

Annual Review of Economics

Uncertainty Spillovers for Markets and Policy

Lars Peter Hansen

Departments of Economics and Statistics and Booth School of Business, University of Chicago,
Chicago, Illinois 60637, USA; email: lhansen@uchicago.edu

Annu. Rev. Econ. 2021. 13:371–96

The *Annual Review of Economics* is online at
economics.annualreviews.org

<https://doi.org/10.1146/annurev-economics-082020-054645>

Copyright © 2021 by Annual Reviews.
All rights reserved

JEL codes: C18, C58, C61, D81, E44, E62, G12

Keywords

uncertainty, valuation, financial markets, policy, climate change, ambiguity, misspecification

Abstract

We live in a world filled with uncertainty. In this essay, I show that featuring this phenomenon more in economic analyses adds to our understanding of how financial markets work and how best to design prudent economic policy. This essay explores methods that allow for a broader conceptualization of uncertainty than is typical in economic investigations. These methods draw on insights from decision theory to engage in uncertainty quantification and sensitivity analysis. Uncertainty quantification in economics differs from uncertainty quantification in most sciences because there is uncertainty from the perspective both of an external observer and of people and enterprises within the model. I illustrate these methods in two example economies in which the understanding of long-term growth is limited. One example looks at uncertainty ramifications for fluctuations in financial markets, and the other considers the prudent design of policy when the quantitative magnitude of climate change and its impact on economic opportunities are unknown.

ANNUAL
REVIEWS **CONNECT**

www.annualreviews.org

- Download figures
- Navigate cited references
- Keyword search
- Explore related articles
- Share via email or social media

1. INTRODUCTION

We must infer what the future situation would be without our interference, and what changes will be wrought by our actions. Fortunately, or unfortunately, none of these processes is infallible, or indeed ever accurate and complete.

—Frank H. Knight (1921, pp. 201–2)

We live in an environment in which uncertainty is pervasive. In the first decade of this century, we experienced a financial crisis. The magnitude of the crisis caught many policy makers, academics, entrepreneurs, and businesses by surprise, given the previous extended period of so-called great moderation. As the COVID-19 pandemic unfolded at the end of the second decade, we have been living through uncertainties such as (a) what exposure we have to the virus when we are in public places, (b) how soon infections will show a permanent decrease, and (c) how soon a reliable vaccine will be available. These two recent crises and other episodes reveal the limits to our knowledge in ways that have a potentially important impact on the design of policy.

I read Knight's comment from 1921 as a serious challenge to quantitative dynamic modeling in economics. His century-old concern continues to apply today. Economists, as analysts external to their models, face "outside-the-model" uncertainty when making inferences on limited data. Economic agents, people and businesses, also face uncertainty when making decisions, such as consumption and investment choices, not knowing for sure how the future will unfold. As discussed in my previous work (Hansen 2014), my interest is in models with agents who are cognizant of this "inside-the-model" uncertainty when making forward-looking decisions. The internal perspective on uncertainty adds an important twist to the external uncertainty quantification that is often addressed in other scientific disciplines.

How, then, can we capture broad notions of uncertainty in formal ways that are amenable to quantitative modeling? In this essay, I draw on results from decision theory under uncertainty extended to dynamic settings. Following this literature, I push beyond the risk analyses that are a familiar tool for most economists. Under what I and others call risk, the probabilities of future outcomes are known with full confidence but not the outcomes themselves. The broader notion of uncertainty that I embrace includes both ambiguity across alternative possible models and the potential misspecification of each. Under the decision theories that I apply, there can be uncertainty-modeling inputs such as subjective probabilities, as in the case of Bayesian priors over alternative models. Instead of having full confidence in a complete probabilistic specification, a decision maker conducts a sensitivity analysis over families of models or over the probabilities assigned to such families. A decision maker explores the consequences of misspecification by substantially broadening the family of models under consideration. In line with the familiar aversion to risk, this decision theory introduces an aversion to ambiguity over models or to their potential misspecifications. This aversion effectively dictates how extensive the sensitivity analysis is: The greater the aversion, the less restrictive the sensitivity analysis is. The approach is made operational by using a decision problem to isolate the most troublesome model or the probabilities over the models subject to the restrictions imposed on the sensitivity analysis.

Rather than delineating axiomatic defenses for the alternative formulations of decision theory under uncertainty, I illustrate their importance in two example economies. Both entail uncertainty about the long-term forces driving the macroeconomy. Section 4 shows how to generate volatility in market valuation by letting investors inside the model be uncertain about future macroeconomic growth prospects. Section 5 studies the social cost of carbon by considering a fictitious social planner assessing the economic consequences of climate change when there is both geoscientific and economic uncertainty.

Before turning to the two applications, in Section 2 I discuss more the quantitative modeling with highly stylized models. In Section 3, I show how to use asset pricing theory to quantify the impact of long-term uncertainty in valuation over alternative cash-flow horizons. In this section, I construct a martingale component to stochastic discount factor processes used for both market and social valuation. This component dominates valuation adjustments for uncertainty over long horizons and can have important implications also over shorter horizons. The uncertainty adjustments to be characterized in Sections 4 and 5 contribute prominently to this martingale component.

2. QUANTITATIVE STORYTELLING

Hominem unius libri timeo [I fear the man of a single book].

—St. Thomas Aquinas

Structural models in economic dynamics are meant to be tools that can help us gain a better understanding of a variety of phenomena and provide guidance for designing and evaluating economic policy challenges. The models that interest me in this essay are models that are quantitative, that allow for random impulses, and that therefore are stochastic in nature. Under rational expectations, the economic decision makers within the models recognize the existence of these shocks in the future and know their probabilities. That is, they confront what I (and many predecessors) formally call risk.

I find it attractive to think of each such model as providing a quantitative story.¹ For a lot of questions in economics, I am reminded of the quote of St. Thomas Aquinas mentioned at the outset of this section. Here, I take the liberty to extend this quote to warn any decision maker who is fully committed to one model and its associated story. Thus, for many purposes, I find it valuable to entertain multiple model specifications (reflecting different economic structures or alternative parameter configurations), which leads me to quantitative storytelling with multiple stories. Looking across model specifications or constructions is often done informally, and potentially across specific constructions, by different research groups.² While it may make good sense for the current review to drill down on the implications of a single model, there are many settings in which a decision maker would be wise to look across models produced by different researchers or research groups.

I explore the consequences of doing this more formally and within the actual model construction. In the first of my two substantive examples, I will have investors inside the model entertain parameter uncertainty, including the possibility of time-varying parameters governing economic growth. In the second example, a hypothetical policy maker aims to address climate change in the presence of limited knowledge about the impact of CO₂ emissions on future economic opportunities.

I use the term “ambiguity” to capture the uncertainty associated with how much weight should be assigned to the alternative models. One elegant solution supported by an axiomatic defense

¹The term “quantitative storytelling” has been used elsewhere in a different but partially overlapping way (see, e.g., Saltelli & Giampietro 2017). I do not intend to address nor do I fully embrace Saltelli & Giampietro’s (2017) critiques of evidence-based policy; but I do share their concern about committing to one model with excessive confidence when drawing on evidence and prior insights to derive optimal policies.

²An obvious exception is the extensive literature on Bayesian estimation of stochastic general equilibrium models under rational expectations. This literature confronts unknown parameters with subjective probability inputs, but typically decision makers within these models are presumed to know the unknown parameters, in contrast to econometricians.

is the subjective probability approach of de Finetti (1937) and Savage (1954). This argues for assigning subjective probabilities or a prior over alternative models, which can be updated by Bayes's theorem based on evidence. But a decision maker who is unsure about these subjective inputs may wish to explore the consequences of this uncertainty. This is the perspective of a robust Bayesian. In some circumstances, the data may be so rich that the posterior probabilities are not very sensitive to the prior inputs, in which case the consequences of ambiguity may be minimal. We might expect this phenomenon to be pervasive in our current big-data environment, with machine learning algorithms helping to assess the data implications; but often the data are not rich along all dimensions of interest. It seems far too severe to limit our research questions to the ones for which sensitivity over a substantively interesting set of priors is of little consequence.³ Both examples I choose in this essay are cases in which, I argue, prior sensitivity remains an important consideration.

My metaphorical use of the term “story” for a model is in part due to the fact that a model is necessarily misspecified. It is a simplification or an abstraction that helps us think about the world or how to design more prudent policies, and it is not intended as a complete accounting of a complex economic environment. Designing omnibus statistical tests for misspecification seems to be a bit beside the point, unless these tests direct researchers to salient and substantive improvements in the model specification. Including a rich variety of random shocks to make models harder to reject is a common practice, but it leaves open the question of whether these are best conceived of as exogenous impulses or as devices to disguise model approximation errors. Excessive shock proliferation can quickly make the model look like a black box. A harder, but I think important, challenge is to explore systematically the potential consequences of model misspecification. What potential consequences concern us? To address this question, it is natural to think in terms of a decision problem.

3. SOME ASSET PRICING THEORY

The manner in which risk operates upon time preference will differ, among other things, according to the particular periods in the future to which the risk applies.

—Irving Fisher (1930, p. 83)

I use an asset pricing perspective to address problems associated with both social and private valuation. Because the primary focus of my essay is on uncertainty with long-term consequences, I will describe a valuation framework that takes this into account. This framework will allow me to address two questions:

- When will long-term uncertainty have short-term implications for markets?
- How does discounting compound over long horizons?

It is well known from the simplest exogenous growth models that both the subjective rate of discount and the intertemporal elasticity of substitution affect how private and social cash flows get discounted optimally or in equilibrium. Uncertainty contributes to discounting as well. A stochastic counterpart replaces the single per-period discount rate because social or private payoffs that are exposed to uncertainty in different ways should be discounted differently. Conceptually broadening the uncertainty that markets and policy makers face alters the stochastic contributions to

³Of course, dogmatic priors cannot be dominated by the data. In my first example, I entertain such priors as possibilities.

valuation. Thus, my push beyond the risk mentioned by Irving Fisher in the quote above will be important for the applications I consider.

