



Stochastic responses and marginal valuation

Lars Peter Hansen^{a,1,2} and Panagiotis Souganidis^{b,2}

Affiliations are included on p. 11.

Contributed by Lars Peter Hansen; received August 1, 2025; accepted October 20, 2025; reviewed by Dennis Epple and Xun Yu Zhou

The analysis of policy impacts in a dynamic and uncertain reality is vital to supporting informed economic policy design and implementation. Dynamic, stochastic economic models used in policy evaluation necessarily simplify the world as we know it. This motivates us to explore, refine, and extend tools aimed at producing marginal valuations that shed light on why some policies are optimal and how others, though suboptimal, can be improved. We present representations of these marginal valuations that embrace uncertainty and support robust implementation—even in environments characterized by “deep uncertainties.” These representations offer a more complete understanding of how interactions among multiple state variables, concerns about model misspecification, and uncertainties surrounding potentially long-term implications contribute to the cogent assessment of policies. We argue that these methods are particularly salient for evaluating the global cost of climate change and the global value of research and development with long-term prospects for success.

policy assessment | valuation | deep uncertainty | robustness | stochastic differential equations

1. Introduction

Marginal valuations play a central role in economic analyses and serve as tools for both interpretation and measurement. This paper develops such constructs in dynamic settings relevant for investment and policy analyses for which uncertainty is pervasive. We imagine a Markov environment with multiple, interacting, endogenous state variables. Small changes in one of these variables at an initial period have marginal outcome consequences in future dates that we seek to value. As an example, an initial change in some form of capital—induced, for example, by an initial investment—has payoff ramifications for several periods in the future. As another example, a marginal increase in climate change has consequences for economic opportunities in future periods. Finally, proposals such as ref. 1 for using net national product to assess policies in the presence of externalities use marginal calculations to deduce the prices relevant for social valuation.

Uncertainty contributions can arise in multiple forms, including randomness captured by repeated small shocks and infrequent big events. In contrast to more narrowly framed risk analyses, the marginal valuations we explore entertain a broad conception of such uncertainties, rich enough to encapsulate model ambiguities and specification concerns.

While intertemporal marginal assessments are constructed to be local in nature, they complement more global assessments. They suggest where actions or policy changes can have large impacts, and they isolate when such modifications are inconsequential. More formally, they provide insights about global solutions to control problems as the first-order conditions for optimization typically depend on marginal valuations. Intertemporal marginal valuations are commonly used in economics and policy making, including some prominent measurements of the social cost of carbon and marginal abatement costs, and are used in assessing the welfare implications of marginal changes in exposures to uncertainty.^{*,†} Some applications rely on explicitly specified models of state dynamics measured as partial derivatives of value functions. Other applications are more empirically grounded but abstract from fully explicit state dynamic modeling including the interactions among state variables.

^{*}There is a wide variety of measurement of the social cost of carbon using different approaches. For some examples, see refs. 2–4 for alternative approaches to measuring or estimating the social cost of carbon. ref. 2's measurements are based on solutions to a planner's problem that include explicit risk considerations. Ref. 4 uses a more modular approach to incorporate a wide variety of alternative sources of risk. To the extent that their analysis is marginal, it is relative to a suboptimal dynamic allocation. Ref. 3 compares multiple approaches but abstracts from risk and “inside the decision problem” uncertainties that we explore here.

[†]Refs. 5 and 6 use intertemporal marginal valuations to measure the importance of aggregate uncertainty. Ref. 5 uses this approach to reveal implications of aversion to model misspecification. Alvarez and Jermann (6) use it to address a quantitative challenge posed by Lucas (7) on the welfare consequences of macroeconomic fluctuations. In this paper, we provide revealing representations of marginal valuation that support interpretation and measurement and open the door to critical assessments of alternative applications.

Significance

Intertemporal marginal valuations are central to economic decision making and policy analysis. They guide decisions that are robustly optimal and, even in nonoptimal settings, offer informative characterizations for policy improvement. Relevant applications include the social impact of climate change and the social value of research and development. There are a wide range of applications in public economics, environmental economics, and macroeconomics. Our formulation disentangles the roles of modeling assumptions and empirical inputs by enabling conceptual clarity and illuminating measurement challenges. Designed for broad applicability, the analysis explicitly accommodates deep uncertainties, including model ambiguity and misspecification. In doing so, it provides a rigorous and transparent foundation for both economic analysis and practical policy evaluation in uncertain and evolving environments.

Author contributions: L.P.H. and P.S. performed research; and wrote the paper.

Reviewers: D.E., Carnegie Mellon University; and X.Y.Z., Columbia University.

The authors declare no competing interest.

Copyright © 2025 the Author(s). Published by PNAS. This open access article is distributed under Creative Commons Attribution-NonCommercial-NoDerivatives License 4.0 (CC BY-NC-ND).

¹To whom correspondence may be addressed. Email: lhansen@uchicago.edu.

²L.P.H. and P.S. contributed equally to this work.

This article contains supporting information online at <https://www.pnas.org/lookup/suppl/doi:10.1073/pnas.2520857122/-DCSupplemental>.

Published November 24, 2025.

In many economic settings, uncertainty would seem to be a first-order consideration. However, typical calculations either abstract from uncertainty or focus on risk specifications with full confidence in the probabilistic input. We suggest ways to relax that confidence and show how it impacts valuation. As some researchers and policymakers have suggested, deep uncertainties pervade dynamical systems in ways that should impact the design of prudent climate change and other policies, a view that we embrace. For example, see refs. 8–11 for discussions of such challenges in the economics literature. Moreover, a plethora of research in macroeconomics treats aggregate uncertainty as a second-order consideration despite the limited understanding of some aspects of the underlying dynamics. Finally, we use the jump diffusion model to include infrequent big events, along with the repeated small shocks (Brownian increments). Importantly, the infrequent jumps have endogenous intensities and uncertain dynamics.

Although we build on ideas, described in refs. 12–14, and elsewhere, justifying decision theory approaches that provide a way to confront such challenges, our emphasis is different. Indeed, we focus on the construction of i) uncertainty-adjusted probabilities used in valuation, and ii) revealing representations and decompositions of these valuations.

In addition to decision theory, we borrow and extend ideas from asset pricing to obtain our representation and decomposition of marginal valuation. We show formally how to represent marginal values as asset prices with implicit marginal payoffs. Depending on the application, these payoffs could be economic outputs or social payoffs and are the counterpart to dividends of cash flows from market-based asset pricing. Our construction of marginal payoffs provides us with ways to find revealing additive decompositions of marginal valuations. These decompositions quantify how dynamic interactions across state variables contribute to the overall valuation. In applications with so-called “structural models” of economic dynamics, the methods will help researchers to deconstruct model implications and to assess roles of the alternative “moving parts.” The resulting valuation formulas use stochastic responses to initial perturbations in ways that are similar to sensitivity characterizations of derivative claims in mathematical finance.

