inputs because they are the ones of which we and many other economists have the most superficial understanding.

We find it important to *feature* uncertainty rather than to *diminish* its importance. We are certainly not the first to recognize the challenge in quantifying uncertainty, but we want to underscore contributing to uncertainty. Uncertainty has implications for both transmissions and intertemporal valuation. Indeed some important steps have already been taken in the literature to quantify the impact of uncertainty. Dynamic models in economics including those with climate components often feature random shocks as a source of uncertainty. But in reality there are other sources. For instance, we may not know parameters within a given model or we may not know which among alternative model specifications gives the best or most reliable answers to the questions at hand. As the models are dynamic and the potential consequence of CO_2 emissions induced by humans play out over long periods of time, the impact of random shocks and model uncertainty can compound over time. Finally, the quantitative models used for computation are approximations and necessarily misspecified.

Our essay discusses the modeling challenges that emerge as we expand the notion of uncertainty to address better the ramifications of the limits to our understanding of climate economic systems. In our exploration of uncertainty, we deliberately choose to adopt a broad perspective on uncertainty and one that is well beyond the risk analysis most fa-

miliar to economists. We find this broad perspective to be particularly germane for models linking economic behavior and climate change. In section 2 we take inventory of some of the evidence from climate science and use this to motivate why a broad perspective on uncertainty is appropriate. This evidence bears on our understanding of impulse response ingredient i). We find decision theory to be a valuable tool to help us understand better the consequences of uncertainty. In section 3 we review some tractable and defensible approaches to uncertainty including ones that target explicitly model ambiguity and potential model misspecification. Decision theory provides a means to integrate concerns about uncertainty into policy analysis using explicit models and to capture the behavior of people and other economic entities within those models who cope with uncertainty. In section 4 we investigate the stochastic discounting ingredient iii) by exploring the ramifications of stochastic discounting in environments with long-term contributions to uncertainty. We consider pedagogically revealing models of dynamic economies to illustrate how stochastic discounting impacts valuation over multiple horizons and is pertinent in present-value analysis. In section 5 we discuss the challenges in quantifying long-term uncertainty by reviewing some of the existing historical evidence and evidence from simulating the output of potentially complex climate models. In section 6 we discuss briefly some spatial models and measurement research to which a more ambitious uncertainty analysis can be contributed in the future.

2 Components of Uncertainty

Following Knight (1921), Arrow (1951) and others, we take a broad perspective on uncertainty and explore the resulting implications. For pedagogical purposes, we distinguish three forms of uncertainty:

- *risk* – what probabilities does a specific model assign to events in the future?
- *ambiguity* – how much confidence do we place in each model?
- *misspecification* – how do we use models that are not perfect?

2.1 Risk

In this essay we use the term risk to represent the uncertainty about outcomes in contrast to uncertainty about probabilities. Specifically, risk captures the probabilities implied by

a model where we use the term model to include knowledge of parameters. Shocks with specific distributions provide a source of risk within the model. Outcomes that are decided by random draws from these distributions and are unknown, but their probabilities are known. Perhaps the simplest example is a random draw from an urn with a known fraction of red balls and white balls.

Many dynamic economic models have specific shock processes with well specified probability distributions and transmission mechanisms. See, for instance, Frisch (1933) for an initial study of the impact of random shocks in a dynamic economic system and Lucas and Prescott (1971) for an initial equilibrium model, and Brock and Mirman (1972) for an initial stochastic growth model with an explicit role for random impulses. This approach to model building has been developed more fully in the literatures on real business cycle models and on dynamic stochastic general equilibrium models. The shocks and their transitions to economic outcomes are examples of risk. Solving the model implies characterizing probabilities of economic outcomes.

Climate models often have a much higher degree of complexity with a less featured role for random shocks. These models range from simple energy balance models to Atmospheric Ocean General Circulation Models (AOGCM's), Regional Climate Models (RCM's), and General Circulation Models (GCM's) containing large numbers of partial differential equations at horizontal grid spatial resolutions as fine as 10 km. See Prein et al. (2015) for a recent discussion of the varied approaches to climate modeling. Even in the most recent attempts to model deep convective processes with spatial resolution down to 4 km at large expenses in computer time, the fact that convective processes take place at a range of scales including scales smaller than 4km leaves a lot of uncertainty residing in unresolved physics. Indeed random impulses do play a central role in the characterization of some of these models. For instance, see North and Cahalan (1981) for an example of a climate model with random impulses as part of the forcings.

2.2 Ambiguity

Ambiguity refers to the uncertainty associated with how to weight alternative models. For instance, we may not know the number of red and white balls in an urn. Not only do we fail to know outcomes, but we may fail to know probabilities. We may be prevented from figuring the fractions of each type of ball with repeated observations because the fractions themselves may change over time, and we are left chasing a moving target. The models we

use in economics and climate science of course are much more complex than urns, but the same issue of how to weight alternative models remains present perhaps in more subtle or complex ways.




Figure 1: Simulations and projections of annual mean GMST 1986-2050 (anomalies relative to 1986-2005). Projections under all RCPs from CMIP5 models (grey and coloured lines, one ensemble member per model), with four observational estimates: Hadley Centre/Climate Research Unit gridded surface temperature data set 4 (HadCRUT4); European Centre for Medium range Weather Forecast (ECMWF) interim reanalysis of the global atmosphere and surface conditions (ERA-Interim); Goddard Institute of Space Studies Surface Temperature Analysis (GISTEMP); National Oceanic and Atmospheric Administration (NOAA): for the period 1986-2012 (black lines). Source: IPCC Climate Change 2013: The Physical Science Basis, Kirtman et al. (2013), Figure 11.25a.⁴

In practice, there are typically multiple models under consideration. For instance, we may not know which among a discrete family of models generates the data that we have observed or will observe going forward. We capture a simplistic but revealing characterization of this uncertainty with Figure 1. This figure shows in a rather dramatic way the consequences of model uncertainty as reflected in the different model projections of climate change out to 2050. Illustrated there is substantial model uncertainty associated in longer-term forecasts. This longer-term uncertainty reflects both modeling differences and uncertainty as to a possible trajectory of fossil fuel emissions in the future. Representative Concentration Pathway (RCP) scenarios are the result of a collaborative effort of the climate research community to build data sets covering the period 1850 to 2100 for four

⁴Web access: https://www.ipcc.ch/pdf/assessment-report/ar5/wg1/WG1AR5_Chapter11_FINAL.pdf and <https://www.aps.org/policy/statements/upload/climate-review-framing.pdf>.

different but plausible emission scenario possibilities each indexed by the projected number of watts per square meter ($\frac{W}{m^2}$) of radiative forcing at 2100. RCP's encompass the ranges of radiative forcing in $\frac{W}{m^2}$ projected at 2100 that had been discussed in scenarios in the literature to date. Van Vuuren et al. (2011) give a detailed discussion of the development of the RCPs.⁵ Even model and RCP differences represent only part of the uncertainty that is pertinent in climate economics. There has been much discussion and debate about the broader consequences of temperature on economic and social welfare.

The elegant de Finetti (1937)-Savage (1954) theory of subjective probability resolves this challenge of how to weight through the specification of prior probabilities combined with Bayesian learning. But both scholars recognized the challenge of assigning such probabilities in practice, especially when an extensive array of potentially complex models is under consideration. In climate science there is such an array of models with distinct implications as illustrated in Figure 1. Data may eventually help us to distinguish sharply among these models and this evidence may overshadow the impact of some of the subjective inputs. But this seems far from the case for climate models and for models of economic damages. Moreover, the historical evidence we have is for episodes with only modest climate impacts. To the extent that there are substantial uncertainties about models and their implications that are outside the range of historical experience, we remain in situations in which the data do not swamp the priors in a meaningful way. This leaves open the question of what weight we assign to the alternative models. Ambiguity about this weighting becomes part of any conversation pertaining to climate actions. We care more than just about the best fitting historical model. We want models that can make credible predictions outside the range of historical experience.

In the language of econometrics, we seek so-called “structural models” that allow us to make counterfactual predictions. Marschak (1953), Hurwicz (1966) and Lucas (1976) discuss what it means for a model of an interdependent system to be structural and why it matters. The aim of such a model is to provide predictions outside the range of the historical experience. We like climate models because they incorporate physical laws that extend their credibility beyond the realm of the historical data on climate impacts. They allow us to make predictions for alternative scenarios capturing a potentially wide array of human inputs. As is apparent from Figure 1, the modeling details in applying physics to climate science matter in important ways giving rise to a range of model outputs when

⁵The 2100 radiative forcing levels of $2.6 \frac{W}{m^2}$ are projections under a slow emission scenario, $4.5 \frac{W}{m^2}$ and $6 \frac{W}{m^2}$ under intermediate emission, and $8.5 \frac{W}{m^2}$ under a “business as usual” large emission scenario.

projecting into the future.

Exploring the sensitivity to alternative weighting schemes is the purview of robust Bayesian statistics. The outcome of this sensitivity analysis is a range of probabilities of future events. See Berger (1984) for an elaboration of this type of analysis, and see Brock et al. (2007) for how to apply this approach to economic policy evaluation. As we will discuss, decision making requires that we confront this range in a meaningful way.

2.3 Model misspecification

The third component to uncertainty is arguably the most challenging to address, but also could have important ramifications. Climate models of interest have the virtue of incorporating explicit physical principles with the aim to be structural in the sense of econometric analysis. The resulting complexity also makes the models harder to use and less transparent in terms of how they work.

Models we use are, by design, approximations. And in some cases they provide rather coarse characterizations or quantifications. Along some dimensions, they are necessarily misspecified. For instance, some observers or critics of climate models take the overstatement of temperature impacts in Figure 1 as evidence of misspecification. Even with the complexity, there remain concerns about model misspecification. For instance, climate models have recently overstated the impact on carbon emissions on temperature, but over an arguably short time horizon. This finding is consistent with each model among a suite of climate models being misspecified because most of the trajectories reported say in Figure 1 run hot relative to the historical record for the period 1986-2012. Fyfe et al. (2016) give a recent discussion of this phenomenon and its ramifications for climate modeling in which they say:

The last notable decadal slowdown during the modern era occurred during the big hiatus. The recent decadal slowdown, on the other hand, is unique in having occurred during a time of strongly increasing anthropogenic radiative forcing of the climate system. This raises interesting science questions: are we living in a world less sensitive to GHG (greenhouse gas) forcing than previously thought, or are negative forcings playing a larger role than expected? Or is the recent slowdown a natural decadal modulation of the long-term GMST (global mean surface temperature) trend? If the latter is the case, we might expect a ‘surge’ back to the forced trend when internal variability flips phase.

While climate models used in generating output plotted in Figure 1 are arguably complex, this complexity has led researchers to explore simpler prototypes to illustrate some of the key mechanisms. Thus the complex models are themselves approximations, and for reasons of tractability we seek further simplifications to the approximations. See Matthews et al. (2009), Matthews et al. (2012), and Pierrehumbert (2014), who suggest a simplified framework for the link between carbon inputs and temperature changes to make communication more transparent and characterizations of uncertainty more tractable. Economists also embrace simplifications for reasons of tractability. The carbon cycle and temperature dynamics in the familiar DICE models are substantial simplifications of more complex depictions of climate science impacts. Despite the transparency of the three reservoir carbon cycle dynamics of the DICE 2013R manual, Glotter et al. (2014) show that neglecting the nonlinearity of ocean uptake as in DICE 2013R leads to biases and understatements of the Social Cost of Carbon (SCC) at time scales relevant to very long term policymaking. Figure 2 reproduces a figure in their paper, which the authors use to suggest a different simplification than DICE. Of interest to us is the comparison of DICE to other climate models and the time scale over which the approximation works well. While the longer time scale in the right-hand side plots would severely stretch the relevance of most economic analyses, the comparisons in the left panel are of more potential interest.




Figure 2: CO_2 (top) and temperature (bottom) anomalies for BEAM and DICE compared to the intermediate complexity models CLIMBER-2 and UVic for the $A2^+$ scenario (all described in Section 5). BEAM CO_2 matches output of the more complex models well for the duration of the simulation. DICE performs well only for the first several decades but then diverges rapidly. The dotted black line shows cumulative emissions (the CO_2 anomaly if no ocean uptake occurred). DICE removes nearly all emitted CO_2 after several hundred years; in more realistic models, half persists for millennia. Source: Figure 1 from Glotter et al. (2014).

Rather than featuring DICE style approximations, recent developments in climate science suggest a different approximation that is of potential value in framing policy discussions. Specifically, Matthews et al. (2009) and Matthews et al. (2012) suggest a “robust approximation” that is valuable in making coarse predictions from a range of more com-

plex models. The argument is well captured by a diagram displayed in Figure 3. This diagram depicts a potential simplicity gained by streamlining the modeling process and focusing on the approximate linkage between emissions and temperature giving rise to the carbon-climate response. Figure 3 illustrates schematically how the response of the earth system to atmospheric carbon including the role of land sinks and ocean sinks in reducing anthropogenic carbon from where it would have been if not for these carbon sinks.




Figure 3: We define ‘carbon sensitivity’ as the increase in atmospheric CO_2 concentrations that results from CO_2 emissions, as determined by the strength of natural carbon sinks. ‘Climate sensitivity’ is shown here as a general characterization of the temperature response to atmospheric CO_2 changes. Feedbacks between climate change and the strength of carbon sinks are shown as the upper dotted arrow (climate-carbon feedbacks). The CCR aggregates the climate and carbon sensitivities (including climate-carbon feedbacks) into a single metric representing the net temperature change per unit carbon emitted. Source: Figure 1 from Matthews et al. (2009).

The rationale for the approximation is based on some offsetting dynamics relating emissions to atmospheric CO_2 concentration and the dynamics relating CO_2 concentration or radiative forcing to temperature increases. This offset is displayed in Figure 4. A temporary increase in emissions has approximately permanent consequences for atmospheric CO_2 with an impact that builds over time. A permanent increase in atmospheric CO_2 will induce an approximate permanent increase in temperature but one that declines over time. The convolution of these two impacts leads to an approximate dynamics whereby an emissions increase leads to a constant limiting increase in temperature that is arguably resolved over a shorter time scale. But this is only part of the story. This is captured by the carbon-climate response (CCR) plotted in Figure 4. There are nonlinearities in both

of these mappings that are approximately offset in yielding a linear relationship.⁶

⁶See Pierrehumbert (2014) for a discussion of this offset.




Figure 4: Idealized model simulations of the CCR. a: Simulation with a 1% per year atmospheric CO_2 increase for 70 years, showing temperature change per unit atmospheric carbon increase ($\Delta T/\Delta C_A$: thin red line, right axis), airborne fraction of cumulative carbon emissions ($\Delta C_A/E_t$: thin blue line, left axis) and CCR (thick red line, right axis). In this simulation, cumulative airborne fraction decreased with time owing to a delayed carbon cycle response to a rapid prescribed rate of atmospheric CO_2 increase. This is consistent with saturating carbon sinks at higher atmospheric CO_2 , which leads to an increased airborne fraction of annual emissions with increasing atmospheric CO_2 . b: Simulations with an instantaneous doubling (solid lines) and quadrupling (dashed lines) of atmospheric CO_2 for 1,000 years (colours as in a). In all cases, the cumulative airborne fraction decreased with time, whereas the temperature change per unit atmospheric carbon increased with time; consequently, the CCR (defined as the product of these two quantities) remained constant in time. Source: Figure 2 from Matthews et al. (2009).