3.1. Prior Literature

The framework I describe was originally inspired by a quantitative literature in macroeconomics and finance that focuses on the contributions of risk and uncertainty to valuation over alternative investment horizons. This entire literature is a direct outgrowth of Rubinstein's (1976) fundamental paper on valuing flows that grow stochastically over time. The quantitative contributions include, for instance, the works by Parker & Julliard (2003), Campbell & Vuolteenaho (2004), Alvarez & Jermann (2005), Lettau & Wachter (2007), Hansen et al. (2008), Bansal et al. (2009), and van Binsbergen et al. (2012). Much of this quantitative/empirical research works with log-linear or log-normal approximations. I will not attempt to survey this literature, but instead I will suggest a framework for understanding the valuation implications over alternative investment horizons that is both mathematically convenient and able to accommodate nonlinearity. I find this to be a valuable starting point for thinking about how growth rate uncertainty can affect valuation.

3.2. Framework

I represent valuation by taking as given the equilibrium solution to a dynamic stochastic equilibrium model. I begin with some notation: Let $X = \{X_t : t = 0, 1, \dots\}$ be a state vector process including, for instance, endogenous capital and exogenous shifts in technology; let $W = \{W_t : t = 1, 2, \dots\}$ denote a k -dimensional vector of independent and identically distributed shocks to the macroeconomy; and let \mathfrak{F}_t represent the information generated by histories of W and the initial state X_0 . For the time being, I presume that the model has been written in the form of a stochastic steady state and that the state vectors are appropriately scaled to capture this. Thus, for now I take the state vector X to be stationary, although the tools I expound do not require this restriction.⁴

From the solution of the economic model, I write

$$\begin{aligned} X_{t+1} &= \psi(X_t, W_{t+1}) \text{ and} \\ Y_{t+1} - Y_t &= \kappa(X_t, W_{t+1}), \end{aligned} \tag{1}$$

where $Y = \{Y_t : t = 0, 1, \dots\}$ is a scalar process with stationary increments, and Y_0 is initialized based on the date-zero information captured by \mathfrak{F}_0 . The first equation gives the state vector dynamics as a first-order Markov process for the state dynamics. I will use the dynamics in this second equation flexibly to represent alternative economic and financial constructs, including the logarithms of equilibrium consumption, investment, or capital as well as the logarithm of stochastically growing cash flows and stochastic discount factors. We take these latter two constructs to be ingredients to valuation.

3.3. Valuation

We exploit the recursive structure of the Markov formulation. For valuation, it is important that we work in levels and not logarithms. Let us consider two applications of $\exp(Y)$, a stochastic factor process S and a stochastic cash-flow process G . The product SG has the same mathematical

⁴For the application to climate change that I develop in Section 5, I necessarily feature the transition dynamics, but I continue to exploit the Markov specification.

structure. That is, if $\log S$ and $\log G$ have representations given by Equation 1 for a κ_s and κ_g , respectively, then so does their sum κ . Consider the price at time t of an asset with payoff $G_{t+\tau}$:

$$\begin{aligned}\pi_{t,t+\tau}(G) &= \mathbb{E} \left[\left(\frac{S_{t+\tau}}{S_t} \right) G_{t+\tau} \mid \mathfrak{F}_t \right] \\ &= \mathbb{E} \left[\left(\frac{S_{t+\tau} G_{t+\tau}}{S_t G_t} \right) \mid \mathfrak{F}_t \right] G_t.\end{aligned}\tag{2}$$

In these formulas, $\frac{S_{t+\tau}}{S_t}$ is the stochastic discount factor over a horizon τ starting from date t . The valuation of equity sums over such prices for $\tau = 1, 2, \dots$

With this computation in mind, we are led to study conditional expectations of the product GS as well as the S and G components over alternative horizons in constructing proportional uncertainty compensations, as I did in previous work (Hansen 2012). I express these in logarithms to obtain

$$\rho_{t,t+\tau}(G) = \left(\frac{1}{\tau} \right) \left(\log \mathbb{E} \left[\left(\frac{G_{t+\tau}}{G_t} \right) \mid \mathfrak{F}_t \right] - \log \mathbb{E} \left[\left(\frac{S_{t+\tau} G_{t+\tau}}{S_t G_t} \right) \mid \mathfrak{F}_t \right] + \log \mathbb{E} \left[\left(\frac{S_{t+\tau}}{S_t} \right) \mid \mathfrak{F}_t \right] \right).$$

The measure $\rho_{t,t+\tau}(G)$ is the logarithm of the ratio of two expected returns, one with payoff $G_{t+\tau}$ and the other one with risk-less payoff normalized to be 1. To see this, observe that the first term on the right is the expected payoff contribution, and the second one is the corresponding cost of a claim to this payoff. Because the second one comes with a negative sign, the two terms taken together give the logarithm of the expected return over horizon τ . The negative of the third term is the logarithm of the risk-free return over horizon τ , and thus it provides a risk-less comparison with the same horizon.

This uncertainty compensation measure, $\rho_{t,t+\tau}(G)$, depends on both the horizon τ and the conditioning information \mathfrak{F}_t available at time t . In particular, the dependence on τ reflects the impact of compounding. The compensation depends on the uncertainty in both the payoff process G and the stochastic discount factor S . The stochastic discount factor contribution reflects the price impact, and this will be of particular interest in this essay.

3.4. A Revealing Factorization

To set the stage for future discussion, I describe a factorization result that was originally formalized by Hansen & Scheinkman (2009) and motivated by previous research by Alvarez & Jermann (2005).⁵ I write the process $\exp(Y)$ as the product of three terms,

$$\exp(Y_t - Y_0) = \exp(\eta t) \left(\frac{M_t}{M_0} \right) \left[\frac{e(X_0)}{e(X_t)} \right],\tag{3}$$

where M is a positive martingale for which $\log M$ has stationary increments with a stochastic evolution given by Equation 1 with an appropriate choice of κ . In comparison, $\log e(X)$ inherits the stationarity of X . Notice that M and e are both defined only up to scale. The real number η gives a long-term growth rate or decay rate depending on the application. For instance, η is negative when $\exp(Y) = S$.

⁵Although Hansen & Scheinkman (2009) derive their result in a continuous-time environment, there is a straightforward extension to discrete-time settings. Qin & Linetsky (2020) relax the Markovian assumption in their extension of Hansen & Scheinkman's (2009) framework.

As stated, this representation may not be unique. For applications and for establishing the uniqueness of the factorization, it is most advantageous to use M to induce a change in the probability measure. Since M is a positive martingale, we have

$$\mathbb{E} \left(\frac{M_{t+1}}{M_t} \mid \mathfrak{F}_t \right) = 1.$$

In light of this outcome, we use an approach that is familiar in the pricing of derivative claims by building a different probability measure with conditional expectation,

$$\tilde{\mathbb{E}}(B_{t+1} | \mathfrak{F}_t) = \mathbb{E} \left[\left(\frac{M_{t+1}}{M_t} \right) B_{t+1} \mid \mathfrak{F}_t \right],$$

where B_{t+1} is a bounded random variable in the conditioning information at time $t + 1$ of decision makers in the economic model. It may be verified that this change in probability measure preserves the same Markov structure as the original probability measure. There is, at most, one such factorization for which this change of measure implies a Markov process that is stochastically stable, by which I mean that the conditional expectations converge in the limit to well-defined unconditional expectations. This stability property is what allows us to use the factorization to provide meaningful long-term characterizations of pricing.

I remark briefly on the long-term pricing implications of factorization (Equation 3) for limiting uncertainty compensation. Let us write

$$\lim_{\tau \rightarrow \infty} \rho_{t,t+\tau}(G).$$

Let η_g be the growth rate of the cash flow, η_s the decay rate of the stochastic discount factor, and η_{gs} the decay rate of the product GS . Then we obtain

$$\lim_{\tau \rightarrow \infty} \rho_{t,t+\tau}(G) = \eta_g + \eta_s - \eta_{gs}.$$

In the special case in which the stochastic discount factor process does not have a martingale component, this limiting compensation is zero. More generally, the probability associated with the martingale component of S is the long-term forward measure, as it absorbs the long-term compensations for exposure to uncertainty. This is also the measure that Ross's (2015) approach would recover in this setting.⁶

In this essay, I am primarily interested in the martingale component to a stochastic discount factor process. There are three ways to justify such a component:

- permanent shocks to the macroeconomy;
- recursive utility preferences that feature forward-looking components; and
- subjective beliefs, ambiguity-averse preferences, or concerns about model misspecification.

I will consider two example economies in which all three components are interconnected and reinforcing. But before doing so, let me elaborate on each and describe why they are related. Speculation about potential macroeconomic growth provides an interesting channel for uncertainty. Macroeconomists debate about the prospects of secular stagnation, and economic historians make different conjectures about the future prospects for technological advances. What are the long-term consequences of climate change on economic opportunities? Will potential future pandemics alter permanently how economies provide goods and services? As emphasized by Bansal & Yaron (2004), the recursive utility models of Kreps & Porteus (1978), Epstein & Zin (1989), and others contribute a forward-looking component even to a one-period stochastic discount factor $\frac{S_{t+1}}{S_t}$.

⁶Readers are referred to Borovička et al. (2016) for further discussion.

Although these are risk-based models, the same implication is often true in models of ambiguity aversion and model misspecification fears. As a decision maker grapples with uncertainty about the model specification, long-term uncertainty can emerge as having the biggest adverse impact on the decision maker's discounted utility. Thus, once again, concerns about the unknown nature of the long-term uncertainty can affect even the valuation of short-term assets such as one-period returns. The pricing outcomes of a broadly conceived uncertainty aversion are conveniently represented by an altered probability specification capturing the uncertainty that affects decision making in the most adverse way.