Our approach to uncertainty is encapsulated in an uncertainty-adjusted probability measure. We derive this adjustment for alternative characterizations of the underlying uncertainties. The use of a change-of-measure to represent asset values, while conceptually distinct, is reminiscent of the risk-neutral change in probability measure used extensively in derivative claim pricing.

We explore this topic for continuous-time stochastic systems, although many of the insights derived are also applicable to discrete-time formulations. We start with Markov diffusion processes analyzed in the familiar risk environment with full confidence in probabilities. Our results in this baseline case are given in Sections 2 and 2.2. In Section 3, we allow for ambiguities and robustness considerations expressed in terms of likelihood and prior uncertainties. Finally, we include infrequent Poisson-type jumps with endogenous intensities and explore their consequences in Section 4.

Many of the basic mathematical ingredients in our analysis were developed and used previously in related forms of parameter sensitivity analyses for continuous-time diffusions. See, for example, refs. 15 and 16, along with the extensive list of references in the latter paper. A substantial focus of this literature is computation, while our motivation is to develop a framework that supports measurement, model interpretation, and deconstruction in the presence of broadly conceived uncertainties.

2. Impulse Responses and Diffusion dynamics

The use of impulse responses has been a common tool in the study of economic dynamical systems. See ref. 17 for an initial application of this approach. For nonlinear stochastic models, there are a variety of alternative approaches. Consistent with our interest in marginal valuation, we follow the approach of refs. 18 and 19 by which responses are deduced for marginal changes in an initial time period. We refer to the computations as local impulse responses. In what follows, we show the dynamic stochastic evolution of these responses for Markov diffusions.

Consider an N -dimensional vector system of the state dynamics given by the stochastic differential equation (SDE):

$$dX_t = \mu(X_t)dt + \sum_{\ell=1}^L \sigma_{\ell}(X_t)dW_t^{\ell} \quad X_0 = x, \quad [1]$$

where W is an L -dimensional, standard Brownian motion with entry ℓ denoted W^{ℓ} , μ and, the σ_{ℓ} 's are N -dimensional column vectors with the same number of entries as there are states. In addition to the standard assumptions that guarantee the well-posedness of Eq. 1, we restrict the matrix $\sigma\sigma'$ to be uniformly elliptic (its determinant is bounded away from zero).[‡] The infinitesimal generator \mathbb{A} for the diffusion is written, for sufficiently smooth $\phi : \mathbb{R}^N \rightarrow \mathbb{R}$, in the familiar form:

$$\mathbb{A}\phi \stackrel{\text{def}}{=} \mu \cdot \frac{\partial \phi}{\partial x} + \frac{1}{2} \text{trace} \left(\sigma' \frac{\partial^2 \phi}{\partial x \partial x'} \sigma \right). \quad [2]$$

Initially, we consider an expected discounted value objective of a decision maker.

$$V(x) \stackrel{\text{def}}{=} \delta \mathbb{E} \left[\int_0^{\infty} \exp(-\delta t) U(X_t) dt \mid X_0 = x \right].$$

This decision maker could be a private investor or a social planner. We adopt this starting point for the sake of pedagogical convenience, although our analysis will have a more broader applicability. Later, we explore two extensions that are central to our analysis.

It is well known that the process V solves a Feynman-Kac (FK) or equivalently a resolvent equation of the form:[§]

$$0 = \delta U - \delta V + \mathbb{A}V. \quad [3]$$

Since we imagine a decision maker in the background, we suppose that he or she has a Markov control law and that posited μ and σ embed this decision rule in their construction. As we are potentially interested in suboptimal controls or policies, we do not necessarily start with an optimal problem.[¶]

2.1. Stochastic Responses. For decision and policy analyses, it is often of interest to study marginal improvements from potentially suboptimal decision rules. Then it becomes important to consider the marginal valuation $\frac{\partial V}{\partial x_i}$ for alternative coordinates. Depending

[‡]This restriction is known to be too strong for some applications of interest and can be relaxed in alternative ways depending on more specific aspects of the dynamic, stochastic system. One such condition that can sometimes be verified is what is called *partial ellipticity*. See, for instance, ref. 16 for a discussion.

[§]When $\sigma\sigma'$ is uniformly elliptic, V is a smooth solution to Eq. 3. If $\sigma\sigma'$ is degenerate elliptic, then V may not be classical and should be taken to be a viscosity solution to Eq. 3.

[¶]It is often the case that there are controls in Eq. 1, in which case the expected discounted value objective is defined as sup over all possible controls and Eq. 3 becomes a nonlinear pde which can be analyzed using the theory of viscosity solutions.

on the application, these could be measures of the shadow value of capital, the social cost of climate change, or the value of research and development supporting large-scale projects.

We start by characterizing what is called the first variational process.[#] A straightforward application of the chain rule yields that the Jacobian $Y_t = D_x X_t$ satisfies the SDE:

$$dY_t = D\mu(X_t)Y_t dt + \sum_{\ell=1}^L D\sigma_\ell(X_t)Y_t dW_t^\ell, \quad Y_0 = I_N.$$

For the local impulse response theory, we introduce a stochastic process (called a variational process) $\Lambda_t = Y_t \lambda$, which gives a vector of marginal responses of X to some date zero marginal change λ in the state vector. Given the SDE of Y , it is immediate that Λ_t evolves according to:

$$d\Lambda_t = D\mu(X_t)\Lambda_t dt + \sum_{\ell=1}^L D\sigma_\ell(X_t)\Lambda_t dW_t^\ell, \quad \Lambda_0 = \lambda. \quad [4]$$

In effect, the Λ evolution replaces the entries of $\mu(X_t)$ and $\sigma(X_t)$ with the dot products between the corresponding gradients evaluated at X_t and Λ_t . The specification of the initial condition λ will be one of the N coordinate vectors when we are interested in the impact of marginal changes in one of the state variables in the initial time period.^{||} Other initial conditions are also of interest, however, as we discuss subsequently.

The previous calculations can also be justified using Malliavin Calculus given the underlying Brownian information structure—we refer to refs. 16, 18, and 19 for an extensive discussion and examples. For the discussion here, there is no need to use Malliavin Calculus. Such methods, however, may help to simplify some of the formulas derived later in concrete examples.

In view of the interactions with the state dynamics, we study (X, Λ) as a vector system with generator \mathbb{B} . We are particularly interested in functions of the form: $\Psi(x, \lambda) = \psi(x) \cdot \lambda$, where ψ is a N -dimensional vector of functions of x . Let \mathbb{A} act on the vector of functions ψ componentwise. Then, we find:

$$\begin{aligned} \mathbb{B}\Psi &= \lambda \cdot (\mathbb{A}\psi) + (D\mu\lambda) \cdot \psi \\ &+ \frac{1}{2} \text{trace} \left([D\sigma_1\lambda \quad \dots \quad D\sigma_L\lambda]' \frac{\partial \psi}{\partial x} \sigma \right. \\ &\left. + \sigma' \frac{\partial \psi'}{\partial x} [D\sigma_1\lambda \quad \dots \quad D\sigma_L\lambda] \right). \end{aligned}$$

It follows that the outcome can be expressed in the form: $\lambda \cdot \tilde{\psi}(x)$ for a N -dimensional vector of functions $\tilde{\psi}$.