The ratio of CO_2 -induced warming realized over an interval of time, say a year, to cumulative carbon emissions over that same time interval has come to be known as the Transient Climate Response (TCRE) to CO_2 Emissions.⁷ This simplified linear characterization continues to provide a simplification by targeting the composite response of the carbon and temperature dynamics instead of the components that induce it. The time scale over which this approximation applies is important to economists aiming to incorporate climate consequences into their analyses. The time scale depicted in the lower panel of Figure 4 is very long relative to the time horizons that economists building dynamic economic models would find credible. The top panel seems more germane.

MacDougall and Friedlingstein (2015) use analytical reasoning to investigate why there is approximate constancy of the TCRE over a range of cumulative emissions up to 2000 Pg of carbon. They say,

The analysis reveals that TCRE emerges from the diminishing radiative forcing from CO_2 per unit mass being compensated for by the diminishing ability of the ocean to take up heat and carbon. The relationship is maintained as long as the ocean uptake of carbon, which is simulated to be a function of the CO_2 emissions rate, dominates changes in the airborne fraction of carbon. Strong terrestrial carbon cycle feedbacks have a dependence on the rate of carbon emission and, when present, lead to TCRE becoming rate dependent. Despite these feedbacks, TCRE remains roughly constant over the range of the representative concentration pathways and therefore maintains its primary utility as a metric of climate change.

They define TCRE as the change in global temperature divided by cumulated emissions of carbon (MacDougall and Friedlingstein (2015), equation (2)). Their analytical argument for the approximate constancy of TCRE over a range of emissions relevant for long term effects of cumulated emissions on temperature change is quite convincing.

While an approximate linear relationship of temperature to cumulative emissions appears to exist, uncertainty of the magnitude of this response remains. Figure 4 depicts a single CCR, but by looking across models shifts the focus to the uncertainty on this CCR response. See Matthews et al. (2012) and MacDougall et al. (2016) for further discussion. In particular, Figure 5 provides a characterization of uncertainty about the CCR for eleven prominent climate models. Each model yields an approximately linear relationship between

⁷See MacDougall et al. (2016) for a pedagogical summary of this literature.

cumulative emissions and temperature increase but the slope, i.e., the CCR's, varies across the eleven models, as does the nature of the linear approximation.




Figure 5: (a) Histogram of TCRE from the perturbed physics ensemble experiment. Mean value is 1.72 K EgC^{-1} . (b) Cumulative emissions versus temperature curves for all 150 model variants. Individual model variants are in grey, solid black line is the mean, and dashed lines are the 5th and 95th percentiles. Source: Figure 3 from MacDougall et al. (2016).

The spread of CCR values across the eleven models in Matthews et al. (2009) is apparently large as confirmed in later studies such as MacDougall et al. (2016) and displayed in Figure 5. Leduc et al. (2015) quantifies the limits of the linear approximation in Matthews et al. (2009) :

We conclude that the TCRE provides a good estimate of the temperature response to CO_2 emissions in RCP scenarios 2.6, 4.5 and 6, whereas a constant TCRE value significantly overestimates the temperature response to CO_2 emissions in RCP 8.5.

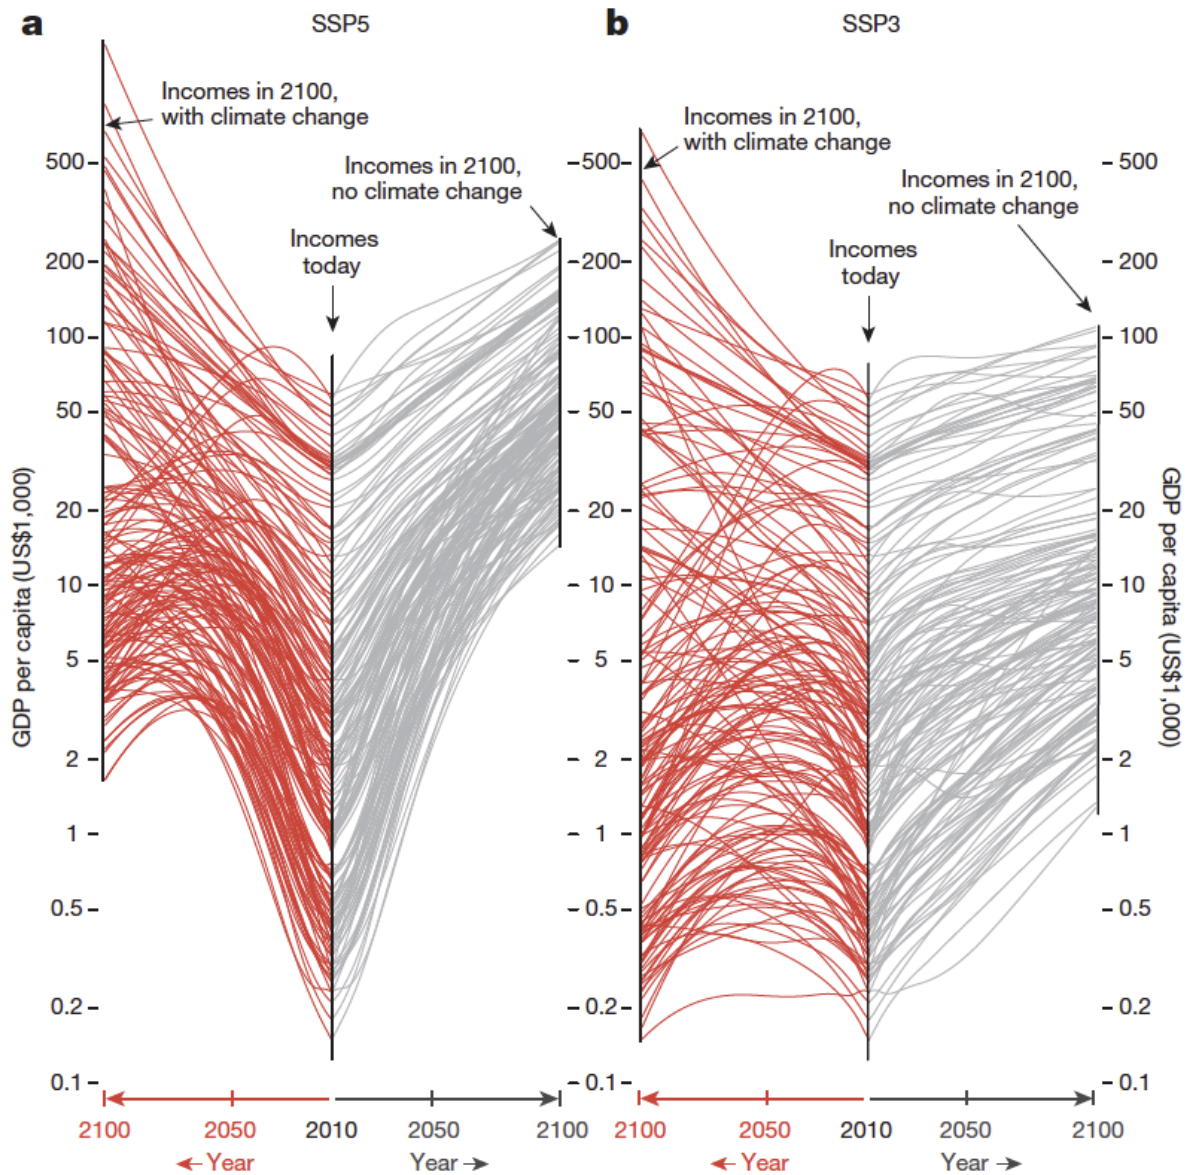


Figure 6: Country-level income projections with and without temperature effects of climate change. a, b, Projections to 2100 for two socioeconomic scenarios consistent with RCP8.5 ‘business as usual’ climate change: a, SSP5 assumes high baseline growth and fast income convergence; b, SSP3 assumes low baseline growth and slow convergence. Centre in each panel is 2010, each line is a projection of national income. Right (grey) are incomes under baseline SSP assumptions, left (red) are incomes accounting for non-linear effects of projected warming. Source: Figure 3 from Burke et al. (2015).

While we have to feature uncertainty in the mechanism that underlies the climate dynamics, how climate changes influence economic opportunity is also critical to any discussion of climate economics and is important to our long term uncertainty perspective here. Burke et al. (2015) study climate economic uncertainties facing rich and poor countries at latitudes ranging from tropical to temperate. Their evidence for country level income projection in Figure 6, with and without projected temperature changes out to the year 2100, around two socioeconomic scenarios consistent with RCP8.5. They argue that humans and crops thrive at a rather narrow range of temperatures with nonlinear declines for temperatures higher than the upper bound of this range. They say:

We show that overall economic productivity is nonlinear in temperature for all countries, with productivity peaking at an annual average temperature of 13 6C and declining strongly at higher temperatures. The relationship is globally generalizable, unchanged since 1960, and apparent for agricultural and non-agricultural activity in both rich and poor countries.

The projected nonlinear decline in economic productivity is especially severe for poorer countries located nearer the equator. There is a lot of uncertainty surrounding such long term projections including the possibility of misspecification. For instance, if adaptation proceeds in the poorer countries nearer the equator along the lines that Barreca et al. (2016) document for the warmer parts of the U.S., the strong decline at higher temperatures projected by Burke et al. (2015) may not occur. Extrapolating the adaptation evidence from Barreca et al. (2016) from U.S evidence is tricky, however. Since locales near the equator are the warmest already, it is not evident what can be copied in adapting agriculture to even warmer conditions. Hence, such strong declines under a high emissions scenario like RCP8.5 may still occur in the case of agriculture, especially for countries nearest to the equator. This rather extreme uncertainty about potential damages in the long term presents a challenge to decision theory, to which we now turn.

3 Decision theory and model slanting

Our discussion so far has featured the interplay between models, approximation, and uncertainty. Climate models are appealing because, in the language of econometrics, they are structural. They build in basic principles from physics that aim to add credibility in studying hypothetical responses to climate change induced by human inputs in the future. As we have noted, the models, while revealing are complex and may be wrong. Moreover, there are potentially important differences across models. The complexity leads naturally to a search for simplification, and this simplification carries over to characterizations of uncertainty and discussions of how we might weight the differing implications across models. To address model approximation and model ambiguity, we are led to decision theory as a framework for prudent policy discourse.

Decision theory, as we use the term, provides a formal framework for confronting uncertainty. This theory helps us understand the potential consequences of various components of uncertainty. Wald (1950)’s initiation of decision theory places the aim of analysis on making decisions that are defensible according to posited objective functions that trade off alternative aims. Advances in decision theory under uncertainty have been substantial and provide us with further guidance in addressing this important policy challenge in the face of incomplete knowledge.

To pose a decision problem requires an objective. Scientific discourse often chooses to avoid explicitly stated objectives, or at least try to, because they introduce preferences, tradeoffs or value judgements. Applications of decision theory require “subjective inputs” that some find objectionable. Assessing formally the consequences of uncertainty, however, makes it impossible to sidestep such issues. In this essay, we will not announce a single objective function but will instead focus on the mapping from objective functions to implications for uncertainty.

3.1 Who is making the decisions?

In economic analyses, decision theory comes into play in multiple ways. Economic models contain “agents”: consumers, firms, or governments, that make decisions. The interactions among these entities determine equilibrium prices and the allocation of resources. As part of the construction of an economic model, decision theory provides a way to approximate decision making in a dynamic environment. To study counterfactuals in a dynamic setting requires that we predict how agents respond to the corresponding change in the environ-

ment. While this can be a challenging task, one particularly powerful approach is to impose rational expectations motivated by a data richness that allows these agents only to confront one of the uncertainty components: risk. With this simplification, the probabilities used in forming expectations are determined as part of the equilibrium in the dynamic model by assuming that agents use the same probabilities as those implied by the model. The dynamic rational expectations approach purposefully rules out policies predicated on systematically fooling people. See, for instance, Lucas and Prescott (1971), Lucas (1976) and Sargent and Wallace (1975). Supposing that agents confront other forms of uncertainty pushes us to depart from the commonly employed rational expectations approach. Contributions from decision theory suggest ways in which we can imagine agents confronting ambiguity and concerns about model misspecification. This is what Hansen (2014) refers to as “inside” (the model) uncertainty in contrast to the “outside” perspective that we consider next.

External analysts and applied researchers, including those that support policy making, solve and assess models. They may confront unknown parameters or multiple models. Using evidence to infer parameters and assess models also benefits from a decision theoretic perspective. This “outside the model” perspective is the typical motivation from statistical decision theory and guided Wald (1950) and Savage (1954) in their initial formalization. From this vantage point, it is the analysts that face the uncertainty while taking the behavioral responses of the people inside the models as given.

3.2 Basic decision theory setup

In this section we draw heavily on presentation in Hansen and Marinacci (2017) in our description of alternative decision theories. Let the unknown models be indexed by a parameter θ that resides in a set Θ . Given θ , a random vector X with realizations $x \in \mathcal{X}$ is described by a probability density $f(x|\theta)$ relative to a measure τ over \mathcal{X} . For instance, θ could be an indicator of the alternative models used to generate Figure 1 including perhaps explicit economic interactions. Model *ambiguity* comes into play as a decision maker considers how much weight to attach to the alternative models when making decisions. Model *misspecification* is germane when all of the parameterized models could be constructed as simplified approximations but could be potentially flawed along some dimensions. A decision maker observes a realization x and takes an action $a \in \mathcal{A}$ that can depend on x . Formally, an *action* (or *decision*) *rule* is a suitably measurable function $A : \mathcal{X} \rightarrow \mathcal{A}$. We let $\mathcal{A}(x) \subset \mathcal{A}$ denote a set of potential actions constrained by state x so that $A(x) \in \mathcal{A}(x)$.

Represent the decision maker's preferences in terms of a utility function $U(a, x, \theta)$. Integrate over x to construct expected utility conditioned on θ :

$$\bar{U}(A|\theta) = \int_{\mathcal{Y}} U[A(x), x, \theta] f(x|\theta) \tau(dx), \quad (1)$$

The expected utility, as we have computed it, conditions on the parameter θ that is often unknown to the decision maker. Since A can depend on x , in what follows we will also condition on $X = x$ and work directly with U and not its *ex ante* counterpart.