Here I propose one possible application of the factorization in Equation 3. I also feature the interaction between all three rationales for the martingale component to valuation and illustrate its quantitative importance.

4. UNCERTAINTY AND FINANCIAL MARKETS

[A] practical theory of the future...being based on so flimsy a foundation...is subject to sudden and violent changes.

—John Maynard Keynes (1937, p. 214)

In this section, I describe an example economy in which investors confront uncertainty about the future of macroeconomic growth. I base this example on a paper by Hansen & Sargent (2021).⁷ This example captures in a stylized way both the worrisome secular stagnation feared by some macroeconomists and the uncertain prospects for technological advances debated by economic historians. The model has a long-run risk component, like the one featured by Bansal & Yaron (2004), entering into the evolution of capital. We include production to provide substantively interesting interpretations of shocks. In the extensive literature that builds on or applies Bansal & Yaron's risk-based formulation of uncertain growth, the risk aversion that is imposed on investors is typically very large. In the example economy, I show how a broader notion of uncertainty aversion, including model ambiguity and misspecification concerns, provides a different perspective on how uncertain growth affects asset valuation. Whereas the long-run risk literature typically introduces stochastic volatility, specified exogenously, I abstract from that in the model I describe. I adopt this simplification to feature what fluctuations in asset valuation can be induced by investors' struggles with pinpointing the precise nature of the uncertainty.

As I will show, the prices of uncertainty fluctuate because investors especially fear high persistence of macroeconomic growth in bad states and low persistence in good ones. Why this asymmetry? When an economy is growing, the salient fear preoccupying investors is that high growth will not persist. But when the economy is stagnant, the salient fear is that sluggishness will persist. We show how these concerns manifest themselves in relation to both model ambiguity and model misspecification. Formally, they take the form of an endogenous nonlinearity in the stochastic discount factor that investors use to evaluate the prospective payout streams that underlie asset evaluations. This nonlinearity compounds itself over time in ways that provide a novel explanation for how asset market fluctuations reflect uncertainty.

4.1. Setup

Suppose that capital, K_t , evolves according to the continuous-time stochastic differential equation as

$$dK_t = K_t \left(\alpha_k + \beta_k Z_t + \frac{I_t}{K_t} \right) dt + K_t \sigma_k \cdot dW_t, \quad 4.$$

⁷Hansen & Sargent (2019) offer a short perspective on this research.

and output is proportional to the capital stock,

$$C_t + K_t \phi \left(\frac{I_t}{K_t} \right) = aK_t,$$

where C_t is consumption, I_t is investment, ϕ is a convex function introduced to capture adjustment costs, and dW_t is a bivariate increment to a Brownian motion. I interpret capital broadly to include potentially both human and organization capital.

For simplicity, I assume that investors have a unitary elasticity of intertemporal substitution in preferences, which implies that, for an optimal resource allocation, consumption and capital should be proportional. The process Z captures exogenous shifts in technology and is modeled as a scalar autoregression,

$$dZ_t = \alpha_z - \beta_z Z_t + \sigma_z \cdot dW_t.$$

While the Z process captures a form of long-run risk, I have purposefully abstracted from including a stochastic volatility component, even though such a component is common in macrofinance models. Although there is evidence for such volatility in the macroeconomy, my aim is to show how concerns about uncertainty can induce fluctuations in the market price of uncertainty endogenously, without imposing it through an external source.

There are two parameters that will be of particular interest to us. The first is $\beta_k > 0$, which determines how responsive the capital stock is to changes in the persistent technology process. Larger values of β_k imply more exposure to technological uncertainty. The second is $\beta_z > 0$, which determines the persistence of the process Z . Larger values of β_z imply more mean reversion and more pull toward the long-run mean. The counterpart to the discrete-time autoregression coefficient is $\exp(-\beta_z)$, which is closer to zero when β_z is large.

I construct the rational expectations competitive solution by following Lucas & Prescott's (1971) approach, in which the equilibrium resource allocation is the outcome of a fictitious planner's problem for a stochastic growth model à la Brock & Mirman (1972). I then compute the three additional solutions:

- planner's problem with a concern for the misspecification of a baseline model used in the rational expectations solution;
- planner's problem with a concern for ambiguity across models, as indexed by the parameter pair (β_k, β_z) ; and
- planner's problem with a concern for both model misspecification and ambiguity across models.

The solution to the first problem illustrates the impact of robust preferences (as in Hansen & Sargent 2001), the second illustrates the impact of ambiguity averse preferences (as in Chen & Epstein 2002), and the third illustrates the impact of combining the two (as in Hansen & Sargent 2021).⁸

I treat the case with a unitary elasticity of substitution for pedagogical simplicity. Capital is the sole source of wealth in this economy, and with this preference restriction, consumption is proportional to capital. Similarly, there is a constant investment–capital ratio. Although changing how investors confront uncertainty alters the equilibrium stochastic discount factor, the equilibrium

⁸The robust preferences used by Hansen & Sargent (2001) and the generalization by Hansen & Sargent (2021) can be viewed as special cases of the dynamic variational preferences of Maccheroni et al. (2006).

allocations for consumption and capital stay the same. I still solve a planner's problem because the minimizing probabilities from the planner's problem, expressed as a function of the state variables and Brownian shocks, contribute an uncertainty adjustment to the stochastic discount factor used to represent asset prices in a competitive equilibrium.⁹ Relative to the rational expectations model, each of the three solutions induces a different martingale contribution. The martingale contributions capture the uncertainty-adjusted probabilities deduced from the planner's problem. Because I posed the model in continuous time, the value functions and worst-case probabilities solve Hamilton–Jacobi–Bellman equations, as in Hansen & Sargent's (2021) article.

Had I not imposed a unitary elasticity of substitution, or had I introduced multiple capital stocks with differential exposure to uncertainty, the quantity allocation would no longer be invariant across the three alternative preference specifications. Instead, changing the preferences as described in the previous paragraph would have altered not only the stochastic discount factor but also the equilibrium allocation of consumption and capital.¹⁰ Such generalizations provide some additional insights into how uncertainty aversion alters precautionary motives for savings as reflected in the evolution of capital. In particular, the ratio of consumption to aggregate capital would be state dependent. Of course, such outcomes are of considerable interest, but here I have chosen to focus on the implications for the equilibrium stochastic discount factor.

4.2. Modeling Details

Given the equivalence of the capital and consumption dynamics, I construct the baseline parameter values from the estimates of $(\hat{\alpha}_k, \hat{\beta}_k, \hat{\sigma}_k, \hat{\alpha}_z, \hat{\beta}_z, \hat{\sigma}_z)$, as done by Hansen & Sargent (2021). They set $\hat{\alpha}_z = 0$, which means that the Z process has mean zero in the implied stationary distribution, and the sum of $\hat{\alpha}_k$ and the investment–capital ratio gives the long-term growth rate in logarithms for consumption and capital. They also set $\hat{\beta}_k = 1$. In the rational expectations equilibrium, the investor has full confidence in the baseline parameter values.

To explore model misspecification, the investor considers alternative probability specifications, each of which can be represented as a Brownian motion with drift. We use U to represent the stochastic process of drifts. Then we obtain

$$d\widehat{W}_t = U_t dt + dW_t^U,$$

where dW_t^U is a Brownian motion under the change of probability measure. Thus, locally I am changing the distribution of the underlying shocks from being multivariate standard normal shocks to being shocks that have a mean that is not zero and instead can be time dependent and state dependent. I use flexible specifications of the drift process U to capture alternative forms of misspecification.

The martingale associated with this drift distortion has a particularly simple dynamic evolution under the baseline $\widehat{\cdot}$ dynamics: We have

$$dM_t^U = M_t^U U_t \cdot d\widehat{W}_t,$$

⁹While it is common to compute the competitive equilibrium allocation by solving a social planner's problem, I use an extended version of that same approach to compute the uncertainty adjustment for the equilibrium asset prices (see Hansen & Sargent 2021 for more elaboration on this point).

¹⁰Eberly & Wang (2012) provide a tractable extension of this model with two capital stocks and risk-based recursive utility.

where the (at least local) martingale property follows, given that the conditional mean or drift is zero. As M^U is a likelihood ratio, the evolution of the logarithm is also of particular interest:

$$\begin{aligned} d \log M_t^U &= -\frac{|U_t|^2}{2} dt + U_t \cdot d\widehat{W}_t \\ &= \frac{|U_t|^2}{2} dt + U_t \cdot dW_t^U. \end{aligned}$$

Thus, the local mean of expected log-likelihood ratio under the M^U probability distribution is one half the squared norm of the drift distortion U . This is the local contribution to relative entropy, which is a well-known discrepancy and a mathematically convenient measure for assessing differences in probability measures. Relative entropy is also often referred to as Kullback–Leibler divergence. This local contribution of relative entropy exploits the underlying Brownian formulation with locally normal shocks. When the squared norm of the drift is large, the implied probability measure is easy to distinguish from the baseline probability.