Remark 2.1: Discrete-time counterparts to linear systems, like, for example, Eq. 1 with μ linear in the state vector and σ state independent, are often used in approximations in macroeconomics. Under these restrictions on the diffusion dynamics, Λ can be solved analytically.

[#]Although the calculations here are straightforward, since we are dealing with a linear uniformly elliptic pde, we present them to familiarize the reader. For the models later in the paper, however, it becomes necessary to use more sophisticated mathematical arguments coming from the theory of nonlinear pdes.

^{||}There is a substantial applied macroeconomics literature that engages in “innovation accounting” to assess the impacts of different initial shocks within a dynamic linear systems. See ref. 20 for an initial and prominent contribution to this literature. For such applications, the counterpart for the initial λ is the impact vector of one of the initial shocks on the initial state evolution. See ref. 19 for an elaboration.

2.2. An Initial Representation. Going forward, we are interested in the process of stochastic (marginal) responses of future $U(X)$ to an initial change in one of the states. We follow the economic convention of referring to $U(x)$ as the “utility” given state realization x , although we could instead follow a convention in control theory of viewing $U(x)$ as a measure of a current period contribution to costs and obtain analogous results.^{**} The marginal responses are given by the process $\left\{ \frac{\partial U}{\partial x}(X_t) \cdot \Lambda_t : t \geq 0 \right\}$. Investigating the behavior of this process requires that we study the joint dynamics (X, Λ) .

We construct a representation for each entry of the gradient DV , and use this representation to perform an intertemporal decomposition that quantifies the importance of state variable interdependencies in representing marginal values.

For convenience, we study $DV \cdot \lambda$ where we are particularly interested in coordinate vectors λ 's that select particular coordinates of DV . Using the process (X, Λ) we find the representation:

$$\begin{aligned} DV(x) \cdot \lambda &= \delta \mathbb{E} \left[\int_0^\infty \exp(-\delta t) DU(X_t) \cdot \Lambda_t dt \mid X_0 = x, \Lambda_0 = \lambda \right]. \end{aligned} \quad [5]$$

Since the initial state variable perturbation can alter future values of the entire state, we include the entire gradient vector $DU(X_t)$ at date t . The dot product between DU and Λ reflects a version of the chain rule.

Differentiating Eq. 3 with respect to each x_j gives a system of equations for the vector DV of functions. We reduce this system to a scalar equation by premultiplying the vector λ . This reduction is of interest precisely because it results in a single FK equation with a substantive interpretation that is revealing. Specifically, it provides the following scalar FK equation for $DV \cdot \lambda$:

$$0 = \delta DU(x) \cdot \lambda - \delta DV(x) \cdot \lambda + \mathbb{B} [DV(x) \cdot \lambda]. \quad [6]$$

Formula Eq. 5 gives the stochastic representation of the solution to FK equation Eq. 6.

Eq. 5 opens the door to an “asset pricing” interpretation for the marginal valuation $\frac{\partial V}{\partial x_j}$ as a discounted expected value of “cash flow” $DU(X_t) \cdot \Lambda_t$ that is the stochastic utility response to an initial perturbation in the i th component of the state vector when initializing Λ_0 at the i th coordinate vector.

Remark 2.2: The study of control problems will sometimes use Hamiltonian methods to deduce first-order conditions for optimization. Early prominent examples of this in economics include (21) in the case of optimal deterministic growth models and ref. 22 in the case of the stochastic counterpart. Specifically, such methods give rise to a costate equation that is included when constructing an optimal solution. There is a known equality between a costate and a derivative of a value function with respect to a corresponding state variable. This costate is often interpreted as a measure of a marginal valuation.

Our characterization provides a way to represent the marginal valuation that is distinct from a costate equation and opens the door to further interpretation. The equation of interest is different from the costate equation system for two reasons.

^{**}In environments with uncertainty, economic measures of costs typically involve stochastic discounting determined endogenously. Constant discounting is instead used to represent preferences expressed as discounted expected utilities. See Remark 2.3 for an illustration of stochastic discounting.

First, we do not exploit the optimality of the control law to allow for a more general applicability of our analysis. Instead, we take the Markov control law or the Markov equilibrium conditions given when deducing our representation of interest to use for interpretation. Costate evolutions are typically derived to be solved as part of constructing a solution to the optimal control problem and the impact of optimization is reflected in the costate dynamics. Of course, we could simply focus on a degenerate case of optimality in which the control law is constrained *ex ante* to equal a prespecified function of the state. The resulting costate evolution can be viewed as an input into our analysis.

Second, we do not directly represent the solution to this vector system of backward SDE's for $DV(X)$. Instead we focus on an interpretable, state-dependent linear combination $\Lambda \cdot DV(X)$. Stochastic response theory gives us the evolution of Λ that supports this reduction.^{††} We then solve the single backward SDE in $\Lambda \cdot DV(X)$, which has the form of the Feynman–Kac equation, taking as input the forward SDE for the joint (X, Λ) process. Our reason for doing this is neither for mathematical nor computational advantages; but it is rather because it results in a substantively interesting and economically motivated representation. *SI Appendix, section 1* gives an illustration of these various points using a linear-quadratic control problem.

Remark 2.3: Asset prices are often represented in terms of stochastic discount factors. To make this connection precise, consider a utility function U of the form:

$$U(x) = \tilde{U}[c(x)],$$

where c denotes consumption, which can depend on the Markov state. By the chain rule,

$$\frac{\partial U}{\partial x} = \frac{d\tilde{U}}{dc} \frac{\partial c}{\partial x}.$$

Thus, the stochastic flow to be valued is the stochastic response of consumption times the marginal utility at the date of the flow. By converting the marginal valuation at date zero into units of consumption at date zero, we obtain the formula:

$$\frac{DV(x) \cdot \lambda}{\frac{\partial \tilde{U}}{\partial c}(C_0)} = \mathbb{E} \left(\int_0^\infty S_t \left[\frac{\partial c}{\partial x}(X_t) \cdot \Lambda_t \right] dt \mid X_0 = x, \Lambda_0 = \lambda \right), \quad [7]$$

where C_t denotes the stochastic consumption for date t . This representation includes a date t stochastic discount factor:

$$S_t \stackrel{\text{def}}{=} \exp(-\delta t) \left[\frac{d\tilde{U}}{dc}(C_t) \right].$$

The division on the left side of formula Eq. 7 converts the marginal valuation into units of consumption. The stochastic discount factor process $\{S_t : t \geq 0\}$ is familiar from asset pricing theory as a measure of the intertemporal marginal substitution rates expressed in units of consumption at the respective dates.^{‡‡}

^{††} See, for instance, ref. 23 for more general discussions of infinite horizon, forward-backward SDE's.