Notice that we have entered θ as an argument in U . When applying decision theory to economics, the unknown parameter or model indicator θ may be an intermediate target, however. For instance, consider a decision maker facing uncertainty captured by a future payoff relevant state. Represent this state as a random vector Y with realized values y in a set \mathcal{Y} . Let $f^*(y|a, x, \theta)$ denote the density relative to a measure τ^* over alternative y 's in \mathcal{Y} conditioned on the current period action a and observed data x . Consider a next period utility function U^* that depends on (y, a) . For instance,

$$U^*(y, a) = v(a, x) + \beta V(y), \quad (2)$$

where $0 < \beta < 1$ is a subjective discount factor and V is tomorrow's value function. Integrate over y to construct:

$$U(a, x, \theta) = v(a, x) + \beta \int_{\mathcal{Y}} V(y) f^*(y|a, x, \theta) \tau^*(dy).$$

Thus the θ dependence of U could be induced by the dependence of f^* on θ . The integration over y while conditioning on θ using the density f^* adjusts for *risk*.

3.3 Bayesian decision theory

As posed so far, this representation of decision theory is incomplete. It confronts risk but not *ambiguity* across alternative models. Following de Finetti (1937) and Savage (1954), we include a subjective prior probability π , and integrate over the posited θ . Since a decision or action a can depend on the realized state x , we complete the specification of a conditional objective as:

$$\int_{\mathcal{Y}} \int_{\Theta} U^*(y, a) f^*(y|a, x, \theta) \tau^*(dy) \pi^*(d\theta|x)$$

where π^* is the familiar posterior formed by updating using Bayes rule:

$$\pi^*(d\theta|x) \propto f(x|\theta)\pi(d\theta).$$

Before we describe formally some alternative approaches, consider the conditional objective for this decision problem as a two-stage lottery. In stage two, a random outcome y is drawn from a given distribution associated with model θ in accordance to the model specific probabilities. In stage one, there is ambiguity as to which is among a family of distributions that will be used in stage two. The Bayesian approach to decision theory “reduces” the two-stage lottery into a single lottery with subjective probabilities providing the inputs for the decision maker and updated via Bayes rule by forming a posterior. The reduced lottery has probabilities represented by:

$$\int_{\Pi} f^*(y|a, x, \theta) \pi^*(d\theta|x) \tau^*(dy)$$

In this specification there is no scope for the expression of an aversion to model ambiguity that is distinct from risk. Both de Finetti and Savage acknowledge the challenge in using subjective probability to address such aversions as noted by Berger (1984) and Watson and Holmes (2016). The decision theories that follow do not reduce the two-step lottery and instead draw distinctions between the alternative components of uncertainty.⁸

While a robust Bayesian statistician characterizes the sensitivity of posterior probabilities by the choice of prior, a decision maker must confront this sensitivity when designing a course of action. This leads us to pose formally decision problems that illustrate alternative approaches that recognize this sensitivity.

3.4 Smooth ambiguity aversion

One tractable approach introduces aversion to prior or posterior ambiguity in a way that is conceptually similar to risk aversion by including a strictly increasing concave function Φ as in the smooth ambiguity model of Klibanoff et al. (2005):

Problem 3.1.

$$\max_{a \in \mathcal{A}(x)} \Phi^{-1} \left(\int_{\Theta} \Phi [U(a, x, \theta)] \pi^*(d\theta|x) \right).$$

⁸See Segal (1990) for a axiomatic rationale for not reducing two-stage lotteries.

Since we are interested in dynamic applications, we use $\pi^*(d\theta|x)$ which is the posterior for θ conditioned on x .⁹ In a dynamic context, yesterday's posterior is today's prior so that π^* may just as well be conceived as the today's prior. The curvature of the function Φ measures how a decision maker responds to ambiguity over models as distinct from risk conditioned on a model. When Φ is not affine, this decision theory ceases to be one in which the two-stage lottery described earlier is simply reduced to a single composite lottery by model averaging. This decision theory, however, continues to feature a single posterior distribution π^* without an explicit scope for assessing the sensitivity to the choice of prior. Nevertheless, the informativeness or lack thereof in the prior does play a role in the decision criterion through curvature in the function Φ .

What follows gives an alternative formulation of concerns about ambiguity across alternative models and is also sufficiently general to include concerns about model misspecification.

3.5 Max-min utility and penalization

An alternative approach addresses model ambiguity through the use of multiple priors or model misspecification by entertaining an extensive set of potential models. These methods impose aversion through finding the prior or model with the most adverse expected utility consequences subject to constraints or penalization. It provides a structured way to perform a sensitivity analysis. It follows Wald (1950)'s approach by relying on the game theoretic analysis of Von Neumann and Morgenstern (1944) to shape an approach to uncertainty.

Introduce a convex cost function C to penalize the exploration of alternative prior/posterior distributions π . This cost function captures how the decision maker confronts ambiguity. Formally, the decision maker solves:

Problem 3.2.

$$\max_{a \in \mathcal{A}(x)} \min_{\pi^*} \int_{\Theta} U(a, x, \theta) \pi^*(d\theta|x) + C(\pi^*|a, x).$$

Penalization methods are well known in both statistics and control theory. The preferences implicit in this decision problem are what Maccheroni et al. (2006a,b) call *variational preferences*. Such preferences nest the multiple priors specification of Gilboa and Schmeidler (1989), where the cost function takes on the extreme form of being equal to infinity if the priors are outside a convex set of priors Π and zero inside.

⁹See Klibanoff et al. (2009) for a dynamic extension of the smooth ambiguity model.

In what follows we sketch three different decision problems that are special cases of Problem 3.2. Two of the special cases have the decision maker be averse to ambiguity and, in the third case, averse to potential model misspecification. The first two are motivated in part by the aims of robust Bayesian methods. Robust Bayesians explore how sensitive inferential conclusions are to changing the prior when there is doubt as to which prior to use in the analysis. Problem (3.2), with its embedded minimization, goes one step further by imposing an aversion to ambiguity over the specification of the prior/posterior. The third one builds a weakly structured set of alternative models vis a vis a benchmark or reference model as way a to allow this initial model to be misspecified.

3.5.1 Smooth ambiguity reconsidered

While the smooth ambiguity Problem 3.1 does not formally entertain prior sensitivity, Hansen and Sargent (2007) point out that sometimes there is an alternative interpretation. For the familiar and commonly used relative entropy formulation of costs, there is a simple connection between a penalization approach for assessing sensitivity and the smooth ambiguity approach.

To illustrate this, let π_o^* denote the corresponding reference posterior. Also let $g \geq 0$ denote a probability density with respect to π_o^* , implying that

$$\int g(\theta)\pi_o^*(d\theta|x) = 1.$$

Finally, let \mathcal{G} denote the family of such densities. Relative entropy is measured by:

$$\int_{\Theta} \log g(\theta)g(\theta)\pi_o^*(d\theta|x)$$

and is nonnegative and convex in g . It is zero in the when $g = 1$ the reference posterior in constructing the conditional objective for the decision problem. Express ambiguity aversion by solving:¹⁰

Problem 3.3.

$$\max_{a \in \mathcal{A}(x)} \min_{g \in \mathcal{G}} \int_{\Theta} U(a, x, \theta)|\theta)g(\theta)\pi_o^*(d\theta) + \kappa \int_{\Theta} \log g(\theta)g(\theta)\pi_o^*(d\theta|x)$$

¹⁰This minimization problem is a special case of an optimization problem with a relative entropy penalty that emerges in a variety of areas of applied mathematics. We will see another application in what follows.

$$= \max_{a \in \mathcal{A}(x)} -\kappa \log \int_{\Theta} \exp \left[-\frac{1}{\kappa} U(a, x, \theta) \right] \pi_o^*(d\theta|x). \quad (3)$$

Thus, a particular form of a smooth ambiguity model emerges from a search over alternative posterior densities subject to a penalization. The minimizing g satisfies:

$$g^*(\theta) = \frac{\exp \left[-\frac{1}{\kappa} U(a, x, \theta) \right]}{\int_{\Theta} \exp \left[-\frac{1}{\kappa} U(a, x, \theta^*) \right] \pi_o^*(d\theta^*|x)}$$

which generates probabilities that are tilted towards θ 's with adverse consequences for the expected utility: $U(a, x, \theta)$. Recall that in our dynamic application, $U(a, x, \theta)$ is itself constructed by integrating over the risk conditioned on a model as in (2).

This provides a concrete illustration of the impact of prior/posterior uncertainty. By targeting a posterior conditioned on x the sensitivity analysis is implicitly over both the likelihood used to represent probabilities over alternative x realizations given θ and some initial prior over θ .

3.5.2 Robust Bayesian method with constraints

Consider a different approach that imposes constraints on π^* , targeting perhaps ambiguity about date zero priors. We formulate the problem in order that it be dynamically consistent in the manner justified by Epstein and Schneider (2003). We start with a parameterized family of such models f^* parameterized by θ and a convex family of potential posterior probabilities $\Pi(x)$ with the state x sufficiently rich to encode relevant past information needed for updating probabilities for any of the initial family of priors. The resulting decision problem is

Problem 3.4.

$$\max_{a \in \mathcal{A}} \left[v(a) + \beta \min_{\pi^* \in \Pi(x)} \int_{\mathcal{Y}} V(y) f^*(y|a, x, \theta) \tau^*(dy) \pi^*(d\theta|x) \right]. \quad (4)$$

3.5.3 Misspecification

We confront the potential misspecification of a reference or benchmark probability model by introducing a broad set of alternative models and considering all possible priors over models in this set. As an example, consider one of the preferences for robustness used in control theory and adapted to economics by Hansen and Sargent (2001). Under these preferences the cost function takes the form of a relative entropy penalty for deviating from

the reference probability model. Formally, suppose that θ is a relative density in a space of relative densities Θ satisfying

$$\int_{\mathcal{Y}} \theta(y|a, x) \tau^*(dy) = 1$$

and f_o^* is a baseline or reference transition density and model θ has transition density $\theta(y|a, x) \tau^*(dy)$. Construct a cost function in terms of the log-likelihood ratio of a θ model relative to the initial reference model:

$$C(\theta|a, x) = \kappa \int_{\mathcal{Y}} [\log \theta(y|a, x) - \log f_o^*(y|a, x)] \theta(y|a, x) \tau^*(dy)$$

Then for specification (2), minimizing over θ gives:

$$\theta^*(y|a, x) = \frac{\exp \left[-\frac{1}{\kappa} V(y) \right] f_o^*(y, |a, x)}{\int_{\mathcal{Y}} \exp \left[-\frac{1}{\kappa} V(y^*) \right] f_o^*(y^*, |a, x) \tau^*(dy^*)}$$

This outcome is analogous to the one in section 3.5.1 but applied to continuation values expressed as a function of y . This is the so-called *exponential tilting* solution in which the adjustment for potential model misspecification slants probabilities towards future states for which the value function is relatively low. This illustrates how adjustments for model misspecification depend on the specifics of the decision problem and the consequences of alternative courses of action as reflected by the value function and the baseline transition density.

The outcome of this minimization is the reduced form:

$$\begin{aligned} & \min_{\theta \in \Theta} \left[v(a) + \beta \int_{\mathcal{Y}} \theta(y) V(y) f_o^*(y|a, x) + C(\theta|a, x) \right] \\ &= v(a) - \beta \kappa \log \int_{\mathcal{Y}} \exp \left[-\frac{1}{\kappa} V(y) \right] f_o^*(y|a, x) \tau^*(dy). \end{aligned} \tag{5}$$

which results in the decision problem:

Problem 3.5.

$$\max_{a \in \mathcal{A}} \left[v(a) - \beta \kappa \log \int_{\mathcal{Y}} \exp \left[-\frac{1}{\kappa} V(y) \right] f_o^*(y|a, x) \tau^*(dy) \right].$$

The second line of (5) has the appearance of a risk adjustment familiar from recursive

utility theory even though it is constructed from a concern that the baseline model is misspecified. Initial demonstrations of the reduced form relationship appeared in the control theory literature by Jacobson (1973) and Whittle (1981).¹¹ Setting $\kappa = \infty$ results in an infinite penalty and with the outcome being:

$$v(a) + \beta \int_{\mathcal{Y}} V(y) f_0^*(y|a, x) \tau^*(dy)$$

3.5.4 Hybrid approaches

While the approach illustrated in Section 3.5.3 to model misspecification allows for a rich family of alternative θ 's, it features a single baseline transition density f_o^* . Suppose that instead we start with a parameterized family of transition densities $f^*(\cdot|a, x, \alpha)$ where we now use the notation α to denote hypothetical parameter values. We continue to let θ denote alternative transition densities.

One possibility is to apply a smooth ambiguity adjustment to:

$$U(a, x, \alpha) = v(a) - \beta \kappa \log \int_{\mathcal{Y}} \exp \left[-\frac{1}{\kappa} V(y) \right] f^*(y|a, x, \alpha) \tau^*(dy)$$

where the right-hand side is the outcome of θ minimization given α as in (5) where the minimization is performed for each α . In effect we have a parameterized family of reference models, each of which could be misspecified. Hansen and Sargent (2007) suggest applying the smooth ambiguity adjustment given by equation (3) motivated by a robust posterior analysis targeting ambiguity about the parameter α .

Another possibility is to start with a robust Bayesian problem with a convex set of posteriors over α in a set $\Pi(x)$ and entertain misspecification of any of these models by again introducing densities $\theta \geq 0$. We form a cost function for θ by

$$C(\theta|a, x) = \min_{\pi^* \in \Pi(x)} \left(\log \theta(y|a, x, \alpha) - \log \left[\int f_{\alpha}^*(y|a, x, \alpha) \pi^*(d\alpha|x) \right] \right) \theta(y) \tau^*(dy)$$

Hansen and Sargent (2016) develop and apply a continuous-time counterpart to this approach building on previous work of Hansen and Sargent (2001) and Chen and Epstein

¹¹They studied risk and robustness using a relative entropy formulation for a linear-quadratic environment. See, for example, Petersen et al. (2000) for a formulation of robustness using relative entropy when applied to a more general control theory environment. While the connection between risk aversion and a concern about model misspecification is self evident when the intertemporal elasticity of substitution is unity, Maenhout (2004) connects these constructs for other elasticity of substitution parameters.

(2002).

These alternative decision problems give examples of responses to uncertainty when model-builders and policy makers push beyond the usual risk model under which probabilities are presumed to be known. As we see from these problems, the quantitative consequences of treating uncertainty more broadly depend on what aspects of uncertainty we choose to feature and how much aversion we impose to that uncertainty. An important task for quantitative research in climate economics is to explore what external evidence will be most revealing in helping to implement these decision problems as part of the construction and application of climate-economic models. Specifically they will help to direct social policy towards climate policy that addresses the potential ruinous outcomes featured by Bettis et al. (2017) and others and avoid featuring social costs of carbon based on overstated knowledge.

3.6 Probability slanting

The max – min approach to decision theory can be viewed equivalently as two-player zero sum game. There are well known circumstances from the theory of zero-sum games that inform us when we can reverse the orders of max and min without altering the implied value function. In such circumstances, there is typically a well defined worst case probability model for the transition density. When the maximizing player optimizes using this model, the resulting decision rule coincides with that from the max – min problem. This entails a conservative adjustment.