Motivated by results from the robust control theory, Hansen & Sargent (2001) include penalty, which is a scaled version of

$$\frac{1}{2} \mathbb{E} \left[\int_0^\infty \exp(-\delta t) \left(\frac{M_t^U}{M_0^U} \right) |U_t|^2 dt \mid \mathfrak{F}_0 \right]$$

in the objective function of the fictitious planner.¹¹ This objective includes the same subjective discount rate, δ , as the one used in discounting the expected utility of the planner. I impose the restriction so that the preferences are dynamically consistent. A larger penalty limits more the decision maker's exploration of the potential consequences of misspecification. Including this term in the planner's objective leads to a particularly simple choice of U_t^* for this example economy, that is,

$$U_t^* \propto - \left(\frac{\widehat{\beta}_k}{\widehat{\beta}_z + \delta} \right) \widehat{\sigma}_z - \widehat{\sigma}_k.$$

The right-hand side of this expression gives the negative of how the increment in the continuation value (expressed in logarithms) responds to the Brownian increment. Notice that both slope parameters contribute to this drift distortion. More exposure to the exogenous growth rate uncertainty, captured by a larger value of $\widehat{\beta}_k$, magnifies the contribution of $\widehat{\sigma}_z$. More persistence in this process captured by a smaller value of $\widehat{\beta}_z$ has the same qualitative impact. The proportionality factor depends on the reciprocal of the scaling factor for the relative entropy penalty. In other words, a large penalty results in a small implied drift distortion. Notice that while the drift distortion can be state dependent, in the solution to the planner's problem, this distortion is constant. The decision maker repeatedly worries about permanent shifts in drift of the Brownian increment. The fact that $-U_t^*$ is constant provides a mechanism for enhancing uncertainty prices but not for causing them to vary. Although this invariance is driven in part by the functional forms I have assumed, more generally I expect the implied state dependence in the uncertainty prices to be small.

I now consider ambiguity aversion relative to what Hansen & Sargent (2021) call a structured set of models. This will also illustrate a version of Chen & Epstein's (2002) formulation of ambiguity aversion. In this computation, to capture parameter ambiguity, I introduce a special case of

¹¹For some related robust control references, readers are referred to Jacobson (1973), James (1992), and Petersen et al. (2000).

the drift distortion,

$$S_t = \eta_0 + \eta_1 Z_t,$$

where η_0 and η_1 are two-dimensional vectors. Now rewrite the state vector dynamics as

$$\begin{aligned} dK_t &= K_t \left(\hat{\alpha}_k + \hat{\beta}_k Z_t + \frac{I_t}{K_t} \right) dt + K_t \hat{\sigma}_k \cdot d\widehat{W}_t \\ &= K_t \left(\alpha_k + \beta_k Z_t + \frac{I_t}{K_t} \right) dt + K_t \hat{\sigma}_k \cdot dW_t^S, \\ dZ_t &= \hat{\alpha}_z - \hat{\beta}_z Z_t dt + \hat{\sigma}_z \cdot d\widehat{W}_t, \text{ and} \\ dZ_t &= \alpha_z - \beta_z Z_t dt + \hat{\sigma}_z \cdot dW_t^S, \end{aligned}$$

where

$$\begin{aligned} \alpha_k &= \hat{\alpha}_k + \hat{\sigma}_k \cdot \eta_0 & \beta_k &= \hat{\beta}_k + \hat{\sigma}_k \cdot \eta_1, \\ \alpha_z &= \hat{\alpha}_z + \hat{\sigma}_z \cdot \eta_0 & \beta_z &= \hat{\beta}_z + \hat{\sigma}_z \cdot \eta_1. \end{aligned}$$

By varying η_0 and η_1 , I obtain alternative parameters of the mean dynamics of the state variables.

Given the role of $\hat{\beta}_k$ and $\hat{\beta}_z$ in the previously constructed U^* , I will feature ambiguity about these two parameters. Thus, I hold fixed $(\alpha_k, \sigma_k, \alpha_z, \sigma_z) = (\hat{\alpha}_k, \hat{\sigma}_k, 0, \hat{\sigma}_z)$ and limit the parameter uncertainty to (β_k, β_z) . This leads me to focus on η_1 and set $\eta_0 = 0$.¹² As Hansen & Sargent (2020) argue, I cannot embed misspecification concerns of the type described previously within this framework, because Chen & Epstein (2002) impose an instant-by-instant constraint. They restrict Brownian motion drifts in this manner to implement a version of Gilboa & Schmeidler's (1989) ambiguity-averse preferences, extended to be dynamically consistent. To implement this approach, we continue to use relative entropy but now in a much more limited way.

Consider for the moment a single alternative model given by a choice of η_1 . I could compute the implied discounted relative entropy ϵ , which will be dependent on the realized state z . My aim, however, is to build a constraint on structured models. Thus, I start with a function ϵ and find values of η_1 with the same relative entropy. This defines the boundary of the set of structured models and is characterized by a quadratic equation in the two-dimensional vector η_1 , or equivalently for (β_k, β_z) . I also include in this ambiguity set the interior points, each of which is a drift distortion associated with a value of η_1 . Given this constraint, the planner will justifiably focus on the boundary when minimizing over S . As Hansen & Sargent (2021) make clear, this constraint on S is much tighter than a date-zero constraint restricting relative entropy to be less than or equal to $\epsilon(Z_0)$.

The fictitious planner's problem is actually a two-player max-min dynamic game that can be expressed and solved using recursive methods. The maximization is over investment, and minimization, conditioned on the investment choice, is over the structured drift distortion. For the recursive minimization, I start with a functional equation for relative entropy for each η_1 on the boundary.¹³ I convert this equation into a constraint on η_1 for each possible value of z by taking as input the relative entropy ϵ as a function of z . For my application, this minimization problem is quadratic given the state z .¹⁴

I depict the resulting constraints on the parameters in **Figure 1**. The decision maker is more ambiguity averse when the ambiguity set is larger. To be in line with Chen & Epstein (2002), I

¹²Hansen & Sargent (2021) also present results when η_0 is different from zero.

¹³This equation is a special case of the so-called Feynman–Kac equation.

¹⁴Readers are referred to Hansen & Sargent (2021) for more details.




Figure 1

Ambiguity sets. These are parameter values constrained by relative entropy, where β_k quantifies the exposure to the macroeconomic growth rate process and β_z quantifies the persistence of that process. The single point is the baseline specification $(\hat{\beta}_k, \hat{\beta}_z)$, and the two regions are implied by a relatively tight (*blue*) and a relatively loose (*orange*) relative entropy constraint. Figure adapted from Hansen & Sargent (2021) (CC-BY-NC-ND).

allow for the minimizing choice of (β_k, β_z) to depend on the alternative realized values z of Z_t . This is, indeed, an outcome of our model solution. Following Hansen & Sargent (2021), I motivate this extra flexibility by allowing a priori the parameters to be time varying but constrained to reside in the ambiguity parameter set. Alternatively, I could allow for a restricted set of nonlinear specifications of the local means for the state vector evolution.

Investors in this economy start with a simple linear first-order autoregressive model of macroeconomic growth rate dynamics. Their concerns about ambiguity induce them to explore other statistically similar specifications, including ones that they especially fear. In **Figure 2**, the linear relation with the negative slope captures their baseline model with so-called mean reversion for the growth rate process Z . The mean reversion is evident because there is a pull from the more extreme growth states toward the center of the growth rate distribution. The vertical axis is the local pull toward the center of the distribution of macroeconomic growth (net of its long-run average growth rate.) In the absence of random shocks, there is a pull toward zero. In **Figure 2a** I also include a dotted red curve that is the ambiguity-averse response when investors confront a set of statistically similar specifications. Formally, they explore the ambiguity set of parameters for different macroeconomic growth states. The flatter slope (a smaller value of β_z) to the left of zero reflects the concerns of investors in bad economic times that the macroeconomy may be stuck with more growth sluggishness than in the original model. The steeper slope to the right of zero reflects opposite forces. Here, good macroeconomic growth outcomes are feared to be shorter-lived than in the original model specification. Finally, in **Figure 2a** I plot the implied drift when




Figure 2

Local dynamics for macroeconomic growth. Solid gray lines represent the baseline linear model. The gray arrows depict the pull of the local dynamics toward the mean growth rate of the macro economy for the baseline model. In panel *a*, the dotted red curve incorporates aversion to model ambiguity, and the dashed blue line incorporates aversion to model misspecification. In panel *b*, the dashed magenta curve incorporates aversions to both model ambiguity and model misspecification. Figure adapted from Hansen & Sargent (2021) (CC-BY-NC-ND).

the only investor aversion is that the baseline model is misspecified, which just shifts the baseline drift proportionally downward. There is a corresponding movement in the minimizing choice of β_k , which governs how exposed the economy is to the uncertain growth process.

Neither the dotted red curve nor the dashed blue line in **Figure 2a** is intended to represent the beliefs of the planner or the investors inside our example economy. The computation of this altered drift is part of deducing the averse response to ambiguity over the state dynamics. The imputed drift distortion S_{thinsp}^* from the planner's problem that implements the competitive equilibrium is a nonlinear function of the growth state Z . Interestingly, it is a continuous-time version of a threshold autoregression model. The martingale associated with this S_{thinsp}^* , with a drift given by

$$dM_t^{S^*} = S_t^* \cdot d\widehat{W}_t,$$

contributes prominently to the martingale component of the stochastic discount factor.

In **Figure 2b**, I look at a specification featured by Hansen & Sargent (2021) that allows each model in the ambiguity set to be misspecified. Here, the structured specifications of the drift distortion S are in the same ambiguity set described previously, but I also include a relative entropy penalty that is proportional to

$$\frac{1}{2} \mathbb{E} \left[\int_0^\infty \exp(-\delta t) \left(\frac{M_t^U}{M_0^U} \right) |U_t - S_t|^2 dt \mid \mathfrak{F}_0 \right].$$

The unstructured misspecification is captured by the drift distortion U , which is penalized relative to the set of structured models. **Figure 2b** shows the resulting minimizing unstructured drift distortion given by the dashed magenta curve. The altered drift inherits a similar shape to that of ambiguity-adjusted drift, but it is shifted downward. This downward shift is no longer proportional, as it was when investors only considered a single specification.