^{‡‡} For pedagogical simplicity, this illustration presumes decision makers' preferences that are time separable. There is a corresponding marginal utility adjustment for temporally dependent preferences, such as those suggested by Ryder and Heal (24) and many others studying economic dynamics.

For example, see refs. 25–27. It follows from the right side of Eq. 7, that the process $\{S_t : t \geq 0\}$ is used to discount the stochastic flow process $\left\{ \frac{\partial c}{\partial x}(X_t) \cdot \Lambda_t : t \geq 0 \right\}$.

When Λ_0 is initialized at e_i (the i th coordinate vector), we obtain the marginal valuation of state i ; for explicit policy, we are interested in other initializations as illustrated in the following remark.

Remark 2.4: Let $\Gamma(x)$ denote a given, possibly suboptimal, control law where:

$$\begin{aligned} \mu(x) &= \hat{\mu}[x, \Gamma(x)] \\ U(x) &= \hat{U}[x, \Gamma(x)] \end{aligned}$$

To investigate the impact of a “small” change in Γ , first write $\Gamma + \epsilon \bar{\Gamma}$. Substitute in the modified control law and differentiate the altered Feynman–Kac equation for an ϵ -dependent value function with respect to ϵ for the given initial value function and its derivatives. By including all of the terms, we obtain a Feynman–Kac equation for the derivative of the value function with respect to ϵ for which the term of interest is the “flow term:”

$$\delta \frac{\partial \hat{U}}{\partial \gamma} \cdot \bar{\Gamma} + \frac{\partial V}{\partial x} \cdot \left[\frac{\partial \hat{\mu}}{\partial \gamma'} \bar{\Gamma} \right]. \quad [8]$$

A solution is a discounted (conditional) expected value of the implied flows. When Γ is chosen optimally, the sum of the two terms in Eq. 8 is zero achieved by setting marginal costs equal to marginal benefits. When the discounted expected value is of the flows implied by Eq. 8 is positive, $\bar{\Gamma}$ isolates a direction for policy improvement for the associated state vector. For a variety of applications, we can interpret separately the two terms in Eq. 8 in terms of marginal costs and benefits. Given the second term, we are led to set the initial condition for the Λ process of stochastic responses to:

$$\Lambda_0 = \frac{\partial \hat{\mu}}{\partial \gamma'} [X_0, \Gamma(X_0)] \bar{\Gamma}(X_0).$$

These formulas represent costs and benefits in terms of the utility numeraire. In economic applications, a further division by a current period marginal utility of a consumption numeraire converts the measures in units of that numeraire.

Example 2.5: *SI Appendix, section 6* gives a climate–economics example. In the resulting model, there are three state variables, temperature, a broadly based capital used in producing output, and a measure of knowledge capital that makes new green technology discovery more likely. The controls include a dirty form of energy that emits CO₂, an investment in the broadly based capital, and an investment in knowledge capital. From the guises of this model, there are two cost–benefit style computations that are of interest.

First, construct $\bar{\Gamma}$ to be a coordinate vector with a one in entry corresponding to the energy input. This energy input enhances current-period product, but it also induces warming. The marginal social cost of this dirty energy source is

$$-\frac{\partial V}{\partial x} \cdot \left[\frac{\partial \hat{\mu}}{\partial \gamma'} \bar{\Gamma} \right].$$

By converting the energy input into units of the implied CO₂ emissions, this measurement becomes the commonly discussed

social cost of carbon. We refer to it as social because it includes costs that are not internalized by decentralized markets in the absence of targeted taxes. The marginal benefit measures the contribution of the energy input into production, which in turn induces increases in consumption. Formally, the resulting marginal benefit is measured by

$$\frac{\partial \widehat{U}}{\partial \gamma} \cdot \bar{\Gamma}$$

using the same choice of $\bar{\Gamma}$. By design, our representation reveals the important components to the social cost of carbon.

As a second computation with this same model, observe that a larger stock of knowledge makes technology discovery of a new clean energy source more likely. This is the source of a benefit to investment in green research and development. This investment is costly because it reduces the output available for other activities such as consumption or investment.

We now let $\bar{\Gamma}$ be a coordinate vector with a one in the entry corresponding to investment in green research and development. In this application,

$$\frac{\partial V}{\partial x} \cdot \left[\frac{\partial \hat{\mu}}{\partial \gamma'} \bar{\Gamma} \right]$$

is a measure of the marginal benefit to investing in the stock of knowledge, whereas $-\frac{\partial \widehat{U}}{\partial \gamma} \cdot \bar{\Gamma}$ captures the marginal cost.

In conclusion, our representation provides insights into what contributes intertemporally to the social value of investment in green research and development.

2.3. An Additive Decomposition of the Marginal Values. The formula for $DV \cdot \lambda$ opens the door to a decomposition by states. Suppose that we initialize $\Lambda_0 = \lambda$ where λ is a coordinate that isolates the initial state of interest for the marginal valuation. Then Λ_t measures how a perturbation in this initial state affects the entire state vector at future dates. Each state variable process thus gives a channel by which the change in an initial state impacts valuation. We use these channels to give an additive decomposition of the marginal valuation.

This valuation decomposition is based on an additive decomposition of the stochastic flows from each of the N different states. Write

$$DU = \left[\frac{\partial U}{\partial x_1} \dots \frac{\partial U}{\partial x_N} \right] \text{ and } \Lambda_t' = [\Lambda_{1,t} \dots \Lambda_{N,t}].$$

Then the stochastic flow is decomposed as

$$DU(X_t) \cdot \Lambda_t = \sum_{j=1}^N \frac{\partial U}{\partial x_j}(X_t) \Lambda_{j,t}.$$

While the choice of initial condition λ , determines the marginal valuation of interest, the subscript j determines the state variable channel that contributes to this marginal valuation.

2.4. Decomposition of Values. This additive decomposition of flows implies a corresponding decomposition of values.

$$DV(x) \cdot \lambda = \sum_{j=1}^N V_j(x) \cdot \lambda, \quad [9]$$

where:

$$V_j(x) \cdot \lambda \stackrel{\text{def}}{=} \delta \mathbb{E} \left[\int_0^\infty \exp(-\delta t) \frac{\partial U}{\partial x_j}(X_t) \Lambda_{j,t} dt \mid X_0 = x, \Lambda_0 = \lambda \right].$$

The terms in this marginal valuation decomposition reflect the impact of changes in an initial state (specified by the choice of λ) on the alternative states in the future, as they contribute to the valuation.

As a motivational example, some measurements of the social cost of climate change leave out the stock of knowledge associated with investments in research and development, or possibly other suboptimal government policy responses to future changes in climate change. This decomposition helps us understand both the empirical challenges and the potential drawbacks of ignoring some of the states that interact with climate change.