Figure 7 illustrates this phenomenon in a model with macroeconomic growth-rate uncertainty analyzed by Hansen and Sargent (2016). Investors face explicit uncertainty about the persistence of macroeconomic growth as well as concerns about model misspecification. In addition to the overall caution, these concerns are reflected by investor worst-case models that show more persistence when growth is sluggish and less persistence when it is vigorous. This figure shows how these impacts compound over time.




Figure 7: Distribution of the logarithm of aggregate consumption growth $Y_t - Y_0$ under the baseline model and worst-case model for a macroeconomic model with growth rate uncertainty. The gray shaded area depicts the interval between the .1 and .9 deciles for every choice of the horizon under the baseline model. The red shaded area gives the region within the .1 and .9 deciles under the worst-case model. Source: Hansen and Sargent (2016)

This framework provides a formal way to operationalize caution in how we use models. This caution is reflected in the endogenously determined worst-case models and they depend on the objective used in the decision problem. They are the outcome of analysis whereby we acknowledge uncertainty more broadly conceived than risk.

This type of analysis suggests the computation of worst case models as way to *char-*

acterize or to *implement* cautious decision making. This type of slanting may also emerge implicitly in attempts to *influence* policy making. While we understand a role for slanting as an outcome of aversion to uncertainty when conceived broadly, the use of slanting should be recognized as tied to a specific decision problem and not the model/evidence based inputs into a decision problem.

4 Long-term risk

Macro asset pricing has featured characterizations of “long-term risk” ranging from uncertain growth to rare disasters. The most common approach presumes rational expectations and endows agents inside the model with preferences for which intertemporal composition of risk matters. By extending the recursive utility framework of Koopmans (1960), Kreps and Porteus (1978) produce tractable preferences that allow for the intertemporal composition of risk to matter. In the macro asset pricing literature, the resulting recursive utility preferences are typically used in conjunction with rational expectations. Epstein and Zin (1989) show that this framework provides a tractable and revealing way to distinguish risk aversion from intertemporal substitution. Using this preference specification for investors, Bansal and Yaron (2004) show how risk about long-term macroeconomic growth can have a quantitatively important impact on even short-term asset pricing.

Measurements of the social cost of carbon often make reference to the discount rates and some discussions of measurement explore the sensitivity to the choice of those rates. From an asset pricing perspective, a reference to even a one-period rate used in the market discounting of cash flows requires modification in the presence of uncertainty. Payoffs are risky and their present values depend on the inherent riskiness. A market-based measure of riskiness depends on how the cash flow covaries with the uncertain macroeconomic outcomes. There is not a single number to use as a discount factor, but instead a stochastic discount factor that is different depending on the future macroeconomic outcomes.¹² Present-discounted-value calculations under uncertainty require cumulating stochastic discount factors over multiple horizons. Thus not only are interest rates compounded, but also the so-called risk prices. These risk prices are the market compensations for exposure to macroeconomic risk. Incorporating an aversion to ambiguity or model misspecification alters the resulting equilibrium stochastic discount factor processes in known ways and

¹²See Ross (1976) for the pricing of nondiversifiable macroeconomic risk and Hansen and Richard (1987) for the stochastic discount factor formulation of the empirical implications in dynamic settings.

changes the market-based valuations.

Why might we be interested in present-value discounting in assessing climate impacts instead of a more explicit global analysis? The best defenses are those that are local in nature. They assess the impact of small changes in policy provided that we start with a stand-in consumer or a Pareto problem that weights the different consumer types. Within such a framework, small policy changes whose impacts play out over time can sometimes be depicted with present-value formulas just like calculations of market-based measures.¹³ Sometimes such local calculations support or provide bounds on more global calculations. Examples of such calculations are the so-called Pigouvian tax rates computed in terms of marginal contributions evaluated at the socially efficient intertemporal allocation of resources. These tax rates are idealized marginal net costs of social externalities. Alternatively, local calculations might bound global responses.

In the next subsection, we illustrate the asset pricing over multiple horizons in the context of simple illustration prior to exploring economic environments with more complexity.

4.1 A revealing example

For pedagogical simplicity, we follow a formulation in Hansen et al. (2008).¹⁴ Consider a consumption process for which the logarithm of consumption at date t is denoted Y_t . We presume that the growth rate in consumption, $Y_{t+1} - Y_t$ has an impulse response of the form: $\{\alpha_\tau \cdot W_{t+1} : \tau = 0, 1, \dots\}$ where W_{t+1} is a vector of multivariate normally distributed shocks that have a linear impact $\alpha_\tau \cdot W_{t+1}$ on $Y_{t+\tau+1} - Y_{t+\tau}$. We refer to $\{\alpha_\tau : \tau = 0, 1, \dots\}$ as the vector of impulse responses. Notice that the impact W_{t+1} on the cumulative stochastic growth of $Y_{t+\tau+1} - Y_t$ is: $\sum_{j=0}^{\tau} \alpha_j$. When

$$\lim_{\tau \rightarrow \infty} \sum_{j=0}^{\tau} \alpha_j = \alpha_\infty \neq 0,$$

there is exposure to so called long-run risk as a shock today that has permanent consequences for consumption arbitrarily far into the future.

As a practical matter, it is challenging to measure this limiting exposure accurately (see Hansen et al. (2008)). But let us put the actual measurement challenge to the side and

¹³See Hansen et al. (1999) and Alvarez and Jermann (2004) for some illustrations of how to use asset pricing formulas in stochastic environments to make welfare statements.

¹⁴This formulation has very similar simplicity to that of Campbell and Vuolteenaho (2004) although the modeling inputs differ.

suppose that investors have confidence in their ability to quantify accurately the impulse responses. To assess the long-term risk, suppose a decision maker has continuation utilities that satisfy the recursion:

$$V_t = (1 - \beta)Y_t + \frac{\beta}{1 - \gamma} \log E(\exp[(1 - \gamma)V_{t+1}] | \mathcal{F}_t) \quad (6)$$

where V_t is the date t continuation value and β is a subjective discount factor. Notice that this recursion includes a risk-adjustment of the next periods continuation value with a risk aversion parameter of γ . In the absence of risk, that is when the next periods continuation value V_{t+1} is perfectly forecastable, the parameter γ drops out of the recursion and $V_t = (1 - \beta)Y_t + \beta V_{t+1}$. Solving this forward gives the discounted utility formula:

$$V_t = (1 - \beta) \sum_{j=0}^{\infty} \beta^j Y_{t+j}.$$

When V_{t+1} is not known one-period in advance, the parameter γ captures an exponential risk adjustment in the future continuation value. This type of risk adjustment is implicit in risk sensitive control initiated by Jacobson (1973) and Whittle (1981).¹⁵

We rewrite recursion (6) to feature growth-rate uncertainty:

$$V_t - Y_t = \frac{\beta}{1 - \gamma} \log E(\exp[(1 - \gamma)(V_{t+1} - Y_{t+1} + Y_{t+1} - Y_t)] | \mathcal{F}_t).$$

Let $\{v_\tau : \tau = 0, 1, \dots\}$ be the impulse responses for $\{V_t - Y_t\}$. The moving-average coefficients solve a corresponding recursion:

$$v_j = \beta(v_{j+1} + \alpha_{j+1})$$

for $j = 0, 1, \dots$. Solving this equation forward:

$$v_\tau = \sum_{j=1}^{\infty} \beta^j \alpha_{j+\tau}.$$

The exposure of V_{t+1} to a date $t + 1$ shock is the combination of exposures of $V_{t+1} - Y_{t+1}$

¹⁵Hansen and Sargent (1995) showed how to formulate risk-sensitive control in a way that makes it a special case of recursive utility.

and $Y_{t+1} - Y_t$ to W_{t+1} and given by:

$$\left(\sum_{j=1}^{\infty} \beta^j \alpha_j + \alpha_0 \right) \cdot W_{t+1} = \alpha(\beta) \cdot W_{t+1},$$

where

$$\alpha(\beta) = \sum_{j=0}^{\infty} \alpha_j \beta^j$$

is the vector of discounted impulse responses for the consumption growth rate.

The multi-period stochastic discount factor for horizon τ , $\frac{S_{t+\tau}}{S_t}$, equals the implied τ period intertemporal marginal rate of substitution where

$$\frac{S_{t+1}}{S_t} = \beta \exp(Y_t - Y_{t+1}) \left[\frac{\exp[(1 - \gamma)V_{t+1}]}{E(\exp[(1 - \gamma)V_{t+1}] | \mathcal{F}_t)} \right] \quad (7)$$

The first term on the right-hand side is the contribution from the subjective rate of discount, the second term is the one-period intertemporal marginal rate of substitution for time separable logarithmic utility for which the elasticity of intertemporal substitution is unity. The third term is the contribution from the more general utility recursion that emerges when $\gamma > 1$. We impose the unitary elasticity assumption for sake of illustration and instead feature the impact of γ for uncertainty prices. Notice that the third term in (7) has conditional expectation one, which we will have more to say about subsequently.

4.2 One-period asset pricing

For assessing investment opportunities over one time period, it is convenient to rewrite this formula:

$$\frac{S_{t+1}}{S_t} = E\left(\frac{S_{t+1}}{S_t} \mid \mathcal{F}_t\right) \exp\left([(1 - \gamma)\alpha(\beta) - \alpha_0] \cdot W_{t+1} - \frac{1}{2} |(1 - \gamma)\alpha(\beta) - \alpha_0|^2\right)$$

The term $E\left(\frac{S_{t+1}}{S_t} \mid \mathcal{F}_t\right)$ is the implied discount factor abstracting from risk exposure at date $t + 1$. It is the reciprocal of the gross one-period interest rate. The exposure of the stochastic discount factor to the shock vector W_{t+1} encodes the one-period risk adjustments. Specifically, this exposure is given by $\rho_0 \cdot W_{t+1}$ where:

$$\rho_0 = (\gamma - 1)\alpha(\beta) + \alpha_0$$

denotes the implied one-period risk price vector for exposure to the vector of shocks W_{t+1} . For instance, consider a one-period log-normal cash flow:

$$\exp(\eta + \nu \cdot W_{t+1}) \quad (8)$$

where ν is a vector of exposures to one-period risk. The logarithm of the proportional risk premia for this cash flow is

$$\begin{aligned} & \log E[\exp(\eta + \nu \cdot W_{t+1}) | \mathcal{F}_t] - \log E\left[\left(\frac{S_{t+1}}{S_t}\right) \exp(\eta + \nu \cdot W_{t+1}) | \mathcal{F}_t\right] \\ & \quad + \log E\left[\left(\frac{S_{t+1}}{S_t}\right) | \mathcal{F}_t\right] \\ & = \nu \cdot \rho_0. \end{aligned}$$

The first two terms on the left-hand side of the equation are equal to the logarithm of the expected return to the cash flow and the third term is the negative of the logarithm of the risk-free return. Thus the entries of ρ_0 are the prices used to represent the proportional compensations for exposures to the alternative components of the shock vector W_{t+1} . These alternative components could be constructed cash flows with different exposure vectors ν .

The long-run risk models in the macro finance literature use models for which some of the entries of $\alpha(\beta)$ are substantially larger than the corresponding α_0 . That is, the long-term growth rate response measured as a discounted sum is larger than the immediate response at least for some components of the shock W_{t+1} . This illustrates that long-term macroeconomic risk exposure can imply large short-term prices in contrast to discounted expected utility ($\gamma = 1$).

4.3 Horizon dependent prices

Valuing intertemporal assets entails discounting and adjusting for risks over multiple horizons. Here we follow Hansen (2012), Borovička et al. (2014), and Borovička and Hansen (2016), but specialized to lognormal models for pedagogical simplicity. The horizon dependent risk prices are the multi-period impulse responses for the cumulative stochastic discount factor process. They are computed by compounding the one-period stochastic

discount factors and are given by:

$$\rho_\tau = (\gamma - 1)\alpha(\beta) + \sum_{j=0}^{\tau} \alpha_j$$

for investment horizon τ . While the recursive utility function informs us about horizon-specific price effects, exposures of the cash flows to these shocks also typically depend on the horizon and contribute to asset valuation. Notice that the limiting price is

$$\lim_{\tau \rightarrow \infty} \rho_\tau = (\gamma - 1)\alpha(\beta) + \alpha(1) \approx \gamma\alpha(1)$$

where the approximation on the right-hand side holds if the subjective discount factor β is close to one.

Let's take this inventory so far. In economic models with random impulses, we may ask which sources of uncertainty are most consequential. The answer to this question typically depends on the horizon over which the shock has an impact. Provided that we can identify the macroeconomic shocks of interest, asset pricing methods provide a way to quantify implied market compensations for horizon specific exposures. We may deduce these prices by constructing empirically relevant structural models or extracting more directly from asset market data.

4.4 Stochastic technology and stochastic volatility

The log-linear environment allows us to produce simple and revealing characterization, but it is also constraining both from an interpretive standpoint and an empirical standpoint. The appendix constructs a formal model with a so-called AK technology and adjustment costs in investment. This model extends a single sector version of a model in Eberly and Wang (2009) by incorporating a predictable component to the physical returns to investment. In this model there are shocks that induce riskiness in the physical returns to investment and shocks that alter the predictable component of these returns. Finally, there are shocks to stochastic volatility. These shocks alter future consumption and investment. With the AK technology, the implied consumption dynamics match those of a so-called long-run risk model of the type featured by Bansal and Yaron (2004).

In the absence of stochastic volatility, the consumption dynamics are fully log-linear as in our previous analysis. Stochastic volatility changes the previous analysis in two

ways. First, the implied compensations for the exposure to macroeconomic uncertainty depend on the volatility state; and second, exposure to the stochastic volatility shock requires compensation. In what follows we illustrate the impact of this state dependent compensation. While the first two shocks both have permanent consequences, we represent them as linear combinations of a permanent shock and an uncorrelated transitory shock and illustrate their distinct impact on valuation. From the standpoint of valuation, this distinction is important for typical calibrations of the long-run risk model.

Figure 8 depicts the impulse responses for temporary and permanent technology shocks. Notice, in particular, that the impact of the permanent shock builds over time and only approximates its peak impact over a long-time horizon.




Figure 8: Impulse responses for the temporary and permanent and temporary shocks. Solid curves depict the median responses. The light blue bands depict .1 and .9 deciles. The responses for the temporary shocks are given in the left panel for the permanent shocks in the right panel.

Figure 9 gives the shock-price elasticities for the temporary and permanent shocks. It compares outcomes for the power utility and the recursive utility specifications of preferences. The power utility model uses the same value of γ as the recursive utility model, but it restricts $\rho = \gamma$. As the subjective discount rate approaches zero, the limiting prices as the investment horizon becomes arbitrarily large converge to the same limit point. For the power utility model, the risk prices are essentially proportional to the impulse responses for consumption. For the recursive utility model, the price compensations are much more substantial for the permanent shock. This occurs because of the forward-looking channel

in the permanent income model. The price trajectory is almost flat for the recursive utility model in contrast to the power utility model. While the shock prices eventually become large for the power utility model, they are large at the outset for the recursive utility model. The bands in Figure 9 show the impact of the initial state on these trajectories. To avoid clutter, we only depict this impact for the recursive utility model.