Empirical evidence suggests that uncertainty compensations for the macroeconomy are larger in bad times relative to good times. The local uncertainty compensation induced by an aversion to




Figure 3

Term structure of uncertainty prices. The gray curve sets the initial growth state at the median of the stationary distribution; the dashed red curve sets the initial state at the 0.9 decile; the dotted blue line sets the initial state at the 0.1 decile. Panels *a* and *b* give the term structure of prices for exposures to the two underlying shocks. Figure adapted from Hansen & Sargent (2021) (CC-BY-NC-ND).

ambiguity or model misspecification is governed by U^* . To explore the pricing effects, we consider a term structure of uncertainty prices captured by the formula

$$\mathbb{E}(M_t^{U^*} U_t^* | Z_0). \quad 5.$$

I think of prices as building blocks of valuation adjustments due to uncertainty. Notice that $-\sigma_k \cdot U^*$, as depicted in **Figure 2b**, is large when the magnitude of the growth state is large. This is true independent of whether growth is relatively high or relatively low. An important asymmetry emerges, however, once we look at the horizon dependence reflected in Equation 5. The steepersloping conditional mean for Z when z is positive means that under the M^{U^*} probability, there is a relatively fast escape from these states vis-à-vis the baseline probability specification. Conversely, there is a slower escape for z negative. This asymmetry is reflected in **Figure 3**, which reports the term structure of uncertainty prices given by Equation 5. These uncertainty prices depend on the current growth state. We consider three values of this state: One is median ($z = 0$), and the others are the 0.1 and 0.9 deciles of the Z stationary distribution under the base probability measure. Whereas both the upper and lower deciles start substantially higher than the median uncertainty prices, the 0.9 decile prices show substantial decay because of the likely escape from the high-growth state. **Figures 3a** and **3b** give the uncertainty prices for the two underlying shocks. Notice that the patterns are very similar. In summary, the term structure reveals an intriguing asymmetry in the uncertainty prices depending on macroeconomic growth.

4.3. Related Literature

In this example economy, I abstracted from learning by including alternative models with parameters whose prospective variations cannot be inferred from historical data. These time-varying

parameter models differ from the decision maker's baseline model, a fixed-parameter model whose parameters can be well estimated from historical data. Other research imposes parameter invariance and instead induces fluctuations using a robust Bayesian perspective. For instance, in work by Hansen (2007) and Hansen & Sargent (2010), a representative investor's robust model averaging induces fluctuations in uncertainty prices. The investor carries along difficult-to-distinguish models of macro growth. The investors make robust adjustments to Bayesian posteriors updated with historical data. This leads the investors to act as if good news is temporary and bad news is persistent, an outcome that is qualitatively similar to what we have found here. The decision makers in their analyses confront dynamic consistency by playing a game against future versions of themselves, although I speculate that similar qualitative insights could be obtained using the continuous-time recursive formulation by Hansen & Miao (2018). Collin-Dufresne et al. (2016) and Andrei et al. (2019) show that a comparable outcome emerges with a risk-based, recursive utility formulation with Bayesian learning in which investors have full confidence in the subjective probabilities and are not concerned about model misspecification. Although it is valuable to see this rational learning counterpart, the risk-based approaches miss part of what I consider to be an important challenge investors face as they speculate about the future.

5. UNCERTAINTY AND POLICY

Even if true scientists should recognize the limits of studying human behavior, as long as the public has expectations, there will be people who pretend or believe that they can do more to meet popular demand than what is really in their power.

—Friedrich A. Hayek (1974)

There is often a tendency to overstate the knowledge base when providing public rationales for policies aimed at addressing social and economic problems. Premising decisions on false knowledge can lead to bad outcomes. Some might argue that unless we know something for sure (or at least with great confidence), we should adopt a wait-and-see approach. But this conclusion does not necessarily follow from the basic tenets of decision theory under uncertainty. Because there can be a substantial cost to the delay, decision theory can justify acting now based on the possibility of bad outcomes in the future should we fail to act. The aversion to uncertainty is meant to resolve a trade-off in assessing decisions between focusing on best guesses and considering possible adverse outcomes even though they might be unlikely. My hope is that tools such as the ones described in this essay can allow policy makers to confront uncertainty in sanguine discussions of policy.

In this section, I explore the impact of the economic uncertainty of climate change by drawing on the analysis by Barnett et al. (2020). In their research, these authors assess a stylized integrated assessment model, incorporating both geoscientific and economic uncertainties. There exist many such integrated assessment models, but my particular focus is on the consequences of uncertainty for prudent policy making. Following Barnett and colleagues, I show that geoscientific and economic uncertainties are in effect multiplicative in quantifying the social cost of carbon (SCC) emissions. I accomplish this by using the tools from asset pricing to study social valuation in contrast to private valuation. A central part of Barnett and colleagues' (and hence my) analysis is the construction of probability measures that adjust the uncertainties pertinent to valuation. Specifically, I construct a martingale component to the stochastic discount factor based on a cautious approach to uncertainty.

5.1. Social Cost of Carbon

The SCC is commonly referred to in policy discussions, but its meaning and implications for measurement differ across applications. In this discussion, I use a well-posed version as an analytical tool to assess the impact of uncertainty. I alter the technology in the previous section to include an exploration sector that leads to increases in the stock of fossil fuel reserves. In addition, I modify preferences to include a demand for emissions. Emissions will alter the climate, which in turn will affect economic opportunities and social well-being in the future. I take the emissions' impact on climate and economic opportunities to be an externality not captured by market prices. I model the SCC as a stochastic process of the implied shadow prices for emissions that is pertinent for a fictitious social planner engaged in socially efficient resource allocation. The SCC in my computations depends on the underlying state variables in the economy, and thus it varies over time. The wedge between the market prices and the SCC is the process of so-called Pigouvian taxes on emissions that would correct the externality.

Pigouvian tax policy could be dismissed as a pie in the sky, and the use of a single social planner as remarkably naive, since it ignores a variety of practical challenges in policy making. I am sympathetic to such criticisms, and I will not address the many political and economic challenges of climate change policy. Instead, I use the SCC as a barometer for gauging the quantitative importance of uncertainty.

I bring in tools from asset pricing by viewing the SCC as a process of asset prices with a corresponding social cash flow. For simplicity, I use temperature as the measure of climate change. The social cash flows of interest have two interconnected contributions:

- Geoscience: How do changes in CO₂ emissions alter temperature in the future?
- Economics: How do changes in temperature alter economic opportunities in the future?

Both of these contributions can be conceived of as nonlinear, local impulse responses, the first being the future responses of temperature to an emissions impulse, and the second being the future responses of economic opportunities to a temperature impulse. In effect, I form a convolution of these two response functions to get the social cash flow. The SCC agglomerates the social cash flows using stochastic discounting, including a martingale adjustment for uncertainty.

The SCC is often defined as the social marginal rate of substitution between emissions and consumption; thus, it is a shadow price of the resource allocation problem for a hypothetical planner. It could be implemented via a Pigouvian tax that would correct the private shadow price for the externality, although we use this method to assess the impact of uncertainty, when conceived broadly. Following the previous literature, we start by representing this social cost in terms of partial derivatives of the value function of the social planner. We then apply an asset pricing perspective to interpret the components to this social cost.

5.2. Prior Literature

As many previous researchers have investigated, the human impact on the climate is a potentially important source of uncertainty that could play out over long horizons (see, for instance, Weitzman 2012, Jensen & Traeger 2014, Cai et al. 2015, Nordhaus 2017, Hambel et al. 2018, and, especially, Cai et al. 2017). Our use of an asset pricing perspective to interpret the SCC follows in part discussions by Golosov et al. (2014) and Cai et al. (2017), who engage in an ambitious exploration of the risk consequences for the SCC. My aim is to show how to extend these analyses to include forms of uncertainty other than risk. Although I use a particular model of the climate system for the sake of illustration, this same perspective also allows researchers to better understand

the components to the social cost in more general settings. Millner et al. (2013) and Lemoine & Traeger (2016) previously used the smooth ambiguity formulation of decision making motivated by climate science. In contrast to the analysis that follows, they did not formally motivate their calibration of ambiguity in terms of robustness and sensitivity considerations.

5.3. Setup

I use an AK model of the type I specified in Section 4 (see Equation 4), except that I modify the output constraint to be

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

where J_t is investment in the discovery of new fossil fuel reserves. The stock of reserves evolves according to

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

where R_t is the date- t stock of reserves, and E_t is the date- t flow of emissions of fossil fuels drawn from the reserve stock. The parameters satisfy the restrictions $\psi_0 > 0$ and $0 < \psi_1 \leq 1$. The well-known Hotelling's (1931) model of resource extraction with a fixed stock of reserves becomes a special case of our reserves by setting $\psi_0 = 0$ and $\sigma_R = 0$ and by restricting $R_t \geq 0$ for all $t \geq 0$.¹⁵

Capturing the externality induced by climate change requires the integration of a climate change component into the model. It is difficult to integrate a complex and high-dimensional climate model within a stylized economic model of the type that I use here. To get off the ground, I adopt some seemingly stark approximations. Climate scientists have performed a variety of model comparisons by using exogenously specified representative concentration pathways and by using emission pulses of different sizes. Examples can be found in the work of Eby et al. (2009), Joos et al. (2013), and MacDougall et al. (2017). The simplified dynamics are often specified in terms of two subsystems: (a) the dynamic response of atmospheric carbon concentration to emissions, and (b) the dynamic response of temperature to atmospheric carbon concentration. The response of temperature to emissions is in effect a convolution of these two responses. While it is common for integrated assessment models to formally specify these two subsystems, there is a substantial literature that studies direct simplifications of the convolution by showing that the temperature change over an interval of time is roughly proportional to the cumulative carbon emissions over the same time period. Matthews et al. (2009), among others, proposed this as a useful approximation. Many subsequent papers have assessed this approximation and have used it for model comparisons and policy discussion. For instance, Ricke & Caldeira (2014) have argued that, in fact, this proportionality emerges over a relatively short timescale. Dietz & Venmans (2019) and Barnett et al. (2020), in independent contributions, have used this approximation within an explicit economic optimization framework. The latter contribution proposes ways to confront uncertainty, broadly conceived, within such a framework. Barnett and collaborators use the cross-model uncertainty in the proportionality relationship as input in their analysis. **Figure 4a** shows the cross-model histogram reported in MacDougall et al.'s (2017) paper for the

¹⁵Our introduction of a technology for investing in the exploration of new reserves follows from previous contributions by Casassus et al. (2018) and Bornstein et al. (2019). While investment in new discoveries is typically absent in Dynamic Integrated Climate Economy (DICE) models, DICE models introduce possible mitigation to offset the economic damages induced by climate change.