3. Robust Marginal Valuation

So far, the only form of uncertainty that we have considered is a narrowly defined notion of *risk*. Risk captures unknown outcomes with known probabilities such as Brownian motion risk. We are influenced by Knight's (28) assertion that:

Uncertainty must be taken in a sense radically distinct from the familiar notion of risk, ... and there are far-reaching and crucial differences in the bearings of the phenomena depending on which of the two is really present and operating.

Such concerns appear in a variety of discussions of policy challenges, including climate change, with a recent example being ref. 11. However, the ramifications for marginal valuation have not been formally addressed, a topic we consider next.

We describe formally how to incorporate uncertainty adjustments that relax confidence in the baseline specification of the Brownian motion and instead allow for baseline probabilities to be misspecified. Conveniently, we will show that these adjustments are captured by a change in the baseline probability distribution, a mathematical device that is familiar to derivative claims valuation and financial engineering, but the uncertainty-adjusted probability is constructed in an entirely different way.

Although there are a variety of lines of attack that have been suggested to confront "deep uncertainty," we focus on three interrelated approaches. In all cases, we adopt a recursive formulation that leads to straightforward modifications of the Feynman-Kac equations. Although the approaches differ in terms of the structure that is imposed on the uncertainty, all three of them use recursive versions of variational preferences. (See (29) for an axiomatic defense of these preferences.) We note that despite the commonality to their mathematical structure that we exploit, there are important conceptual differences among the three.

3.1. Robust Valuation with Model Uncertainty. We follow refs. 30–32 by exploring changes in the baseline probability specification, which is assumed to be a multivariate standard Brownian motion W .

Following the well-known Girsanov characterization of positive martingales, we parameterize the local evolution with the process H and consider a martingale M such that

$$dM_t = M_t H_t \cdot dW_t \quad M_0 = 1.$$

The resulting probability measure changes W to satisfy

$$dW_t = H_t dt + dW_t^H, \quad [10]$$

where $\{W_t^H : t \geq 0\}$ is a standard Brownian motion under the probability specification induced by the martingale.

The decision maker entertains these alternative probabilities subject to a relative entropy penalization. Expressed with the new Brownian motion, the dynamics of the state process are

$$dX_t = [\mu(X_t) + \sigma(X_t)H_t] dt + \sigma(X_t)dW_t^H.$$

Since the alternative probability remains locally normally distributed, the local contribution to the log-likelihood ratio between the processes evolves as

$$d \log M_t = H_t \cdot dW_t - \frac{1}{2} H_t \cdot H_t = \frac{1}{2} H_t \cdot H_t dt + H_t \cdot dW_t^H$$

which has a local mean of $\frac{1}{2} H_t \cdot H_t$ under the alternative distribution. This gives the local contribution to the relative entropy.

To make a robust adjustment to the valuation, we consider the value function:

$$V(x) = \min_H \mathbb{E} \left(\int_0^\infty \exp(-\delta t) M_t \left[\delta U(X_t) + \frac{\xi}{2} H_t \cdot H_t \right] dt \mid X_0 = x, M_0 = 1 \right), \quad [11]$$

where the expectation is taken under the original probability distribution with the change in probability measure embedded in the positive martingale M . The positive parameter, ξ , determines the magnitude of the penalty for deviating from the baseline probability specification. Smaller values of ξ imply more aversion to misspecification uncertainty on the part of the decision maker. The state variable M induces a linear scaling of the value function, which we write as mV . Specifically, mV , solves the Hamilton–Jacobi–Bellman (HJB) equation:

$$0 = \delta mU - \delta mV + mAV + \min_b \left\{ (mDV) \cdot (\sigma b) + m \frac{\xi}{2} b' b \right\}.$$

Since the positive variable m scales the HJB equation proportionately, we instead focus on the equation for V , which also satisfies an HJB equation^{§§}

$$0 = \delta U - \delta V + AV(x) + \min_b \left\{ DV \cdot (\sigma b) + \frac{\xi}{2} b' b \right\}. \quad [12]$$

Notice that the minimization problem has the quasi-analytic solution:

$$b^* = -\frac{1}{\xi} \sigma' DV. \quad [13]$$

With this in hand, we find that the minimized objective solves the HJB:

$$0 = \delta U - \delta V + AV(x) + (\sigma b^*) \cdot DV + \frac{\xi}{2} b^* \cdot b^*. \quad [14]$$

Next, we propose two alternative ways to apply our previous analysis to characterize $DV \cdot \lambda$.

3.1.1. Approach 1. We show here how to use the Envelope Theorem in conjunction with the first-order condition Eq. 13 to eliminate some of the “indirect terms” involving Dh^* when we differentiate the HJB-equation. Indeed, it follows from Eq. 13 that

$$0 = (Dh^*)' [\sigma' DV + \xi b^*] = [\sigma Dh^*]' DV + \xi (Dh^*)' b^*. \quad [15]$$

As a consequence, the stochastic flow term is $DU \cdot \lambda$, as it was without robustness considerations, and the drifts or local means for X and Λ are respectively:

$$\mu(x) + \sigma(x)b^*(x) \text{ and } D\mu(x)\lambda + [D\sigma_{\cdot 1}(x)\lambda \dots D\sigma_{\cdot L}(x)\lambda] b^*(x).$$

However, with the simplification of the envelope Theorem, the resulting Λ process ceases to encode the stochastic responses that provide the most substantively interesting interpretation. By imposing the first-order conditions, we dropped the drift contribution σDh^* in the evolution of Λ when imposing b^* . As a consequence, the resulting Λ process is different and no longer can be interpreted as encoding the stochastic response vectors of interest.

One possibility is that we do not impose Eq. 15 and include $[\sigma Dh^*]'$ as part of the evolution for λ . In this case, we would also need to add $\xi (Dh^*)' b^*$ to the stochastic flow term $DU \cdot \lambda$.

As a substantively interesting alternative, we consider a second approach, one that takes advantage of the envelope condition and supports the construction of an uncertainty-adjusted probability distribution.

3.1.2. Approach 2. Suppose that the decision maker considers drift distortions that are perceived not to depend on endogenous state variables. In other words, the minimization explores altered probability specifications that are beyond the control of the decision maker. To represent this result, we set $\bar{H}_t^* = b^*(\bar{X}_t)$ and construct an artificial process $\{\bar{X}_t : t \geq 0\}$ evolving according to

$$d\bar{X}_t = [\mu(\bar{X}_t) + \sigma(\bar{X}_t) b^*(\bar{X}_t)] dt + \sigma(\bar{X}_t) dW_t^{\bar{H}^*} \quad \bar{X}_0 = \bar{x}. \quad [16]$$

The original process $\{X_t : t \geq 0\}$ is now replaced by

$$dX_t = [\mu(X_t) + \sigma(X_t) b^*(\bar{X}_t)] dt + \sigma(X_t) dW_t^{\bar{H}^*} \quad X_0 = x. \quad [17]$$

This construction captures the idea that the drift distortion b^* is specified exogenously and contributes to the marginal valuation only through a change in the probability measure. Although the process H^* still depends on the outcome of the minimization problem, for the purposes of evaluation, the decision maker does not view it as a function of the endogenous state variables.^{§§} By design, the resulting marginal valuation is the same as in Approach 1, but the interpretation is altered. See [SI Appendix, section 3](#) for a formal elaboration.