Figure 9: Shock price elasticities. The solid curve depicts the median risk price for recursive utility and the dashed curve depicts the median risk price for power utility. The bands depict a range between both the .1 and .9 deciles for the recursive utility risk prices. These plots presume that $\gamma = 8$. The intertemporal elasticity of substitution is unity for the recursive utility model calculations.

4.5 Incorporating temperature dynamics

So far we have abstracted from the human impact on climate change. Figure 4 reveals and the CCR approximation presumes that random shifts in carbon emissions will have long-term consequences on temperature. Indeed, Archer et al. (2009) indicate that

There is a strong consensus across models of global carbon cycling, as exemplified by the ones presented here, that the climate perturbations from fossil fuel CO_2 release extend hundreds of thousands of years into the future.

Measurements of climate damages typically depend on temperature where these damages might alter output or directly influence the utility functions of stand-in consumers. Thus

climate economics offers one source of long-term risk, or risk that can be magnified over time.

In the appendix, we introduce a simplified climate model imposing the Matthews et al.-style proportionality relation between human CO_2 emissions and temperature. While this is a longer-term approximation, for pedagogical simplicity we impose it in formulating the climate dynamics. In this simplified model, temperature τ_t evolves in continuous time as:

$$d\tau_t = \lambda \mathcal{E}_t dt + \sqrt{\mathcal{E}_t} \sigma_\tau \cdot dW_t + d\tau_t^*$$

where $\{\tau_t^*\}$ is a stationary process and $\sigma_\tau \cdot W_t$ is a scalar Brownian motion. Notice that when \mathcal{E}_t is zero, the temperature τ_t and τ_t^* agree subject to initialization. As we have made abundantly clear, the Matthews et al.-approximation disguises a much more complex climate-energy dynamical system with feedbacks. We impose a Hotelling-like stock constraint on the total amount of fossil fuels which in turn limits the overall consequences for global warming. We include an exponential damage function that captures adverse temperature impacts represented as proportional reductions in productivity: $\exp(-\gamma\tau_t)$ where τ_t is the current period temperature. There are many reasons to entertain other specifications for external economic damages inflicted by human activity. Much has already been written on the challenges of measuring damages. We make these modeling simplifications in large part for analytical convenience. We now let the current period utility contribution to the continuation value recursion be

$$Y_t = (1 - \alpha) \log C_t + \alpha \log \mathcal{E}_t$$

where \mathcal{E}_t is a measure of “dirty energy” that emits CO_2 into the atmosphere.

By design this model exploits the previously presented model of long-term productivity risk, but now we have included a temperature adjustment with adverse consequences for production. In light of this adjustment, temperature dynamics now come into play. The model continues to have a quasi-analytical solution, one that we derive and display in the appendix.

4.6 Other related examples

Dietz et al. (2017) use an asset pricing perspective to explore the social consequences in which the logarithm of a consumption process has independent and normally distributed

increments. Let $\mu_c + \sigma_c \cdot W_{t+1}$ denote the increment in the logarithm of consumption between dates t and $t+1$. As a consequence, the vector of shock price elasticities is simply: $\gamma\sigma_c$ independent of the states and the horizon. Given this simple structure, we define the composite scalar shock

$$\gamma W_{t+1}^* = \frac{1}{|\sigma_c|} \sigma_c \cdot W_{t+1}$$

which we have normalized to have unit standard deviation. It is the exposure to this composite shock that requires compensation. The price elasticity for the composite shock is $\gamma|\sigma_c|$.¹⁶

Dietz et al. (2017) consider a local net benefit process B_t where the localization allows them to employ asset pricing methods for assessing its social consequences. The process $\{\log B_t : t = 0, 1, \dots\}$ has a time varying response to the composite shock $\{W_{t+1}^* : t = 0, 1, \dots\}$ where the response of $\log B_{t+\tau}$ to W_{t+1}^* is given by

$$\beta_{t+\tau} |\sigma_c|$$

which depends on both the time period t and the horizon τ to which the net benefit accrues, albeit in a special way. While the shock price elasticities are constant, the exposure elasticities depend on the horizon τ and this dependence is reflected in the implied risk adjustments. Dietz et al. (2017) target their analysis towards quantifying the sequence $\{\beta_t : t = 0, 1, \dots\}$ through approximations from climate models. They explore counter-vailing influences on this measurement. Positive shocks to factor productivity growth are associated with increases in emissions and induce their exposure sequence $\{\beta_t : t = 0, 1, \dots\}$ to exceed unity whereas the damages from the climate component from the increases in emissions push this sequence towards numbers less than one. Dietz et al. (2017) use this as a framework for quantification by fitting simple approximations to some climate economic models. In their study of some of the existing models, they find that the combined effects net out to produce an exposure sequence $\{\beta_t : t = 0, 1, \dots\}$ close to one. Like our other illustrations, this is a pedagogically interesting example that warrants further scrutiny through the use of more ambitious models and measurements.

Bansal et al. (2016) also use a similar setup to the stylized model given here. They also have an explicit linkage to temperature. In their model, temperature alters both

¹⁶Dietz et al. (2017) assume a power utility model with $\gamma = \rho$. With the assumption of independent increments, when it comes to risk pricing there is a well known observational equivalence within the family of recursive utility preferences. See Kocherlakota (1990) for a discussion.

the intensity of an adverse change in the logarithm of a “climate good” as well as the distribution given that an adverse event occurs.¹⁷ While temperature is formally included in their analysis, there is no attempt to connect to the transition dynamics reported in Figure 4 and instead the focus is on so-called disaster events. But their model serves as a valuable illustration of the asset pricing impacts of long-run risk models with climate shocks.

4.7 Making decisions robust

The risk sensitive control theory of Jacobson (1973) and Whittle (1981) noted a connection between making control objective more sensitive to risk and including a robust adjustment for model misspecification. As we have noted, this connection carries over to some recursive formulations of preferences. It is evident by comparing the reduced form outcome from (5) to the recursive representation of continuation values in (6). The recursive utility adjustment in (6) could be the outcome of a concern about misspecified dynamics as in (5) provided that we set

$$\frac{1}{\kappa} = \gamma - 1$$

for $\gamma > 1$. The implied worst-case model entails a relative density adjustment that is the exponential tilting:

$$\frac{\exp[(1 - \gamma)V_{t+1}]}{E(\exp[(1 - \gamma)V_{t+1}] | \mathcal{F}_t)},$$

the same term that captures the recursive utility adjustment to one-period stochastic discount factor in (7). As we argued earlier, probabilities are tilted towards states in which continuation values are relatively large.

In the simple production based model, the formula for the adverse shift in the conditional mean for capital is given by:

$$\text{capital evolution distortion} = -\frac{.01}{\kappa} \sigma_k(z)' \left[(.01) \sigma_k(z) + \sigma_z(z) \frac{\partial \nu}{\partial z}(z) \right]$$

In this formula z is a potential realization of the exogenous state vector, Z_t , $\sigma_k(Z_t) \cdot dW_t$ is the shock to a capital investment and ν is the exogenous state contribution to the value

¹⁷Bansal et al. (2016) have a composite good that we could take to be Y_t with a common long-run risk dynamics. They use this same good as the numeraire in the construction of the stochastic discount factor in contrast to the formula we provide. The important difference is the explicit use of temperature to trigger an extreme event. They also alter the unitary elasticity of substitution assumption.

function, $\sigma_z(z)'dW_t$ gives the local impact of the shock vector dW_t to the exogenous state evolution. Note the potential dependence of σ_k and σ_z on the exogenous state stochastic volatility in the macro economy. They show up in the implied worst-case model because more volatility makes statistical discrimination as captured by relative entropy more challenging. When both of these objects are constant and the conditional mean for dZ_t is linear in Z as in the case of a first-order vector autoregression, ν is linear and hence its partial derivative is constant. The implied adverse shift in the conditional mean dynamics is constant as well.

Quasi analytical formulas are also available when the σ_k and σ_z depend on the square root of the one of the exogenous states as in a so-called Feller square root process. The value function ν remains linear in z as is the case in the production economy that underlies the computations for Figures 8 and 9. Now the implied worst case model depends linearly on the associated volatility state.

Next we introduce uncertain temperature dynamics. Misspecification concerns now emerge for both the capital evolution equation and the temperature equation:

$$\begin{aligned} \text{capital evolution distortion} &= -\frac{.01(1-\alpha)}{\kappa} \sigma_k(z)' \left[(.01)\sigma_k(z) - \gamma\sigma_{\bar{\tau}}(z) + \sigma_z(z)' \frac{\partial \nu}{\partial z}(z) \right] \\ \text{temperature evolution distortion} &= \frac{\gamma(1-\alpha)}{\kappa} (|\sigma_{\tau}|^2 e + |\sigma_{\bar{\tau}}(z)|^2). \end{aligned}$$

For simplicity, let's suppose that $\sigma_{\bar{\tau}}$ is constant. The robust adjustment to the temperature dynamics is affine, with an altered slope coefficient for \mathcal{E}_t . In particular, the implied worst-case model is one for which the λ in the Matthews approximation is enhanced to:

$$\lambda^* = \lambda + \frac{1}{\kappa} |\sigma_{\tau}|^2 (1-\alpha) \gamma.$$

This augmentation depends not only on the penalty parameter κ but also on the local variance $|\sigma_{\tau}|^2$ for the temperature dynamics along with the preference parameter α and the damage parameter γ . It is a conservative adjustment induced by an effort to make more cautious decisions in the face of potential model misspecification. The variance contribution $|\sigma_{\tau}|^2$ is present because of our use of a relative entropy and its connections to likelihood functions. Alternative specifications that are hard to distinguish statistically are entertained as alternative models. The damage parameter γ is present because of the long-term adverse consequences of temperature. This conservative adjustment for the temperature dynamics carries over directly to enhance the implied Pigouvian tax on emissions.

We include these calculations as an illustration and not as a full fledge analysis. Our aim is to illustrate the impact of the concern for robustness in a highly stylized example. A more ambitious approach is called for to produce a credible quantification, and it will necessarily be computational in nature. While pedagogically revealing, the example we sketched is too special. In a more ambitious model the finite energy stock constraint would be replaced by a backstop green technology. Concerns about model misspecification should carry over to the specification of damages. Moreover, the analysis of tipping points in climate features a particular form of nonlinearity that could have big impacts on measures of economic and social damage and important consequences for valuation. But tipping points are only one potential source of nonlinearity that could be sizable over long horizons. More general versions of this stylized model have been used in climate economics by several researchers who characterize some forms of nonlinearity and climate risk exposures including tipping point uncertainty. Most of this research abstracts from adjustments for robustness, however.

4.8 Examples of previous economic research that confront uncertainty

State of the art examples of climate economic models using recursive utility specifications are Cai et al. (2015) and Cai et al. (2016b). Cai et al. (2015) provide displays of the probability density of the social cost of carbon and other important economic quantities under stochastic factor productivity growth and stochastic climate outcomes showing how probabilistic uncertainty compounds over a time horizon for a variety of processes including atmospheric carbon and the social cost of carbon. Moreover, they document the sensitivity of model outputs to a range of preference, technology, and other parameters. This provides a nice example of risk analysis and computations that could support the impact of weighting results across alternative parameter configurations. Economic agents within the model have rational expectations as an equilibrium outcome. Given that some models with robustness concerns by economic agents sometimes have very similar implications to those in which agents have recursive utility preferences, some of their conclusions may extend to this decision making setting.

We now speculate on how worst-case probabilities would emerge if there was ambiguity across models and concerns about misspecification were added. We can only guess, but our Figure 7 is suggestive. For atmospheric carbon, the corresponding fan chart under the worst case model is likely to be shifted upwards and widen as a function of time horizon looking

forward. Our intuition behind this guess is that under the worst case climate model, policy makers face explicit uncertainty about the persistence of growth of atmospheric carbon as well as concerns about model misspecification of the earth system response to emissions. But such speculation can be checked and quantified by careful computation.

Hennlock (2009) and Anderson et al. (2016) consider robust planning problems in which the planner has concerns about model misspecification. They use a recursive version of exponential tilting described previously to represent the endogenous response to misspecification concerns, and they use the Cumulative Carbon Response (CCR) approach to cutting through the complexity of modeling the carbon cycle and the impact of emissions. By adopting a simplified but misspecified model, their approach reduces the policy problem to carbon budgeting using the approximate linear relation between cumulative emissions and planetary temperature.

Anderson et al. (2016) estimate the CCR parameter they use from data and, Anderson et al. (2003), link the preferences for robustness to statistical discrimination challenges on the part of the decision maker. But as we have seen the CCR parameter in this approximation varies dramatically across the alternative climate models simulated by MacDougall and Friedlingstein (2015), MacDougall et al. (2016) and others. Thus, in addition to model misspecification, there is a reason to explore robust model averaging by incorporating the style of quantitative policy analysis proposed by Brock et al. (2007) into a formal decision problem with an aversion to ambiguity along the lines we described previously. One illustration of such an approach applied to monetary policy is given in Cogley et al. (2008).

Millner et al. (2013) and others suggest application of the smooth ambiguity preferences as a way to confront ambiguity aversion to model averaging. Recall that the smooth ambiguity preferences, like those implied by the subjective utility model, rely on a prior across models. By adopting this starting point such preferences are not designed to address the concerns of robust Bayesian decision theory. On the other hand, they allow for a decision-maker's aversion to this uncertainty induced by this prior weighting to be distinct from that of the risk conditioned on alternative models. As we argued previously, however, some smooth ambiguity specifications can be interpreted as the outcome of a prior sensitivity analysis starting with a reference prior.

5 Long-term uncertainty

The rational expectations model presumes that investors know the probability model generating future outcomes but not the actual outcomes. It is the workhorse framework in macroeconomics, finance and in structural models linking economics and climate change with uncertainty.¹⁸

The rational expectations long-run risk literature opens the door to discussions of weak, or statistically subtle evidence, for these long-term components. Long-term macroeconomic exposures are challenging to quantify with much accuracy using time series statistical methods. If *outside* econometricians or applied researchers have limited confidence in this evidence, then how do economic agents inside the model become instilled with precise knowledge of the probabilities that underly these components?

The long-run risk literature in asset pricing has largely abstracted from what the source is for long-run risk beyond macroeconomic growth-rate risk. Climate risk offers a tangible source for this risk, albeit one that is hard to quantify. As we will see, the challenge of identifying random shocks with long-term components comes with two other challenges. One is the substantial statistical uncertainty in potentially important sensitivity parameters, and the other is models used both to assess the information in historical data and to provide revealing stylized approximations to more complex climate models. Such approximations are purposefully misspecified.