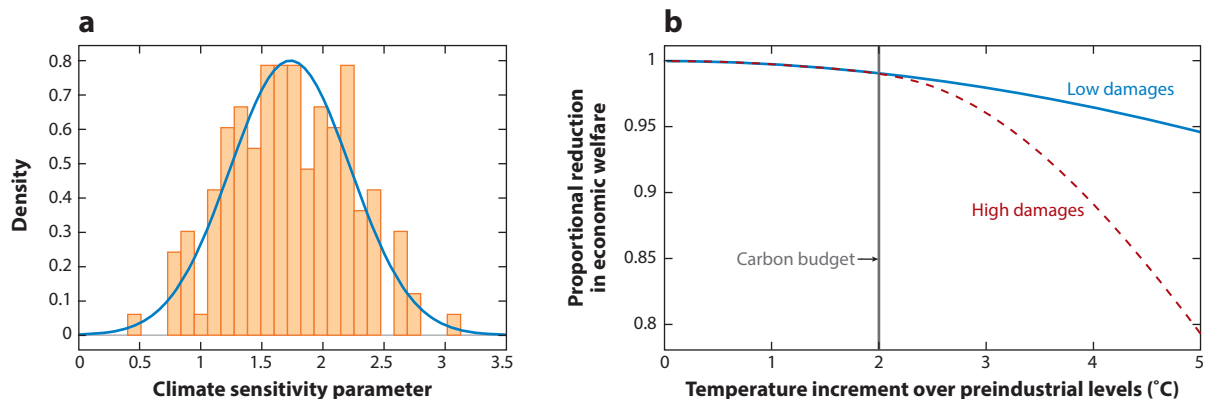


Figure 4

Uncertainty inputs. (a) Histograms and density for the climate sensitivity parameter across models. Evidence is from MacDougall et al. (2017). I use the blue curve as a baseline specification of probabilities across models instead of an orange histogram to simplify computation. (b) Two damage functions specified as proportional reductions in consumption due to climate change. Climate change is measured as the climate sensitivity parameter times cumulative emissions. Figure adapted from Barnett et al. (2020) (CC-BY-NC).

proportionality coefficient β . Taking this histogram as a direct characterization of uncertainty means that all of the inputs into the construction of the histogram are given equal weighting. Although this might be an interesting starting point, we shall examine sensitivity to this choice.

From the economic side, I specify a damage function meant to summarize the economic consequences of climate change. While there are good reasons to consider dynamic responses to climate change through alternative forms of adaptation, like many scholars, including Dietz & Venmans (2019) and Barnett et al. (2020), I posit a static relationship between temperature and damages, denoted by N_t . I follow Barnett and colleagues by letting N_t capture a proportional reduction in consumption due to climate change. Formally, they assume that

$$\log N_t = \Gamma \left(\beta \int_0^t E_u du \right) + \zeta_n(Z_t),$$

where

$$\Gamma(y) = \begin{cases} -\gamma_1 y - \frac{\gamma_2}{2} y^2 & y \leq 2 \\ -\gamma_1 y - \frac{\gamma_2}{2} y^2 - \frac{\gamma_2^+}{2} (y - 2)^2 & y > 2 \end{cases}.$$

The parameters γ_2 and γ_2^+ give nonlinear damage adjustments. Specifically, $\gamma_2^+ > 0$ gives a smooth alternative to a carbon budget at 2°C. Notice that the function Γ is zero at zero, decreasing, and concave.¹⁶ We include an exogenous shifter in this specification of damages via $\zeta_n(Z_t)$.

There is substantial divergence in the environmental economics literature on the magnitude of damages. To illustrate this divergence, **Figure 4b** depicts $\exp[\Gamma(y)]$ for two different possible γ_2^+ : One is zero and the other is positive. The construction of the Γ with $\gamma_2^+ = 0$ is motivated by the evidence reported by Nordhaus & Moffatt (2017), and the second more curved specification of Γ with $\gamma_2^+ > 0$ is loosely motivated by Weitzman's (2012) arguments for much more severe

¹⁶As Barnett et al. (2020), I posit Γ as a function of the climate change component to temperature fluctuations by omitting other exogenous components of temperature change that are not induced by economic activity.

damages beyond 2°C. Given this specification of damages, we see that γ_2^+ is multiplied by either β or β^2 , implying that the uncertainty across the economic and climate inputs is multiplicative. In the illustrative calculation, I impose a baseline subjective probability of equal weight on the two values of γ_2^+ and explore the consequences of making arguably small changes in these probabilities.

The fictitious planner in our analysis confronts ambiguity in the distribution of the climate sensitivity parameter β and the damage parameter γ_2^+ . The planner uses baseline subjective probabilities for both. We use Hansen & Miao's (2018) framework for a recursive implementation of the smooth ambiguity model in continuous time. The discrete-time version of smooth ambiguity was originally axiomatized by Klibanoff et al. (2005). Hansen & Sargent (2007) provide a robust Bayesian interpretation of an important special case of the discrete-time smooth ambiguity model, whereby the impacts of changes in the baseline posterior probabilities are explored recursively, subject to relative entropy penalty.¹⁷ Hansen & Miao's (2018) formulation preserves this robustness interpretation. Like Hansen & Sargent (2007), Hansen & Miao (2018) also show how to extend this decision theory representation to include concerns about model misspecification. In the climate economic computations that I report, I feature only the model ambiguity component. This is in contrast to the examples of the impact of uncertainty on financial markets that I described in Section 4. The robust Bayesian interpretation provides an altered probability as part of the planner solution for resource allocation, giving rise to a martingale component to a stochastic discount factor that is pertinent for social valuation.¹⁸

5.4. Implications of Ambiguity for the Social Cost of Carbon

Although stochastic discounting has been extensively applied in market valuation, it is also pertinent in social valuation. It justifiably redirects the discussion about what should be the discount rate for social valuation. Just as in market valuation, in social valuation there is a stochastic discount factor that adjusts valuation for exposures to uncertainty in ways that can depend on both the horizon and the current Markov state. Again, the martingale component is of particular interest. As in market valuation under uncertainty, there is a measure that adjusts for the aversions to ambiguity and model misspecification associated with this martingale component.

The SCC is an asset price associated with a particular social cash flow that reflects the adverse impact of climate change on economic and social outcomes.¹⁹ This social cash flow is the impulse response function of emissions today on damages in the future, where damages incorporate marginal utility adjustments from the different time periods. We then apply an asset pricing perspective to interpret the components to this social cost. This social cash flow depends on the interacting uncertainties about economic damages and climate change. Following Barnett et al.

¹⁷Whereas the example I report considers two potential damage functions, Li et al. (2016) consider a parameterized family of log-linear damage functions with a baseline exponential distribution over the unknown slope parameter. They follow Hansen & Sargent (2007) by making a robustness adjustment to this baseline distribution.

¹⁸Initial applications of smooth ambiguity models to the economics of climate change can be found in work by Millner et al. (2013) and Lemoine & Traeger (2016), but without reference to robustness considerations for the subjective probabilities.

¹⁹Such a discussion follows in part from Golosov et al.'s (2014) work. Cai et al. (2017) engage in a more ambitious exploration of the risk consequences for the SCC. Barnett et al. (2020) also embrace an asset pricing interpretation, but in particular, they show how to extend the analysis of the SCC to include aversions to ambiguity and model misspecification.

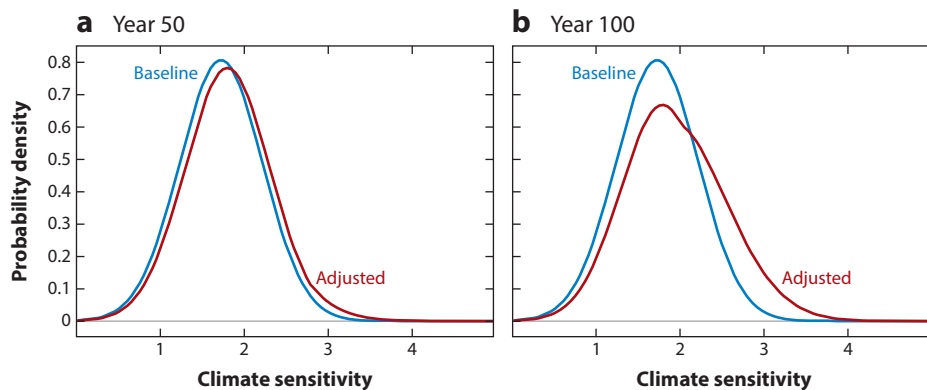


Figure 5

The initial weighting is the same for both the low- and high-damage specifications. The ambiguity-adjusted probabilities remain very close to this initial weighting at year 50 (panel *a*) but are tilted toward the high-damage specification (probability 0.59) at year 100 (panel *b*). Figure adapted from Barnett et al. (2020) (CC-BY-NC).