With this in mind, the martingale used to represent the change in probability evolves as

$$d\bar{M}_t^* = \bar{M}_t^* \bar{H}_t^* dW_t \quad \bar{M}_0^* = 1.$$

^{§§}Dynamic economic modelers might see a very rough analogy here between this construction and that of the big K , little k trick sometimes used in representing equilibrium outcomes. See ref. 33, pages 226, 470, and 489, for a textbook discussion of the big K , little k approach.

The value function $\bar{V} = \bar{V}(x, \bar{x})$ given by

$$\bar{V}(x, \bar{x}) = \mathbb{E} \left(\int_0^\infty \exp(-\delta t) \bar{M}_t^* [\delta U(X_t) + \frac{\xi}{2} b^*(\bar{X}_t)' b^*(\bar{X}_t)] dt \mid X_0 = x, \bar{X}_0 = \bar{x}, \bar{M}_0^* = 1 \right) \quad [18]$$

solves the modified HJB equation:

$$0 = \delta U(x) - \delta \bar{V}(x, \bar{x}) + \bar{\mathbb{A}} \bar{V}(x, \bar{x}) + \frac{\xi}{2} b^*(\bar{x})' b^*(\bar{x}), \quad [19]$$

where $\bar{\mathbb{A}}$ is constructed with a composite drift and Brownian exposure matrices given by

$$\bar{\mu}(x, \bar{x}) = \begin{bmatrix} \mu(x) + \sigma(x) b^*(\bar{x}) \\ \mu(\bar{x}) + \sigma(\bar{x}) b^*(\bar{x}) \end{bmatrix} \text{ and } \bar{\sigma}(x, \bar{x}) = \begin{bmatrix} \sigma(x) \\ \sigma(\bar{x}) \end{bmatrix}.$$

Since we are only interested in the stochastic responses associated with X captured by Λ , and not those \bar{X} , we modify the stochastic response evolution Eq. 4:

$$d\Lambda_t = D\mu(X_t)\Lambda_t dt + [D\sigma_1(X_t)\Lambda_t \quad \dots \quad D\sigma_L(X_t)\Lambda_t] b^*(\bar{X}_t) dt + [D\sigma_1(X_t)\Lambda_t \quad \dots \quad D\sigma_L(X_t)\Lambda_t] dW_t^{\bar{H}^*}, \quad [20]$$

which does not depend on the stochastic responses of \bar{X} , only on the process of \bar{X} itself.

This construction and analysis justifies the robust extension of Eq. 5:

$$\begin{aligned} & \frac{\partial \bar{V}}{\partial x}(x, x) \cdot \lambda \\ &= \mathbb{E} \left[\int_0^\infty \bar{M}_t^* \exp(-\delta t) \frac{\partial U}{\partial x}(X_t) \cdot \Lambda_t dt \mid X_0 = \bar{X}_0 = x, \Lambda_0 = \lambda \right] \\ &= \tilde{\mathbb{E}} \left[\int_0^\infty \exp(-\delta t) \frac{\partial U}{\partial x}(X_t) \cdot \Lambda_t dt \mid X_0 = x, \Lambda_0 = \lambda \right], \quad [21] \end{aligned}$$

where $\tilde{\mathbb{E}}$ is an uncertainty-adjusted expectation implied by the joint stochastic evolution given in Eq. 16 and 20 induced by the martingale \bar{M}^* along with the initial condition that $\bar{X}_0 = x$. Notice that we dropped $\xi [Dh^*(\bar{x})]' b^*(\bar{x})$ from the discounted objective, which is permitted because $b^*(\bar{x})$ depends only on \bar{x} and not on x .

This probability measure does not represent the beliefs of the decision maker, but rather a *mathematical construct* that is used as a device to support robust decision-making and the corresponding marginal valuation. The outcome of the minimization problem given in Eq. 11 provides the ingredients to construct this probability measure. We refer to this as an *uncertainty-adjusted* probability measure. The decomposition described in Section 2.4 applies here as well, but with this change in the probability measure.

Remark 3.1: Although we are not claiming that the state dynamics emerge from an optimal control problem, this development does extend to robust control problems. Here, we appeal to results

from ref. 34 on two-player stochastic differential games and the existence of a corresponding dynamic programming principle.## Importantly, we do not impose the simplification from Envelope Theorem counterpart to the maximization in supporting the resulting representation. Instead, we assume that μ and σ in the state dynamics include the recursive, robustly optimal solution for the maximizing controls.

Remark 3.2: The analysis described here can be extended to provide a robust interpretation of a well-known family of recursive utility models initiated by Kreps and Porteus (36) and Epstein and Zin (37). They are constructed from a homogeneous-of-degree one continuation value recursion including a risk adjustment to the continuation value of consumption. Our analysis is directly applicable to the case with a unitary elasticity of consumption. In addition, one may impose a unitary risk aversion in conjunction with a preference for robustness, as we describe in this subsection, to reinterpret the risk aversion presumed in the recursive utility literature. This alternative interpretation of risk aversion has an antecedent in the link in control theory between risk sensitive and robust control, first suggested by Jacobson (30). This reinterpretation carries over to recursive utility specifications in which the elasticity of intertemporal substitution differs from one, but not starting from a discounted expected utility function. Nevertheless, our approach extends in a straightforward manner to this case by altering δ to be state dependent in the FK equation for DV . See *SI Appendix, section 2* for an elaboration.

3.2. Adding Structure to the Analysis. Building on an approach suggested by Chen and Epstein (38) and Hansen and Sargent (13), we show how to add more structure to the exploration of potential misspecification.

Consider the following relative entropy contribution to the objective Eq. 11:

$$\frac{\delta}{2} E \left(\int_0^\infty \exp(-\delta t) M_t H_t \cdot H_t dt \mid X_0 = x \right), \quad [22]$$

where the process: $\{M_t : t \geq 0\}$ evolves as in Eq. 10, and where we set $\xi = \delta$ as a convenient normalization. We refer contribution Eq. 22 to the objective as exponentially weighted, relative entropy. The limiting version of this as δ goes to zero typically gives a measure of relative entropy pertinent for the continuous-time Donsker-Varadhan Large Deviation theory.