5.1 Historical time series evidence

One way to construct measurements based on historical data is to use vector autoregressions. For instance, see Hansen et al. (2008), but this evidence is itself fragile, and it does not include time series on emissions, radiative forcings, and temperature. While permanent shocks are identified by statistical methods, the rationale for these shocks is left implicit. On the other hand, as we have seen, it is the permanent shocks that are most consequential in terms of valuation. These same methods could also be applied to time series data on temperature along with different components of radiative forcing (natural and anthropogenic) as well as emissions, and could draw on some of the existing research.

In simplified climate models, random changes in emissions have permanent consequences

¹⁸See Collin-Dufresne et al. (2016) for a recent exception with parameter learning within a recursive utility model where unknown invariant parameters become a source of long-term uncertainty. They use subjective probabilities without consideration of ambiguity aversion.

for temperature. This is most evident in the stark Matthews approximation, but it is evident in many of the highly stylized models of the climate used in the economics literature. Given our interests, the multivariate time series approach is more pertinent for model purposes than the flexible fitting of time trends in temperature. Indeed, it is the sources of secular changes in temperature that are essential for policy relevant characterizations of the data. While aiming to be “robust,” the time trend approach skirts the source of the secular growth. We comment briefly on three recent time series papers that we find to be revealing, and we make no pretense at providing an exhaustive survey of the literature. Stern and Kaufmann (2014) fit a vector autoregression with time series data on temperature and components of radiative forcing to make some assessments of Granger-causality. Recall that Granger-causality measures dynamic feedback effects between time series. Perhaps not surprisingly, they find that fluctuations in the radiative forcing processes have important impacts on temperature. They find more modest evidence for feedback effects in the other direction whereby temperature has a notable impact on greenhouse gasses. This research purposefully avoids characterizing specifically how the various series grow together through so-called cointegration relations and instead features methods of inference that have some flexibility *vis a vis* stochastic trends. But as Hansen et al. (2008) found in their empirical investigation, understanding of the cointegration can be an important source for identify shocks with permanent consequences. Moreover in other applications, the coefficients that identify this relation can have interesting structural interpretations.

Poppick et al. (2016) use a time series regression framework with temperature as the left-hand side variable and components of radiative forcing on the right-hand side. This approach is justified when there are no feedback effects from temperature to greenhouse gasses or from temperature to greenhouse gasses in contrast to the evidence documented by Stern and Kaufmann (2014). While the absence of such effects may be difficult to justify in a more complete analysis of the climate system, their impact could be so small as to be inconsequential to some of their key measurements. After all, the Matthews approximation, in its aim to provide a useful and simple characterization of human impacts on temperature, also abstracts from these feedback effects. To their credit, Poppick et al. (2016) motivate and include some formal structure in their analysis, which allows them to isolate interpretable parameters. Since anthropogenic forcing displays secular growth, they in effect impose a cointegration relation used in a key sensitivity parameter relating changes in anthropogenic forcing to temperature. More generally, their formal modeling allows them to show how much information there is in the historical record for some climate sensitivity

parameters and document the practical challenge in inferring these from evidence. The imprecision of these estimates would seem to challenge the naive application of a rational expectations approach, one that endows the private sector with full knowledge of such parameters.

As a third approach, Storelvmo et al. (2016) also focus on the temperature evolution, but they bring in spatial or locational data as well. A cross-sectional average of their location model gives the temperature equation of a multivariate time series model.¹⁹ By allowing for aggregates to alter the location specific evolution equation, the statistical model used in this study introduces a common (across location) scalar variable that is presumed to absorb aggregate influences on the location specific model including stochastic time trends. Although the focus is on a temperature evolution equation, this approach allows for feedback effects without attempting to formally identify them. Full characterizations of long-run uncertainty are left to other studies, however. While Storelvmo et al. (2016) add a spatial dimension to the analysis, they also target a key sensitivity parameter with aggregate implications. Importantly, they show what the empirical evidence has to say about uncertainty in their estimation of this parameter.

5.2 Using climate models to generate data

Climate models are complex and costly to simulate. As we remarked earlier, this has led several authors to seek simpler representations including linear time series models represented as transfer functions from radiative forcing to temperature. The resulting models are represented as difference or differential equations. They can be fit to time series data, but also to data from climate models when researchers seek a good approximation of simple linear model to a highly complex nonlinear model.

Li et al. (2009) argue that in spite of the approximation, such models have interpretable coefficients. This is entirely consistent with what Poppick et al. (2016) argue when they estimate a linear time series with historical data. Li and Jarvis (2009) fit to simulations from the HADCM3 model using a dynamic regression model viewed as a mapping from CO_2 forcing history to temperature. They consider a single equation formulation without feedback and fit the resulting model via least squares. Presumably the errors in this analysis include model approximation errors induced by the simplification. Li and Jarvis (2009) allow these errors to be autocorrelated, which certainly seems well motivated. What is

¹⁹Their locations are on land, so the actual cross sectional average omits ocean locations, but the ocean contribution is included in the aggregates that enter into their econometric model.

more challenging to defend is that the model approximation errors are uncorrelated with current, past, and future values of the forcing input process.²⁰ Perhaps least squares without regard to serial correlation is better as a model approximation criterion. Relatedly, when the underlying model is nonlinear, the calibration output can be particularly sensitive to the numerical experiment chosen to generate the data from the model.

The challenge about how best to calibrate simple approximations to complicated but imperfect climate models is a nontrivial one. Li and Jarvis (2009) are to be commended for taking it more seriously than many economic calibrators do. But it is hard to determine what is a good model approximation without asking what the purpose is of the approximation. Thus decision theory can contribute to this challenge once we are specific as to the final goal of the analysis.

These studies abstract from the dynamics connecting emissions to radiative forcing. Similar issues emerge in the study of this dynamic relation. In fact it is the convolution of these two dynamic mappings that is typically taken as modeling inputs into economic analyses.

5.3 Uncertainty in model inputs

These and other findings lead us to embrace the broader perspective on uncertainty described previously. Shocks with long-term consequences can be challenging to identify, and their construction and measurement requires knowledge of sensitivity parameters that we only have limited knowledge about. Highly stylized models help to preserve tractability and support empirical investigation, but they are also acknowledged to be misspecified. All of this leads us to conceive of uncertainty in broader terms. By taking such a perspective, applied both to agents *inside* the model and to policy evaluation, the possibility of long-term adverse consequences remains a concern even if we are unsure as to the magnitude of the adverse consequences and when they might be realized. We need not endow investors and policy makers with confident quantifications but instead suggest they entertain the potential for sizable long-term adverse consequences as a possibility. Thus we aim to replace long-term risk with long-term uncertainty more broadly conceived.

²⁰This is needed to justify the pre-whitening approach used in Li and Jarvis (2009).

6 Heterogenous regional climate analysis

For simplicity, our exposition in this essay has focused on the aggregative planetary scale, but a major part of the literature in climate science is focused on regional impacts and the role of spatial transport of heat and moisture from the low latitudes to the high latitudes on regional climate change. See Alexeev and Jackson (2013), Hollesen et al. (2015) and Leduc et al. (2016) for further discussion.

Indeed Leduc et al. (2016) write

Ensemble mean regional TCRE values range from less than 1 degree Celcius per TtC for some ocean regions, to more than 5 degrees Celcius per TtC in the Arctic, with a pattern of higher values over land and at high northern latitudes. We find also that high-latitude ocean regions deviate more strongly from linearity as compared to land and lower-latitude oceans. The strong linearity of the regional climate response over most land regions provides a robust way to quantitatively link anthropogenic CO_2 emissions to local-scale climate impacts.

The larger higher latitude regional TCRES compared to the smaller lower latitude TCRES displayed by Leduc et al. (2016). Many of the potential tipping points, for instance potential permafrost melt, are in the high latitudes. Such tipping points may occur earlier than forecasted with models that neglect heat transport.

Some recent research has explored some of the potential policy implications of these findings. Brock and Xepapadeas (2017) built a small two region model with spatial heat and moisture transport and explore potential biases in optimal carbon taxes and the social cost of carbon. Xepapadeas and Yannacopoulos (2017) formally integrate robustness concerns into a spatial analysis and Cai et al. (2016a) explore potential biases from neglecting spatial heat and moisture transport using a spatial extension of the DSICE model of Cai et al. (2015).

These papers are just the start of a promising research program. Of course, much more work must be done to evaluate the quantitative importance of spatial and moisture heat transport to regional damages. Hsiang et al. (2017) write in their study of damages to the U.S. at the county level,

... The combined value of market and nonmarket damage across analyzed sector - agriculture, crime, coastal storms, energy, human mortality, and labor -

increases quadratically in global mean temperature, costing roughly 1.2 percent of gross domestic product per +1 degree Celsius on average. Importantly, risk is distributed unequally across locations, generating a large transfer of value northward and westward that increases economic inequality. By the late 21st century, the poorest third of counties are projected to experience damages between 2 and 20 percent of county income (90 percent chance) under business-as-usual emissions (Representative Concentration Pathway 8.5).

Hsiang et al.'s finding that warmer areas of the U.S. are likely to suffer more damage relates to left-hand side of Figure 6 (Burke et al. (2015), Figure 3), which emphasizes the nonlinear effects of damages under Representative Concentration Pathway 8.5.

While these conclusions remain highly speculative, they show a potential for substantial increases in regional inequality of climate change impacts, especially in heavily populated poorer low latitude coastal areas. They are suggestive of important future research for which uncertainty, broadly conceived, will have much to contribute.

7 Conclusions

Long-term uncertainty can have a big impact on both market determined stochastic discount factors and on the design of prudent policies for climate damage mitigation. Understanding better and acknowledging this uncertainty in climate economics will improve scientific discourse and help to nurture valuable quantitative research in the future. Using insights from modern decision theory to integrate the various components to uncertainty in climate economics will elevate discourse about research implications.

If there was only one model that was credible in climate science, and where the quality of the linear approximation was quite good, then economists could put most of their focus on uncertainties from the economic side and avoid having to do their own modeling on the climate side. Although the spread of TCRE values across the MacDougall et al. (2016) set of models shown in Figure 5 is large, the virtue is that this level of uncertainty is documented by expert climate scientists. So it makes sense for economists to use such characterizations of model uncertainty on the climate side and focus on uncertainties on the economics side where our relative expertise lies. It remains an open question in climate science as to how accurate both the large scale models are as well as their simplified approximations.

As we have repeatedly stressed in the essay, concerns about model misspecification are real, and the consequences in decision making are left to be explored. Defenses for

policies that combat climate damage externalities induced by human activity need not require precise knowledge of the magnitude or timing of the potential adverse impacts. The possibility of long-term damages that are extremely difficult if not impossible to reverse can justify a call to action. Waiting for precise knowledge of the eventual consequences of continued or expanded human induced CO_2 emissions could make mitigation or adaptation extremely costly.

A Production economy

Consider an economy with an AK technology that makes output be proportional to capital. Output can be allocated between investment and consumption. Suppose that there are adjustment costs to capital that are represented as the product of capital times a quadratic function of the investment-capital ratio. A robust planner chooses consumption-capital and investment-capital ratios. Given the constraint on output imposed by the AK technology, it suffices to let the planner choose the consumption-capital ratio. Capital optimally evolves as

$$dK_t = .01K_t \left[\mu_k(Z_t) - \vartheta_1 \frac{C_t}{K_t} dt - \frac{\vartheta_2}{2} \left(\frac{C_t}{K_t} \right)^2 dt + \sigma_k(Z_t) \cdot dW_t \right]$$

where K_t is the capital stock and the multiplication by .01 is included so that we view the term in square brackets as the percentage rate of growth. The capital evolution expressed in logarithms is

$$d \log K_t = .01 \left[\mu_k(Z_t) - \vartheta_1 \frac{C_t}{K_t} dt - \frac{\vartheta_2}{2} \left(\frac{C_t}{K_t} \right)^2 dt - \frac{|.01\sigma_k(Z_t)|^2}{2} dt + .01\sigma_k(Z_t) \cdot dW_t, \right]$$

where K_t is the capital stock. To interpret the right-hand-side of this evolution equation, notice that the zero consumption solution in which all of output is reinvested and not consumed,

$$dK_t = .01K_t\mu_k(Z_t)K_tdt + .01K_t\sigma_k(Z_t) \cdot dW_t.$$

With the quadratic adjustment cost specification, consumption reduces the drift in the capital stock by:

$$.01K_t \left[\vartheta_1 C_t + \frac{\vartheta_2}{2} \left(\frac{C_t}{K_t} \right)^2 \right].$$

As we will see, this specification will allow us to derive a production counterpart to the long-run risk specification featured in the consumption-based asset pricing literature. See Bansal and Yaron (2004). We allow for stochastic volatility in the macroeconomy, and will obtain quasi analytical formulas for square-root specifications as in Hansen (2012).

We suppose that temperature shifts proportionately consumption and capital by a multiplicative factor $\exp(-\gamma\tau_t)$ where τ_t is temperature relative to some preindustrial initial state. For instance the damage adjusted consumption is $C_t = \tilde{C}_t \exp(-\gamma\tau_t)$ and the damage adjusted capital is $K_t = \tilde{K}_t \exp(-\gamma\tau_t)$ where \tilde{C}_t and \tilde{K}_t are determined as before.

Let temperature evolve as:

$$d\tau_t = \lambda \mathcal{E}_t dt + \sqrt{\mathcal{E}_t} \sigma_\tau \cdot dW_t + d\tilde{\tau}_t$$

where $\tilde{\tau}_t = \iota \cdot Z_t$ and ι selects one component of the stationary $\{Z_t\}$ process. This process is simplified along the lines suggested by Matthews. Temperature does have a martingale component, which is consistent with Kaufmann et al. (2013), but this component becomes less important as the energy stock \mathcal{E}_t is reduced. There is a finite stock constraint for the energy:

$$R_0 = \int_0^\infty \mathcal{E}_u du.$$

This for simplicity, and in fact it makes good sense to include a backyard technology that becomes operational when the price of energy becomes sufficiently high.

Finally, the exogenous forcing process is:

$$dZ_t = \mu_z(Z_t)dt + \sigma_z(Z_t)'dW_t$$

We allow for robustness by allowing changes in probabilities that imply drift distortions $H_t dt$ in the Brownian increment dW_t . As the solutions to the stochastic differential equations are functions of the underlying Brownian motions, the probability measure changes for the Brownian motions change the probabilities implied by the solutions to the stochastic differential equations. As part of the robust decision problem, we provide recursive characterizations of the implied drift distortions that are most consequential for the decision maker.

Suppose that the instantaneous utility $\delta(1 - \alpha) \left[\log \left(\frac{C_t}{K_t} \right) + \log K_t - \gamma \tau_t \right] + \delta \alpha \log \mathcal{E}_t + \frac{\kappa}{2} |H_t|^2$. There are three controls \mathcal{E}_t and $\frac{C_t}{K_t}$ and H_t . In what follows we let lower case variable denote potential realized values controls and states except that c is use denote a possible value of the ratio: $\frac{C_t}{K_t}$.