(2020), I include the ambiguity-adjusted probabilities from the robust planner's problem. This change of probability measure contributes a martingale component to the valuation.²⁰

In what follows, I report some calculations by Barnett et al. (2020). They abstract from learning about unknown parameters and assume that the baseline probabilities for the climate sensitivity parameter β , and those for the damage parameter γ_2^+ , are invariant over time. The damage uncertainty pertains to damages induced by temperature for which we do not have historical data. This severely limits the possibility of Bayesian learning. For the parameter β , learning is also particularly challenging. The cross-model uncertainty shown in **Figure 4** seems unlikely to be noticeably reduced in the near future. Although abstracting from learning is a reasonable starting point, it may well be that if we start to experience even more extreme temperatures and more severe damages, future evidence could become much more informative and helpful in resolving this uncertainty. Even though the baseline probabilities are time invariant for the calculations I report, the probability adjustments are time varying. This follows from the recursive formulation of the decision problem. Specifically, the consequences of changing baseline probabilities become more pronounced as damages increase with warmer temperatures, as shown in **Figure 5**. This is evident from the adjusted densities reported at year 50 versus year 100. By year 100, the high-damage specification $\gamma_2^+ > 0$ is assigned probability 0.59, in contrast to 0.5 for the baseline probability and for the year 50 probability.

Although the probability adjustments displayed in **Figure 5** may seem modest, they have a quantitatively important impact on the SCC, as depicted in **Figure 6**. This figure reports the SCC using stochastically discounted social cash flows, with expectations computed using two distinct probabilities: uncertainty-adjusted probabilities and baseline probabilities. The former is the correct calculation, and it includes a martingale component needed to adjust for ambiguity about how to subjectively weigh the alternative climate model outcomes and how to weigh the two alternative specifications of damages. The increases in the SCC for both trajectories, while

²⁰While I use mathematical methods that are familiar from asset pricing theory, the implied shadow prices are computed using a socially efficient allocation and not necessarily an observed allocation determined by suboptimal policies and competitive markets.




Figure 6

Social cost of carbon (SCC) computed under two probability measures. The solid blue curve gives the SCC for the ambiguity-averse planner. The orange curve gives the counterpart computation with the same social cash flow but using the benchmark probabilities without the martingale adjustment. Figure adapted from Barnett et al. (2020) (CC-BY-NC).

reminiscent of Hotelling's (1931) model, are induced by the potential impact of damages to economic opportunities in the future and not to a finite stock of fossil fuel resources.

The model used in this section is purposely stylized to show how uncertainty can affect the SCC. An enhanced version of the model, while computationally more demanding, could address the uncertainty consequences for potential future policy and private sector responses to climate change. These could include the transition from fossil fuel to renewable energy or green technologies, an increase of the sink capacity through biodiversity conservation, and endogenous economic responses to climate change, including adaptation. These extensions will open additional channels through uncertainty that will affect social and economic outcomes, making the tools discussed in this essay all the more relevant.

6. CONCLUSION

For a variety of economic phenomena, researchers, policy makers, and economic agents are hard-pressed to come up with a single model to use for analysis or for the design of a prudent policy. The tools that I described preserve tractability while allowing for decision making when there is ambiguity over models and each model has the potential to be misspecified. I used the term “quantitative storytelling” metaphorically to capture the process of using a model with empirically motivated inputs to make quantitative predictions without carrying the burden of being correctly specified. This review addresses the need to balance off the implications of multiple stories in decision making.

I presented two examples of quantitative storytelling with multiple stories to show how this uncertainty affects the analysis and the conclusions. Empirical asset pricing researchers will have little difficulty finding asset pricing anomalies that the first example fails to address. But my point in developing this example is to expound a novel mechanism without the external clutter needed to address a broader set of empirical evidence from financial markets. While the aim of my second example is to quantify the SCC, other scholars have questioned the value of using integrated assessment models for this purpose. For instance, Pindyck (2013) and Morgan et al. (2017) argue

that SCC calculations from integrated assessment models are of limited value for setting policy because of the modeling simplifications and the limited understanding of climate change and its potential for damaging the economy. The methods I describe are meant to allow us to incorporate a better characterization of the uncertainties and knowledge limits we encounter. With the simplified modeling inputs that I used in the illustration, some important policy-relevant questions are off the table. Nevertheless, the methods that I illustrate open the door to a variety of important extensions that will enlarge the scope of the analysis while continuing to acknowledge the existing uncertainties.

I offer up the methods described here to reveal the sensitivity of quantitative conclusions to model ambiguity and potential misspecification. The penalty parameters needed as inputs tie directly to (decision makers') uncertainty aversions. A researcher using these methods need not choose a single configuration of these parameters. Here, we can start with a broad array of model specifications and subjective probabilities. The decision problem, in conjunction with credible subjective inputs and empirical evidence, reveals where the subjective inputs have the biggest impact on the decision making of economic agents or policy makers. The quantitative implications depend on penalization parameters, but the approach ensures that dependence will be of low dimension. Inspired by Good (1952), I apply robust Bayesian methods, modified to be applicable to the study of economic dynamics. A methodological achievement of the research I describe is providing a framework for showing how parsimonious penalized probabilities reveal sensible choices on how to restrain the search over subjective inputs.

The two examples that I explored align the uncertainty perspectives of private agents in the model with those of a fictitious planner. In the first example, this alignment is used to simplify the computation of the competitive equilibrium prices. In presenting the second application, I suggest, but do not formally derive, a competitive equilibrium counterpart subject to Pigouvian taxes deduced from a planner's problem. More generally, however, the planner and the private agents may have different perspectives on what is uncertain. It is common in macroeconomics to study the design and implementation of policies by formulating Ramsey-type problems as dynamic games played between a government and private agents. Hansen & Sargent (2012) formulate and compare such problems with three different types of ambiguity. In two of the three formulations, the heterogeneity between a government and private agents in terms of the ambiguities they confront is reflected in the resulting equilibria. Karantounias (2013) uses one of these formulations to study how a government that is fully confident in the probability model of its expenditures conducts fiscal policy when facing a skeptical public. In a dynamic, hidden-action contracting problem, Miao & Rivera (2016) posit a principal who has doubts about probabilities when engaging a better-informed and more confident agent. Importantly, these examples open the door to the study of the impact that heterogeneous perspectives on ambiguity among agent types in an economy have on contract and policy design.

DISCLOSURE STATEMENT

The author is not aware of any affiliations, memberships, funding, or financial holdings that might be perceived as affecting the objectivity of this review.

ACKNOWLEDGMENTS

I thank William "Buz" Brock, Doug Diamond, Yiran Fan, Jian Li, Marco Loseto, Lubos Pastor, Diana Petrova, Tom Sargent, Jose Scheinkman, and Grace Tsiang for helpful comments. Financial support for this project was provided by the Alfred P. Sloan Foundation (grant G-2018-11113).

LITERATURE CITED

- Alvarez F, Jermann UJ. 2005. Using asset prices to measure the persistence of the marginal utility of wealth. *Econometrica* 73:1977–2016
- Andrei D, Hasler M, Jeanneret A. 2019. Asset pricing with persistence risk. *Rev. Econ. Stud.* 32:2809–49
- Bansal R, Dittmar R, Kiku D. 2009. Cointegration and consumption risks in asset returns. *Rev. Financ. Stud.* 22:1343–75
- Bansal R, Yaron A. 2004. Risks for the long run: a potential resolution of asset pricing puzzles. *J. Finance* 59:1481–509
- Barnett M, Brock WA, Hansen LP. 2020. Pricing uncertainty induced by climate change. *Rev. Financ. Stud.* 33:1024–66
- Bornstein G, Krusell P, Rebelo S. 2019. *Lags, costs, and shocks: an equilibrium model of the oil industry*. NBER Work. Pap. 23423
- Borovička J, Hansen LP, Scheinkman JA. 2016. Misspecified recovery. *J. Finance* 71:2493–544
- Brock W, Mirman L. 1972. Optimal economic growth and uncertainty: the discounted case. *J. Econ. Theory* 4:479–513
- Cai Y, Judd KL, Lenton TM, Lontzek TS, Narita D. 2015. Environmental tipping points significantly affect the cost-benefit assessment of climate policies. *PNAS* 112:4606–11
- Cai Y, Judd KL, Lontzek TS. 2017. *The social cost of carbon with climate risk*. Tech. Rep., Hoover Inst., Stanford, CA
- Campbell JY, Vuolteenaho T. 2004. Bad beta, good beta. *Am. Econ. Rev.* 94:1249–75
- Casassus J, Collin-Dufresne P, Routledge BR. 2018. Equilibrium commodity prices with irreversible investment and non-linear technologies. *J. Bank. Finance* 95:128–47
- Chen Z, Epstein L. 2002. Ambiguity, risk, and asset returns in continuous time. *Econometrica* 70:1403–43
- Collin-Dufresne P, Johannes M, Lochstoer LA. 2016. Parameter learning in general equilibrium: the asset pricing implications. *Am. Econ. Rev.* 106:664–98
- de Finetti B. 1937. La prévision: ses lois logiques, ses sources subjectives. *Ann. Inst. Henri Poincaré* 7(1):1–68
- Dietz S, Venmans F. 2019. Cumulative carbon emissions and economic policy: in search of general principles. *J. Environ. Econ. Manag.* 96:108–29
- Eberly JC, Wang N. 2012. *Reallocating and pricing illiquid capital: two productive trees*. Tech. Rep., Columbia Univ., New York
- Eby M, Zickfeld K, Montenegro A, Archer D, Meissner KJ, Weaver AJ. 2009. Lifetime of anthropogenic climate change: millennial time scales of potential CO₂ and surface temperature perturbations. *J. Climate* 22:2501–11
- Epstein LG, Zin SE. 1989. Substitution, risk aversion and the temporal behavior of consumption and asset returns: a theoretical framework. *Econometrica* 57:937–69
- Fisher I. 1930. *Theory of Interest*. New York: Macmillan
- Gilboa I, Schmeidler D. 1989. Maxmin expected utility with non-unique prior. *J. Math. Econ.* 18:141–53
- Golosov M, Hassler J, Krusell P, Tsyvinski A. 2014. Optimal taxes on fossil fuel in general equilibrium. *Econometrica* 82:41–88
- Good IJ. 1952. Rational decisions. *J. R. Stat. Soc. Ser. B* 14:107–14
- Hambel C, Kraft H, Schwartz E. 2018. *Optimal carbon abatement in a stochastic equilibrium model with climate change*. NBER Work. Pap. 21044
- Hansen LP. 2007. Beliefs, doubts and learning: valuing macroeconomic risk. *Am. Econ. Rev.* 97:1–30
- Hansen LP. 2012. Modeling the long run: valuation in dynamic stochastic economies. *Econometrica* 80:911–67
- Hansen LP. 2014. Nobel Lecture: uncertainty outside and inside economic models. *J. Political Econ.* 122:945–87
- Hansen LP, Heaton JC, Li N. 2008. Consumption strikes back? Measuring long-run risk. *J. Political Econ.* 116:260–302
- Hansen LP, Miao J. 2018. Aversion to ambiguity and model misspecification in dynamic stochastic environments. *PNAS* 115:9163–68
- Hansen LP, Sargent TJ. 2001. Robust control and model uncertainty. *Am. Econ. Rev.* 91:60–66
- Hansen LP, Sargent TJ. 2007. Recursive robust estimation and control without commitment. *J. Econ. Theory* 136:1–27