For a prespecified $H_t = \hat{h}(X_t)$, exponentially weighted discounted relative entropy could be computed by solving the FK equation:

$$0 = \frac{\delta}{2} \hat{h}' \hat{h} - \delta K + \mathbb{A}K + DK \cdot (\sigma \hat{h}). \quad [23]$$

Instead of an equation that determines the function, K , we use K to constrain potential drift distortions. We require that there be at least one \hat{h} that solves Eq. 23, but there will typically be many such choices. We expand the possibilities by converting this equation into an inequality and construct the constraint set:

$$\mathcal{S}(x) = \left\{ s : \frac{1}{2} s' s + \frac{1}{\delta} DK(x) \cdot [\sigma(x)s] \leq K(x) - \frac{1}{\delta} \mathbb{A}K(x) \right\},$$

Thus, when considering alternative models, we constrain the time derivative of discounted relative entropy. Structure is added to

See ref. 35 for a further discussion.

the robust adjustment by using the function K to represent the constraint set. For a given state realization, x , the set of resulting $s(x)$'s is an ellipsoid centered at $-\sigma' DK/\delta$ and computationally easy to impose.

With this extra structure, this robust valuation is reflected in the HJB equation:

$$\begin{aligned} 0 &= \delta U - \delta V + \mathbb{A}V + \min_{s \in \mathcal{S}(x)} DV \cdot (\sigma s) \\ &= \delta U - \delta V + \mathbb{A}V + DV \cdot (\sigma s^*), \end{aligned} \quad [24]$$

where s^* solves the minimization problem. At this juncture, we may imitate the less structured robust control approach developed in Section 3.1 by setting

$$b^*(x) = s^*(x),$$

and using the approach in Section 3.1.2 to construct an alternative probability measure that adjusts the marginal valuations for the model uncertainties. There are minor changes in the value function constructions because we are now solving a constraint problem and not penalization problem. But the underlying argument still applies. The decomposition of Subsection 2.4 continues to hold as well. See *SI Appendix, section 4* for a formal elaboration. As is described there, we make some additional restrictions when establishing the counterpart to the approach in Section 3.1.2.

Since the constraint set $\mathcal{S}(x)$ changes as an explicit function of the state, it “separates” over different calendar dates. As a consequence, this is a special case of a more general approach proposed by Chen and Epstein (38). The robust adjustment to marginal valuation applies to other types of state-dependent constraints as well. As in the robustness formulation described in Section 3.1, the resulting separation over time allows the underlying preferences of decision makers to be dynamically consistent. As we have seen, the example here has a direct connection to the approach in the previous subsection through the use of relative entropy or Kullback-Liebler divergences.

3.3. Adding Structure via Learning or Filtering. We describe another approach to add structure based on learning or filtering. We work with a Markov model with a hidden state vector. We write this system as

$$dY_t = v(Y_t, Z_t)dt + \zeta(Y_t)dW_t,$$

where the process Y is observed by the decision maker, the process Z is hidden and the matrix $\zeta\zeta'$ is nonsingular (uniformly elliptic). We assume that W is Brownian motion adapted to the filtration \mathfrak{G} generated by the composite process (Y, Z) and we denote by \mathfrak{F} the smallest filtration generated by the Y process of observable states which is pertinent to the decision maker.

The hidden state variable Z adds structure to the robustness adjustments by solving a filtering or learning problem with drift:

$$\hat{v}_t = \mathbb{E}[v(Y_t, Z_t) | \mathfrak{F}_t]$$

and “innovation” process \hat{W} , which is the multivariate Brownian motion adapted to the \mathfrak{F} filtration given by

$$\begin{aligned} d\hat{W}_t &= [\zeta(Y_t)\zeta(Y_t)']^{-1/2} [dY_t - \hat{v}_t dt] \\ &= [\zeta(Y_t)\zeta(Y_t)']^{-1/2} ([v(Y_t, Z_t) - \hat{v}_t] dt + \zeta(Y_t)dW_t). \end{aligned}$$

3.3.1. Simple illustration. We consider the case of a time-invariant hidden state that takes on a finite number of values to illustrate the robustness adjustment. We may think of this as the case of a finite number of possible models, equivalently, a parameter vector that takes on a finite number of values.

Without loss of generality, we may think of realizations of Z as coordinate vectors with dimension I . Moreover, we may write:

$$v(y, z) = \kappa(y)z$$

Note that:

$$\hat{Z}_t \stackrel{\text{def}}{=} \mathbb{E}(Z_t | \mathfrak{F}_t)$$

is the vector of conditional probabilities for a set of I possible models. By essentially applying a conditional version of least squares, we obtain the following recursive learning equation:

$$d\hat{Z}_t = [\text{diag}(\hat{Z}_t) - \hat{Z}_t\hat{Z}_t']\kappa(Y_t)'[\zeta(Y_t)\zeta(Y_t)']^{-1}[dY_t - \kappa(Y_t)\hat{Z}_t dt],$$

where $\text{diag}(\hat{z})$ is the diagonal matrix with the entries of \hat{z} on the diagonal. The state vector for the decision maker now becomes

$$X_t = \begin{bmatrix} Y_t \\ \hat{Z}_t \end{bmatrix},$$

where:

$$dY_t = \kappa(Y_t)\hat{Z}_t dt + [\zeta(Y_t)\zeta(Y_t)']^{1/2} d\hat{W}_t$$

and the constructed Brownian increment $d\hat{W}_t$ under the \mathfrak{F}_t satisfies:

$$d\hat{W}_t = [\zeta(Y_t)\zeta(Y_t)']^{-1/2} [\kappa(Y_t)Z_t dt - \kappa(Y_t)\hat{Z}_t dt + \zeta(Y_t)dW_t].$$

Building on the smooth ambiguity decision model of refs. 39–41, we next use robustness considerations in conjunction with variational utility as developed in ref. 29. The value function $V = V(x)$ solves the HJB equation:

$$0 = -\delta V(x) + \delta U(y) + \mathbb{A}V(x) + \mathbb{M}(DV, x),$$

where \mathbb{A} is the infinitesimal generator for composite X process and

$$\begin{aligned} \mathbb{M}(p, x) &\stackrel{\text{def}}{=} \min_{\pi_i \geq 0, \sum \pi_i = 1} \left[p' \sigma(y) [\zeta(y)\zeta(y)']^{-1/2} \kappa(y) (\pi - \hat{z}) \right. \\ &\quad \left. + \xi \sum_i (\log \pi_i - \log \hat{z}_i) \pi_i \right]. \end{aligned} \quad [25]$$

Conveniently, the minimization problem on the right side of Eq. 25 has a quasi-analytical solution, which we report in *SI Appendix, section 5.1*.

To relate this to our previous analysis, we introduce the drift distortion:

$$h(x, \pi) = [\zeta(y)\zeta(y)']^{-1/2} \kappa(y) (\pi - \hat{z}),$$

parameterized by π with entries that sum to one, and view:

$$\xi \sum_i (\log \pi_i - \log \hat{z}_i) \pi_i \quad [26]$$

as divergence measure in π . This perspective allows us to imitate the approach in Section 3.1.2 to construct an alternative probability measure that adjusts the marginal valuations for the model uncertainties. The decomposition outlined in Subsection 2.4 continues to hold as well.