To illustrate implications of robustness concerns, we consider a special case. Suppose that

$$\begin{aligned} \sigma_k \cdot \sigma_\tau &= 0 \\ (\sigma_z)' \sigma_\tau &= 0 \\ (\sigma_z)' \sigma_{\tilde{\tau}} &= 0. \end{aligned}$$

Guess a separable specification of the value function:

$$V(\log k, \tau, r, z) = \xi_k \log k - \xi_\tau \tau + v(r) + \nu(z)$$

Then the HJB equation after the h minimization is:

$$\begin{aligned} 0 = & -\delta [\xi_k \log k - \xi_\tau \tau + v(r) + \nu(z)] + \\ & \max_{c,e} \delta(1-\alpha) [\log c + \log k - \gamma\tau] + \delta\alpha \log e \\ & + \xi_k \left(.01 \left[\mu_k(z) - \vartheta_1 c - \frac{\vartheta_2}{2} c^2 \right] - \frac{|.01\sigma_k(z)|^2}{2} \right) \\ & - \xi_\tau [\lambda e + \iota \cdot \mu(z)] - \frac{dv}{dr}(r)e \\ & + \frac{\partial \nu}{\partial z}(z) \mu_z(z) + \frac{1}{2} \text{trace} \left[\sigma_z(z)' \frac{\partial^2 \nu}{\partial z \partial z'}(z) \sigma_z(z) \right] \\ & - \frac{e}{2\kappa} |\sigma_\tau|^2 (\xi_\tau)^2 \\ & - \frac{1}{2\kappa} \begin{bmatrix} \xi_k & -\xi_\tau & \frac{\partial \nu}{\partial z}(z)' \end{bmatrix} \begin{bmatrix} .01\sigma'_k(z) \\ \sigma_{\tilde{\tau}}(z)' \\ \sigma_z(z)' \end{bmatrix} \begin{bmatrix} .01\sigma_k(z) & \sigma_{\tilde{\tau}}(z) & \sigma_z(z) \end{bmatrix} \begin{bmatrix} \xi_k \\ -\xi_\tau \\ \frac{\partial \nu}{\partial z}(z) \end{bmatrix}. \end{aligned}$$

where $\sigma_{\tilde{\tau}} = (\sigma_z)' \iota$. The implied minimizing h is

$$h^* = -\frac{1}{\kappa} \begin{bmatrix} .01\sigma_k(z) & \sqrt{e}\sigma_\tau + \sigma_{\tilde{\tau}}(z) & \sigma_z(z)' \end{bmatrix} \begin{bmatrix} \xi_k \\ -\xi_\tau \\ \frac{\partial \nu}{\partial z}(z) \end{bmatrix}$$

In particular, the distortion for the capital and temperature dynamics are:

$$\textbf{capital evolution } .01\sigma_k(z) \cdot h^* = -\frac{.01}{\kappa} \sigma_k(z)' \left[(.01)\xi_k \sigma_k(z) - \xi_\tau \sigma_{\tilde{\tau}}(z) + \sigma_z(z) \frac{\partial \nu}{\partial z}(z) \right]$$

$$\textbf{temperature evolution } \sqrt{e}\sigma_\tau \cdot h^* + \sigma_{\tilde{\tau}}(z) \cdot h^* = \frac{\xi_\tau}{\kappa} (|\sigma_\tau|^2 e + |\sigma_{\tilde{\tau}}(z)|^2).$$

There are three contributions to this slope coefficient adjustment for the temperature equation. The term $\sigma_\tau \cdot \sigma_\tau$ is the conditional variance in the temperature equation. The term ξ_τ reflects how consequential temperature is as state variable for the value function. Finally $\frac{1}{\kappa}$ captures how important the concern is for robustness.

The first-order conditions for c are:

$$\delta(1 - \alpha)\frac{1}{c} - .01\xi_k(\vartheta_1 + \vartheta_2c) = 0.$$

Multiply the first-order condition for c by c . This gives a quadratic equation for c and hence two solutions. Only one of these solutions is positive. This can be seen because the quadratic function is positive and zero and the coefficient on the squared term is negative.

The first-order conditions for e are:

$$\frac{\delta\alpha}{e} - \frac{1}{2\kappa}|\sigma_\tau|^2(\xi_\tau)^2 - \xi_\tau\lambda - \frac{dv}{dr}(r) = 0.$$

Thus

$$\frac{1}{e^*} = \frac{1}{\delta\alpha} \left[\frac{dv}{dr}(r) + \xi_\tau\lambda + \frac{1}{2\kappa}|\sigma_\tau|^2(\xi_\tau)^2 \right].$$

Terms in the HJB equation involving v :

$$\delta\alpha(\log\delta + \log\alpha) - \delta\alpha \log \left[\frac{dv}{dr}(r) + \xi_\tau\lambda + \frac{1}{2\kappa}|\sigma_\tau|^2(\xi_\tau)^2 \right] - \delta v(r) - \delta\alpha = 0.$$

which is first-order differential equation. When $\xi_\tau = 0$,

$$\alpha \log \delta + \alpha \log \alpha - \alpha - \alpha \log \left[\frac{dv}{dr}(r) \right] = v(r).$$

then

$$v(r) = \alpha(\log\delta - 1 + \log r)$$

is a solution. In this case $e^* = \delta r$ and

$$\frac{dR}{dt} = -\delta R_t$$

which is the Hoteling rule with exponential decay in the resource stock.

Rewrite the differential equation as:

$$\frac{dv}{dr}(r) = \exp \left[-\frac{v(r)}{\alpha} + \log\delta + \log\alpha - 1 \right] - \xi_\tau\lambda - \frac{1}{2\kappa}|\sigma_\tau|^2(\xi_\tau)^2$$

From the HJB equation:

$$\begin{aligned}\xi_k &= 1 - \alpha \\ \xi_\tau &= \gamma(1 - \alpha)\end{aligned}$$

We may obtain affine solutions for ν by adopting an affine in z specification for drift coefficients. In addition, one may include a stochastic volatility state that evolves as a square root process as in Hansen (2012).

A.1 An Interesting Subproblem

Let

$$\gamma^* = \gamma\lambda + \frac{\gamma}{\kappa}|\sigma_\tau|^2\xi_\tau = \gamma\lambda + \frac{1}{\kappa}|\sigma_\tau|^2(1 - \alpha)\gamma^2$$

Consider a deterministic control problem solved at date zero

$$w(r) = \max_{\{\mathcal{E}_t\}} \delta \int_0^\infty \exp(-\delta t) \left[\alpha \log \mathcal{E}_t - (1 - \alpha) \int_0^t \gamma^* \mathcal{E}_s ds \right] dt$$

subject to

$$r = \int_0^\infty \mathcal{E}_t dt$$

where $R_0 = r$. Note that

$$\delta \int_0^\infty \exp(-\delta t) \int_0^t \mathcal{E}_s ds dt = \int_0^\infty \exp(-\delta s) \mathcal{E}_s ds$$

Thus we rewrite the control problem as:

$$w(r) = \max_{\{\mathcal{E}_t\}} \int_0^\infty \exp(-\delta t) [\delta \alpha \log \mathcal{E}_t - (1 - \alpha) \gamma^* \mathcal{E}_t] dt$$

subject to

$$r = \int_0^\infty \mathcal{E}_t dt.$$

The HJB equation is:

$$\delta w(r) = \delta \max_e \alpha \log e - (1 - \alpha) \gamma^* e - \frac{dw}{dr}(r) e$$

Thus

$$\frac{\delta\alpha}{e} - (1 - \alpha)\gamma^* - \frac{dw}{dr}(r) = 0.$$

Solving for e

$$e^* = \frac{\delta\alpha}{(1 - \alpha)\gamma^* + \frac{dw}{dr}(r)}.$$

Now solve the resource problem as a deterministic optimization problem in $\{E_t\}$ as a function of time.

$$\exp(-\delta t) \frac{\delta\alpha}{E_t} = \exp(-\delta t)(1 - \alpha)\gamma^* + \mu = 0$$

where μ is the multiplier on the resource constraint. Thus

$$E_t^* = \frac{\exp(-\delta t)\delta\alpha}{\exp(-\delta t)(1 - \alpha)\gamma^* + \mu}$$

Next, integrate and impose the resource constraint:

$$r = \int_0^\infty E_t^* dt = \int_0^\infty \left[\frac{\delta\alpha \exp(-\delta t)}{\exp(-\delta t)(1 - \alpha)\gamma^* + \mu} \right] dt.$$

Notice that

$$\int_0^\infty \left[\frac{\delta(1 - \alpha)\gamma^* \exp(-\delta t)}{\exp(-\delta t)(1 - \alpha)\gamma^* + \mu} \right] dt = \log[(1 - \alpha)\gamma^* + \mu] - \log \mu.$$

Thus

$$r = (\log[(1 - \alpha)\gamma^* + \mu] - \log \mu) \left[\frac{\alpha}{(1 - \alpha)\gamma^*} \right].$$

Solving for the multiplier μ ,

$$\mu = \frac{(1 - \alpha)\gamma^*}{\exp \left[\frac{(1 - \alpha)\gamma^*}{\alpha} r \right] - 1}$$

Moreover,

$$\frac{dw}{dr}(r) = \mu.$$

It remains to integrate $\frac{dw}{dr}(r)$ subject to an initial condition. The implied second derivative of w is concave as expected from the optimization problem.

Guess:

$$w(r) = \mathbf{a}_{-1} \log \left(\exp \left[\frac{(1 - \alpha)\gamma^*}{\alpha} r \right] - 1 \right) + \mathbf{a}_1 r + \mathbf{a}_0$$

Compute:

$$\begin{aligned}\frac{dw}{dr}(r) &= \mathbf{a}_{-1} \left[\frac{(1-\alpha)\gamma^*}{\alpha} \right] \frac{\exp \left[\frac{(1-\alpha)\gamma^*}{\alpha} r \right]}{\exp \left[\frac{(1-\alpha)\gamma^*}{\alpha} r \right] - 1} + \mathbf{a}_1 \\ &= \mathbf{a}_{-1} \left[\frac{(1-\alpha)\gamma^*}{\alpha} \right] \frac{\exp \left[\frac{(1-\alpha)\gamma^*}{\alpha} r \right]}{\exp \left[\frac{(1-\alpha)\gamma^*}{\alpha} r \right] - 1} + \mathbf{a}_1 \frac{\exp \left[\frac{(1-\alpha)\gamma^*}{\alpha} r \right] - 1}{\exp \left[\frac{(1-\alpha)\gamma^*}{\alpha} r \right] - 1}.\end{aligned}$$

Impose that

$$\mathbf{a}_{-1} \left[\frac{(1-\alpha)\gamma^*}{\alpha} \right] + \mathbf{a}_1 = 0,$$

and

$$\mathbf{a}_1 = -(1-\alpha)\gamma^*.$$

Therefore,

$$\mathbf{a}_{-1} = \alpha.$$

This gives a function with the correct derivative, but it leaves the constant \mathbf{a}_0 undetermined.

We restrict the coefficient \mathbf{a}_0 by looking at the constant term of the HJB equation. This gives:

$$\delta \mathbf{a}_0 = \delta \alpha \log(\delta \alpha) - \delta \alpha \log[(1-\alpha)\gamma^*] - \delta \alpha$$

or

$$\mathbf{a}_0 = \alpha [\log(\delta \alpha) - \log[(1-\alpha)\gamma^*] - 1].$$

This constructed w gives the v of interest for the original value function.

A.2 Pigouvian Tax

One measure of the social cost of carbon is the optimal Pigouvian tax:

$$\begin{aligned}& \left[\left(\lambda + \frac{1}{\kappa} |\sigma_\tau|^2 \right) \xi_\tau + \frac{dv}{dr}(r) \right] \frac{(1-\alpha)\tilde{C}_t}{\alpha \mathcal{E}_t} \\ &= \left[\left(\lambda + \frac{1}{\kappa} |\sigma_\tau|^2 \right) \gamma(1-\alpha) + \frac{dv}{dr}(r) \right] \frac{(1-\alpha)c^* K_t \exp(-\gamma \tau_t)}{\alpha \mathcal{E}_t}.\end{aligned}$$

The first term in the square brackets captures the contribution of emissions on temperature and the second term the contribution of emissions on the resource stock depletion.

Robustness considerations, as we have modeled them here, augment the presumed impact of emissions on temperature with the magnitude depending on $\frac{1}{\kappa}|\sigma_\tau|^2$. Setting $\kappa = \infty$ eliminates concerns about model misspecification. Thus in this simplified model, the coefficient λ governing the emissions impact on climate is effectively increased.

References

- Alexeev, Vladimir A and Craig H Jackson. 2013. Polar Amplification: Is Atmospheric Heat Transport Important? *Climate dynamics* 41 (2):533–547.
- Alvarez, F. and U. J. Jermann. 2004. Using Asset Prices to Measure the Cost of Business Cycles. *Journal of Political Economy* 112 (6):1223–1256.
- Anderson, E., 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 (1):68–123.
- Anderson, E. W., W. A. Brock, and A. Sanstad. 2016. Robust Consumption and Energy Decisions. Working paper, SSRN.
- Archer, D., M. Eby, V. Brovkin, A. Ridgwell, L. Cao, U. Mikolajewicz, K. Caldeira, K. Matsumoto, G. Munhoven, A. Montenegro, et al. 2009. Atmospheric Lifetime of Fossil Fuel Carbon Dioxide. *Annual Review of Earth and Planetary Sciences* 37.
- Arrow, K. J. 1951. Alternative Approaches to the Theory of Choice in Risk-Taking Situations. *Econometrica* 19 (4):404–437.
- Bansal, R. and A. Yaron. 2004. Risks for the Long Run: A Potential Resolution of Asset Pricing Puzzles. *The Journal of Finance* 59 (4):1481–1509.
- Bansal, R., M. Ochoa, and D. Kiku. 2016. Climate Change and Growth Risks. Tech. rep., National Bureau of Economic Research.
- Barreca, A., K. Clay, O. Deschenes, M. Greenstone, and J. S. Shapiro. 2016. Adapting to Climate Change: The Remarkable Decline in the US Temperature-Mortality Relationship over the Twentieth Century. *Journal of Political Economy* 124 (1):105–159.
- Berger, J. O. 1984. The Robust Bayesian Viewpoint. In *Robustness of Bayesian Analysis*, vol. 4 of *Studies in Bayesian Econometrics*, edited by Joseph B. Kadane, 63–144. North-Holland, Amsterdam.
- Bettis, O. D., S. Dietz, and N. G. Silver. 2017. The Risk of Climate Ruin. *Climatic Change* 140 (2):109–118.