- Hansen LP, Sargent TJ. 2010. Fragile beliefs and the price of uncertainty. *Quant. Econ.* 1:129–62
- Hansen LP, Sargent TJ. 2012. Three types of ambiguity. *J. Monet. Econ.* 59:422–45
- Hansen LP, Sargent TJ. 2019. Acknowledging and pricing macroeconomic uncertainties. *VoxEU*, July 22.
<https://voxeu.org/article/acknowledging-and-pricing-macroeconomic-uncertainties>
- Hansen LP, Sargent TJ. 2020. Structured ambiguity and model misspecification. *J. Econ. Theory*. In press.
<https://doi.org/10.1016/j.jet.2020.105165>
- Hansen LP, Sargent TJ. 2021. Macroeconomic uncertainty prices when beliefs are tenuous. *J. Econometr.* 223:222–50
- Hansen LP, Scheinkman JA. 2009. Long-term risk: an operator approach. *Econometrica* 77:177–234
- Hayek FA. 1974. *Prize lecture: the pretence of knowledge*. Lecture to the memory of Alfred Nobel, December 11.
<https://www.nobelprize.org/prizes/economic-sciences/1974/hayek/lecture/>
- Hotelling H. 1931. The economics of exhaustible resources. *J. Political Econ.* 39:137–75
- Jacobson DH. 1973. Optimal stochastic linear systems with exponential performance criteria and their relation to deterministic differential games. *IEEE Trans. Autom. Control* 18:1124–31
- James MR. 1992. Asymptotic analysis of nonlinear stochastic risk-sensitive control and differential games. *Math. Control Signals Syst.* 5:401–17
- Jensen S, Traeger C. 2014. Optimal climate change mitigation under long-term growth uncertainty: stochastic integrated assessment and analytic findings. *Eur. Econ. Rev.* 69:104–25
- Joos F, Roth R, Fuglestad JS, Peters GP, Enting IG, et al. 2013. Carbon dioxide and climate impulse response functions for the computation of greenhouse gas metrics: a multi-model analysis. *Atmos. Chem. Phys.* 13:2793–825
- Karantounias AG. 2013. Managing pessimistic expectations and fiscal policy. *Theor. Econ.* 8:193–231
- Keynes JM. 1937. The general theory of employment. *Q. J. Econ.* 51:209–23
- Klibanoff P, Marinacci M, Mukerji S. 2005. A smooth model of decision making under uncertainty. *Econometrica* 73:1849–92
- Knight FH. 1921. *Risk, Uncertainty, and Profit*. Boston: Houghton Mifflin
- Kreps DM, Porteus EL. 1978. Temporal resolution of uncertainty and dynamic choice. *Econometrica* 46:185–200
- Lemoine D, Traeger CP. 2016. Ambiguous tipping points. *J. Econ. Behav. Organ.* 132:5–18
- Lettau M, Wachter J. 2007. Why is long-horizon equity less risky? A duration-based explanation of the value premium. *J. Finance* 62:55–92
- Li X, Narajabad B, Temzelides T. 2016. Robust dynamic energy use and climate change. *Quant. Econ.* 7:821–57
- Lucas RE, Prescott EC. 1971. Investment under uncertainty. *Econometrica* 39:659–81
- Maccheroni F, Marinacci M, Rustichini A. 2006. Dynamic variational preferences. *J. Econ. Theory* 128:4–44
- MacDougall AH, Swart NC, Knutti R. 2017. The uncertainty in the transient climate response to cumulative CO₂ emissions arising from the uncertainty in physical climate parameters. *J. Climate* 30:813–27
- Matthews HD, Gillett NP, Scott PA, Zickfeld K. 2009. The proportionality of global warming to cumulative carbon emissions. *Nature* 459:829–32
- Miao J, Rivera A. 2016. Robust contracts in continuous time. *Econometrica* 84:1405–40
- Millner A, Dietz S, Heal G. 2013. Scientific ambiguity and climate policy. *Environ. Resour. Econ.* 55:21–46
- Morgan MG, Vaishnav P, Dowlatabadi H, Azevedo IL. 2017. Rethinking the social cost of carbon dioxide. *Issues Sci. Technol.* 33:43–50
- Nordhaus WD. 2008. *A Question of Balance: Weighing the Options on Global Warming Policies*. New Haven, CT: Yale Univ. Press
- Nordhaus WD. 2017. Revisiting the social cost of carbon. *PNAS* 114:1518–23
- Nordhaus WD, Moffatt A. 2017. *A survey of global impacts of climate change: replication, survey methods, and a statistical analysis*. NBER Work. Pap. 23646
- Parker JA, Julliard C. 2003. Consumption risk and expected stock returns. *Am. Econ. Rev. Pap. Proc.* 93:376–82
- Petersen IR, James MR, Dupuis P. 2000. Minimax optimal control of stochastic uncertain systems with relative entropy constraints. *IEEE Trans. Autom. Control* 45:398–412
- Pindyck RS. 2013. Climate change policy: What do the models tell us? *J. Econ. Lit.* 51:860–72
- Qin L, Linetsky V. 2020. Long-term risk: a martingale approach. *Econometrica* 85:299–312

- Ricke KL, Caldeira K. 2014. Maximum warming occurs about one decade after a carbon dioxide emission. *Environ. Res. Lett.* 9:1–8
- Ross SA. 2015. The recovery theorem. *J. Finance* 70:615–48
- Rubinstein M. 1976. The valuation of uncertain income streams and the pricing of options. *Bell J. Econ.* 7:407–25
- Saltelli A, Giampietro M. 2017. What is wrong with evidence based policy, and how can it be improved? *Futures* 91:62–71
- Savage LJ. 1954. *Foundations of Statistics*. New York: Wiley
- van Binsbergen J, Brandt M, Koijen R. 2012. On the timing and pricing of dividends. *Am. Econ. Rev.* 102:1596–618
- Weitzman ML. 2012. GHG targets as insurance against catastrophic climate damages. *J. Public Econ. Theory* 14:221–44



Contents

Emmanuel Farhi, Economist Par Excellence <i>Jean Tirole</i>	1
The Political Economy of Deep Integration <i>Giovanni Maggi and Ralph Ossa</i>	19
Large Games: Robustness and Stability <i>Ronen Gradwohl and Ehud Kalai</i>	39
Does Vote Trading Improve Welfare? <i>Alessandra Casella and Antonin Macé</i>	57
What Shapes the Quality and Behavior of Government Officials? Institutional Variation in Selection and Retention Methods <i>Claire S.H. Lim and James M. Snyder Jr.</i>	87
The Elusive Peace Dividend of Development Policy: From War Traps to Macro Complementarities <i>Dominic Rohner and Mathias Thoenig</i>	111
Why Does Globalization Fuel Populism? Economics, Culture, and the Rise of Right-Wing Populism <i>Dani Rodrik</i>	133
Systemic Risk in Financial Networks: A Survey <i>Matthew O. Jackson and Agathe Pernoud</i>	171
The International Aspects of Macprudential Policies <i>Kristin J. Forbes</i>	203
Estimating DSGE Models: Recent Advances and Future Challenges <i>Jesús Fernández-Villaverde and Pablo A. Guerrón-Quintana</i>	229
Firm Dynamics and Trade <i>George Alessandria, Costas Arkolakis, and Kim J. Ruhl</i>	253
The Economics of Currency Risk <i>Tarek A. Hassan and Tony Zhang</i>	281
Empirical Models of Industry Dynamics with Endogenous Market Structure <i>Steven T. Berry and Giovanni Compiani</i>	309

The Macroeconomics of Financial Speculation <i>Alp Simsek</i>	335
Uncertainty Spillovers for Markets and Policy <i>Lars Peter Hansen</i>	371
A Helicopter Tour of Some Underlying Issues in Empirical Industrial Organization <i>Ariel Pakes</i>	397
The Story of the Real Exchange Rate <i>Oleg Itskhoki</i>	423
Choice in Insurance Markets: A Pigouvian Approach to Social Insurance Design <i>Nathaniel Hendren, Camille Landais, and Johannes Spinnewijn</i>	457
The Econometrics of Early Childhood Human Capital and Investments <i>Flavio Cunha, Eric Nielsen, and Benjamin Williams</i>	487
Sufficient Statistics Revisited <i>Henrik J. Kleven</i>	515
The Blossoming of Economic Epidemiology <i>David McAdams</i>	539
Directed Technical Change in Labor and Environmental Economics <i>David Hémous and Morten Olsen</i>	571
Inflation Inequality: Measurement, Causes, and Policy Implications <i>Xavier Jaravel</i>	599
Theoretical Foundations of Relational Incentive Contracts <i>Joel Watson</i>	631

Indexes

Cumulative Index of Contributing Authors, Volumes 9–13	661
--	-----

Errata

An online log of corrections to *Annual Review of Economics* articles may be found at
<http://www.annualreviews.org/errata/economics>