- Borovička, J. and L.P. Hansen. 2016. Term Structure of Uncertainty in the Macroeconomy. In *Handbook of Macroeconomics*, vol. 2, edited by John B. Taylor and Harald Uhlig, chap. 20, 1641 – 1696. Elsevier.
- Borovička, J., L. P. Hansen, and J. A. Scheinkman. 2014. Shock Elasticities and Impulse Responses. *Mathematics and Financial Economics* 8 (4):333–354.
- Brock, W. and L. Mirman. 1972. Optimal Economic Growth and Uncertainty: The Discounted Case. *Journal of Economic Theory* 4 (3):479–513.
- Brock, W. and A. Xepapadeas. 2017. Climate Change Policy Under Polar Amplification. *European Economic Review* 94:263–282.
- Brock, W. A., S. N. Durlauf, and K. D. West. 2007. Model Uncertainty and Policy Evaluation: Some Theory and Empirics. *Journal of Econometrics* 136:629–664.
- Burke, M., S. M. Hsiang, and E. Miguel. 2015. Global Non-linear Effect of Temperature on Economic Production. *Nature* 527 (7577):235–239.
- Burns, A. F. and W. C. Mitchell. 1946. *Measuring Business Cycles*. National Bureau of Economic Research, Inc.
- Cai, Y., K. L. Judd, and T. S. Lontzek. 2015. The Social Cost of Carbon with Economic and Climate Risks. *ArXiv e-prints* .
- Cai, Y., W. Brock, and A. Xepapadeas. 2016a. Climate Change Economics and Heat Transport Across the Globe: Spatial-DSICE. In *2017 Allied Social Science Association (ASSA) Annual Meeting, January 6-8, 2017, Chicago, Illinois*, 251833. Agricultural and Applied Economics Association.
- Cai, Y., T. M. Lenton, and T. S. Lontzek. 2016b. Risk of Multiple Interacting Tipping Points Should Encourage Rapid CO2 Emission Reduction. *Nature Clim. Change* 6 (5):520–525.
- Campbell, J. Y. and T. Vuolteenaho. 2004. Bad Beta, Good Beta. *American Economic Review* 94 (5):1249–1275.
- Chen, Z. and L. Epstein. 2002. Ambiguity, Risk, and Asset Returns in Continuous Time. *Econometrica* 70 (4):1403–1443.

- Cogley, T., R. Colacito, L. P. Hansen, and T. J. Sargent. 2008. Robustness and US Monetary Policy Experimentation. *Journal of Money, Credit and Banking* 40 (8):1599–1623.
- Collin-Dufresne, P., M. Johannes, and L. A. Lochstoer. 2016. Parameter Learning in General Equilibrium: The Asset Pricing Implications. *American Economic Review* 106 (3):664–98.
- Dietz, S., C. Gollier, and L. Kessler. 2017. The Climate Beta. Tech. Rep. 215, Centre for Climate Change Economics and Policy.
- Eberly, J. and N. Wang. 2009. Capital Reallocation and Growth. *American Economic Review* 99 (2):560–566.
- Epstein, L. G. and M. Schneider. 2003. Recursive Multiple-Priors. *Journal of Economic Theory* 113 (1):1–31.
- Epstein, Larry G. and Stanley E. Zin. 1989. Substitution, Risk Aversion, and the Temporal Behavior of Consumption and Asset Returns: A Theoretical Framework. *Econometrica* 57 (4):937–969.
- de Finetti, B. 1937. La Prevision: Ses Lois Logiques, Ses Sources Subjectives. *Annales de l'Institut Henri Poincaré* 7:1 – 68. English translation in Kyburg and Smokler (eds.), *Studies in Subjective Probability*, Wiley, New York, 1964.
- Frisch, R. 1933. Propagation Problems and Impulse Problems in Dynamic Economics. In *Economic Essays in Honour of Gustav Cassel*, 171–205. Allen and Unwin.
- Fyfe, J. C., G. A. Meehl, M. H. England, M. E. Mann, B. D. Santer, G. M. Flato, E. Hawkins, N. P. Gillett, S. Xie, and Y. Kosaka. 2016. Making Sense of the Early-2000s Warming Slowdown. *Nature Climate Change* 6 (3):224–228.
- Gilboa, I. and D. Schmeidler. 1989. Maxmin Expected Utility with Non-Unique Prior. *Journal of Mathematical Economics* 18 (2):141–153.
- Glottter, M. J., R. T. Pierrehumbert, J. W. Elliott, N. J. Matteson, and E. J. Moyer. 2014. A Simple Carbon Cycle Representation for Economic and Policy Analyses. *Climatic Change* 126 (3-4):319–335.
- Golosov, M., J. Hassler, P. Krusell, and A. Tsyvinski. 2014. Optimal Taxes on Fossil Fuel in General Equilibrium. *Econometrica* 82 (1):41–88.

- Hansen, L. P. 2012. Dynamic Valuation Decomposition within Stochastic Economies. *Econometrica* 911-967 (3). Fisher-Schultz Lecture at the European Meetings of the Econometric Society.
- . 2014. Nobel Lecture: Uncertainty Outside and Inside Economic Models. *Journal of Political Economy* 122 (5):945 – 987.
- Hansen, L. P. and S. F. Richard. 1987. The Role of Conditioning Information in Deducing Testable. *Econometrica* 55 (3):587–613.
- Hansen, L. P. and T. J. Sargent. 1995. Discounted Linear Exponential Quadratic Gaussian Control. *IEEE Transactions on Automatic Control* 40 (5):968–971.
- . 2001. Robust Control and Model Uncertainty. *The American Economic Review* 91 (2):60–66.
- . 2007. Recursive Robust Estimation and Control Without Commitment. *Journal of Economic Theory* 136:1–27.
- . 2016. Prices of Macroeconomic Uncertainties with Tenuous Beliefs. Tech. rep., University of Chicago.
- Hansen, L. P., T. J. Sargent, and T. D. Tallarini. 1999. Robust Permanent Income and Pricing. *Review of Economic Studies* 66 (4):873–907.
- Hansen, L. P., J. C. Heaton, and N. Li. 2008. Consumption Strikes Back? Measuring Long-run Risk. *Journal of Political economy* 116 (2):260–302.
- Hansen, L.P. and M. Marinacci. 2017. Ambiguity Aversion and Model Misspecification: An Economic Perspective. *Statistical Science* .
- Hennlock, M. 2009. Robust Control in Global Warming Management: An Analytical Dynamic Integrated Assessment. Tech. rep., Resources for the Future.
- Hollesen, J., H. Matthiesen, A. B. Møller, and B. Elberling. 2015. Permafrost Thawing in Organic Arctic Soils Accelerated by Ground Heat Production. *Nature Climate Change* 5 (6):574–578.

- Hsiang, S., R. Kopp, A. Jina, J. Rising, M. Delgado, S. Mohan, D.J. Rasmussen, R. Muir-Wood, P. Wilson, M. Oppenheimer, et al. 2017. Estimating Economic Damage From Climate Change in the United States. *Science* 356 (6345):1362–1369.
- 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 of *Studies in Logic and the Foundations of Mathematics*, edited by Patrick Suppes Ernest Nagel and Alfred Tarski, 232 – 239. Elsevier.
- Jacobson, D. H. 1973. Optimal Stochastic Linear Systems with Exponential Performance Criteria and Their Relation to Deterministic Differential Games. *IEEE Transactions for Automatic Control* AC-18 (2):1124–131.
- Kaufmann, R. K., H. Kauppi, M. L. Mann, and J. H Stock. 2013. Does Temperature Contain a Stochastic Trend: Linking Statistical Results to Physical Mechanisms. *Climatic Change* 118 (3-4):729–743.
- Kirtman, B., S. B. Power, A. J. Adedoyin, G. J. Boer, R. Bojariu, I. Camilloni, F. Doblas-Reyes, A.M. Fiore, M. Kimoto, G. Meehl, et al. 2013. *IPCC Climate Change 2013: The Physical Science Basis*, chap. 11 Near-Term Climate Change: Projections and Predictability. Cambridge University Press.
- Klibanoff, P., M. Marinacci, and S. Mukerji. 2005. A Smooth Model of Decision Making Under Uncertainty. *Econometrica* 73 (6):1849–1892.
- . 2009. Recursive Smooth Ambiguity Preferences. *Journal of Economic Theory* 144:930–976.
- Knight, F. H. 1921. *Risk, Uncertainty, and Profit*. Houghton Mifflin.
- Kocherlakota, N. R. 1990. Disentangling the Coefficient of Relative Risk Aversion from the Elasticity of Intertemporal Substitution: An Irrelevance Result. *The Journal of Finance* 45 (1):175–190.
- Koopmans, T. C. 1947. Measurement Without Theory. *The Review of Economics and Statistics* 29 (3):161–172.
- . 1960. Stationary Ordinal Utility and Impatience. *Econometrica* 28 (2):287–309.

- Kreps, D. M. and E. L. Porteus. 1978. Temporal Resolution of Uncertainty and Dynamic Choice Theory. *Econometrica* 46 (1):185–200.
- Leduc, M., H. D. Matthews, and R. de Elia. 2015. Quantifying the Limits of a Linear Temperature Response to Cumulative CO2 Emissions. *Journal of Climate* 28 (24):9955–9968.
- Leduc, M., H. D. Matthews, and R. de Elía. 2016. Regional Estimates of the Transient Climate Response to Cumulative CO2 Emissions. *Nature Climate Change* 6 (5):474–478.
- Li, S. and A. Jarvis. 2009. Long Run Surface Temperature Dynamics of an A-OGCM: the HadCM3 4XCO2 Forcing Experiment Revisited. *Climate Dynamics* 33 (6):817–825.
- Li, S., A. J. Jarvis, and D. T. Leedal. 2009. Are Response Function Representations of the Global Carbon Cycle Ever Interpretable? *Tellus B* 61 (2):361–371.
- Lucas, R. 1976. Econometric Policy Evaluation: A Critique. *Carnegie-Rochester Conference Series on Public Policy* 1 (1):19–46.
- Lucas, R. E. and E. C. Prescott. 1971. Investment Under Uncertainty. *Econometrica* 39 (5):659–681.
- Maccheroni, F., M. Marinacci, and A. Rustichini. 2006a. Dynamic Variational Preferences. *Journal of Economic Theory* 128 (1):4–44.
- Maccheroni, F., M. Marinacci, and A. Rustinchini. 2006b. Ambiguity Aversion, Robustness, and the Variational Representation of Preferences. *Econometrica* 74 (6):1147–1498.
- MacDougall, A. H. and P. Friedlingstein. 2015. The Origin and Limits of the Near Proportionality Between Climate Warming and Cumulative CO2 Emissions. *J. Climate* 28 (10):4217–4230.
- MacDougall, A. H., N. C. Swart, and R. Knutti. 2016. The Uncertainty in the Transient Climate Response to Cumulative CO2 Emissions Arising from the Uncertainty in Physical Climate Parameters. *Journal of Climate* forthcoming (0):null.
- Maenhout, P. J. 2004. Robust Portfolio Rules and Asset Pricing. *Review of Financial Studies* 17 (4):951–983.

- Marschak, J. 1953. Economic Measurements for Policy and Prediction. In *Cowles Commission Monograph 14: Studies in Econometric Methods*, edited by W. C. Hood and T. C. Koopmans, 1–26. Wiley.
- Matthews, H. D., N. P. Gillett, and K. Zickfeld. 2009. The Proportionality of Global Warming to Cumulative Carbon Emissions. *Nature* 459:829–832.
- Matthews, H. D., S. Solomon, and R. Pierrehumbert. 2012. Cumulative Carbon as a Policy Framework for Achieving Climate Stabilization. *Philosophical Transactions: Mathematical, Physical and Engineering Sciences* 370 (1974):4365–4379.
- Millner, A., S. Dietz, and G. Heal. 2013. Scientific Ambiguity and Climate Policy. *Environmental Resource Economics* 55 (21-46).
- Nordhaus, W. D. 2017. Revisiting the Social Cost of Carbon. *Proceedings of the National Academy of Sciences* 114 (7):1518–1523.
- North, G. R. and R. F. Cahalan. 1981. Predictability in a Solvable Stochastic Climate Model. *Journal of the Atmospheric Sciences* 38 (3):504–513.
- Petersen, I. R., M. R. James, and P. Dupuis. 2000. Minimax Optimal Control of Stochastic Uncertain Systems with Relative Entropy Constraints. *IEEE Transactions on Automatic Control* 45:398–412.
- Pierrehumbert, R. T. 2014. Short-Lived Climate Pollution. *Annual Review of Earth and Planetary Science* 42:341–379.
- Poppick, A., E. J. Moyer, and M. L. Stein. 2016. Estimating Trends in the Global Mean Temperature Record. *Working Paper* .
- Prein, A. F., W. Langhans, G. Fosser, A. Ferrone, N. Ban, K. Goergen, M. Keller, M. Tölle, O. Gutjahr, F. Feser, E. Brisson, S. Kollet, J. Schmidli, N. P. M. van Lipzig, and R. Leung. 2015. A Review on Regional Convection-Permitting Climate Modeling: Demonstrations, Prospects, and Challenges. *Reviews of Geophysics* 53 (2):323–361. 2014RG000475.
- Ross, S. A. 1976. The Arbitrage Theory of Capital Asset Pricing. *Journal of Economic Theory* 13 (3):341 – 360.

- Sargent, T. J. and N. Wallace. 1975. "Rational" Expectations, the Optimal Monetary Instrument, and the Optimal Money Supply Rule. *Journal of Political Economy* 83 (2):241–254.
- Savage, L. J. 1954. *The Foundations of Statistics*. New York: John Wiley and Sons.
- Segal, U. 1990. Two-Stage Lotteries Without the Reduction Axiom. *Econometrica: Journal of the Econometric Society* 349–377.
- Stern, D. I. and R. K. Kaufmann. 2014. Anthropogenic and Natural Causes of Climate Change. *Climatic Change* 122 (1-2):257–269.
- Storelvmo, T., T. Leirvik, U. Lohmann, P. C. B. Phillips, and M. Wild. 2016. Disentangling Greenhouse Warming and Aerosol Cooling to Reveal Earth's Climate Sensitivity. *Nature Geoscience Letter* 9 (4):286–289.
- Van Vuuren, D. P., J. Edmonds, M. Kainuma, K. Riahi, A. Thomson, K. Hibbard, G. C. Hurtt, T. Kram, V. Krey, J.F. Lamarque, et al. 2011. The Representative Concentration Pathways: An Overview. *Climatic change* 109:5–31.
- Von Neumann, J. and O. Morgenstern. 1944. *Theory of Games and Economic Behavior*. Princeton University Press.
- Wald, A. 1950. *Statistical Decision Functions*. New York: John Wiley and Sons.
- Watson, J. and C. Holmes. 2016. Approximate Models and Robust Decisions. *Statistical Science* forthcoming.
- Whittle, P. 1981. Risk Sensitive Linear Quadratic Gaussian Control. *Advances in Applied Probability* 13 (4):764–777.
- Xepapadeas, A. and A. Yannacopoulos. 2017. Spatially Structured Deep Uncertainty, Robust Control, and Climate Change Policies. Tech. rep., Athens University of Economics and Business.