3.3.2. Smooth ambiguity and misspecification. Learning is absent from the previous two formulations with the view that future misspecifications may not be reflected in past evidence. The perspective in this subsection is different. For example, the particular featured example incorporates a form of learning through the use of realizations \hat{z} of \widehat{Z}_t in the divergence criterion Eq. 26. The decision maker uses the baseline probability specification to solve the parameter learning problem and then entertains a restricted form of a misspecification again captured by a drift distortion along with corresponding divergence penalty.

While this example has considerable pedagogical simplicity, it has direct extensions to accommodate a more general parameter space or to allow for time variation as a discrete-state Markov process, as illustrated in *SI Appendix, section 5.2*. It also has extensions to other hidden-state specifications, including both Kalman-Bucy filters and the Zakai equations.

Remark 3.3: The smooth ambiguity model takes the log – exp contribution to $\mathbb{M}(DV, x)$ as a direct adjustment for ambiguity aversion expressed as a certainty equivalent and implies a robust adjustment to the learning solution yielding an uncertainty adjustment to the probabilities that is pertinent for marginal valuation. (See ref. 42 for a complementary approach to linking smooth ambiguity to robustness in statistics.)

Other forms of structured uncertainty with analogous implications are also possible depending upon the application. Moreover, structured and unstructured valuation adjustments can be combined, as suggested by Hansen and Sargent (13) to form a composite uncertainty-adjusted probability measure.

4. Infrequent Jumps with Endogenous Intensities

We next consider jump processes. We are particularly interested in jumps that occur infrequently and with intensities that depend on potentially endogenous state variables. These jumps could be tipping points in an economic, social, or political system; or they could be dramatic successes in technological advancement induced by research and development that are highly uncertain with payoffs that could occur in many years or decades in the future.

We take the view that the jumps are terminal dates for diffusion specification with payoffs, should a jump occur, that depend on postjump continuation values. In applications, these postjump continuation values may themselves be deduced endogenously; but we take them as input into the prejump analysis. In summary, we adopt two conventions:

- a) We focus on representing values in a current jump state, taking continuation value functions for other states as input into our analysis.
- b) We use a baseline probability without jumps and incorporate jump probabilities as a form of state-dependent discounting in the future.

Although focusing on the “current jump state,” we could repeat the analysis for each possible jump state. The second convention allows us to preserve our previous analysis except that we modify the marginal utility contribution DU based on the marginal jump possibilities.

Example 4.1: We revisit the climate–economics example in *SI Appendix, section 6*. The example model specification has two

types of Poisson jumps, a damage curve realization jump, for which there are $L - 1$ possibilities, and technology jump that eliminates the need for carbon emissions. The $L - 1$ damage realization curves are treated as equally likely under baseline probabilities, given that a damage curve realization jump takes place. All L jumps have intensities that depend on endogenous state variables: either temperature or the stock of knowledge.

4.1. Representing Marginal Valuations Abstracting from Robustness. For simplicity, we consider a finite number of state-dependent jump intensities, $\mathcal{J}^\ell(x)$ for $\ell = 1, \dots, L$. Associated with each ℓ is a postjump value function, V^ℓ . This specification nests as a special case one for which there is a single jump but a distribution over jump locations. We treat the probability conditioned on a jump times the jump intensity as an intensity for each of the possible jump locations. The FK-equation for the jump-diffusion model is

$$0 = \delta U - \delta V + \mathbb{A}V + \sum_{\ell=1}^L \mathcal{J}^\ell (V^\ell - V). \quad [27]$$

Arguing as before, we find that $DV \cdot \lambda$ solves the equation:

$$\begin{aligned} 0 = \delta DU \cdot \lambda - (\delta + \sum_{\ell=1}^L \mathcal{J}^\ell(x)) DV \cdot \lambda \\ + \sum_{\ell=1}^L \frac{\partial \mathcal{J}^\ell}{\partial x}(x) \cdot \lambda [V^\ell(x) - V(x)] \\ + \sum_{\ell=1}^L \mathcal{J}^\ell(x) \frac{\partial V^\ell}{\partial x}(x) \cdot \lambda, \end{aligned}$$

and admits the representation:

$$DV(x) \cdot \lambda = \mathbb{E} \left[\int_0^\infty \Delta_t \Phi_t \mid X_0 = x, \Lambda_0 = \lambda \right], \quad [28]$$

with the expectation in Eq. 28 computed with diffusion dynamics.

In Eq. 28,

$$\Delta_t \stackrel{\text{def}}{=} \exp \left(- \int_0^t \left[\delta + \sum_{\ell=1}^L \mathcal{J}^\ell(X_u) \right] du \right),$$

is the modification of the discount factor process in a potentially state-dependent way to adjust for the jump probabilities. The stochastic flow term Φ_t has three components:

$$\Phi_t = \Phi_t^1 + \Phi_t^2 + \Phi_t^3$$

where, with as before, $\Lambda_t = \frac{\partial X_t}{\partial x} \lambda$,

$$\Phi_t^1 \stackrel{\text{def}}{=} \delta \Lambda_t \cdot \frac{\partial U}{\partial x}(X_t),$$

$$\Phi_t^2 \stackrel{\text{def}}{=} \Lambda_t \cdot \sum_{\ell=1}^L \frac{\partial \mathcal{J}^\ell}{\partial x}(X_t) [V^\ell(X_t) - V(X_t)],$$

$$\text{and } \Phi_t^3 \stackrel{\text{def}}{=} \sum_{\ell=1}^L \mathcal{J}^\ell(X_t) \Lambda_t \cdot \frac{\partial V^\ell}{\partial x}(X_t).$$

The expression for Φ_t^1 captures the utility potential, the one for Φ_t^2 the marginal impact of a jump, and the one for Φ_t^3 the marginal valuation should a jump occur. The flow contribution Φ_t^2 comes into play when the jump intensities depend on endogenous state variables. While Φ_t^1 measures the contributions of the current period to the valuation through the contribution of the flow utility to preferences, Φ_t^3 quantifies the future contributions reflected in the postjump value functions.

4.2. Additive Value Decomposition. The additive flow contribution implies a corresponding additive value decomposition:

$$DV(x) \cdot \lambda = [V^1(x) + V^2(x) + V^3(x)] \cdot \lambda,$$

where:

$$V^j(x) \cdot \lambda \stackrel{\text{def}}{=} \mathbb{E} \left[\int_0^\infty \Delta_t \Phi_t^j \mid X_0 = x, \Lambda_0 = \lambda \right].$$

Again, the flow terms can be decomposed by alternative state contributions, which can be particularly revealing in dynamic stochastic systems with state interactions. The two new jump contributions to the stochastic flow process $\{\Phi_t : t \geq 0\}$ can also be decomposed additively by jump types ℓ .

Remark 4.2: In applications for which the continuation value functions V^ℓ are computed by solving equations with the same mathematical structure, we can obtain a corresponding representation for the partial derivatives of the continuation value functions.

4.3. Robust Adjustment to the Jump Intensities. Following ref. 43, we add robustness to the potential misspecification of the jump intensities by entertaining alternative scalings, g^ℓ , of these intensities where $g^\ell \geq 0$ alters the intensity of type ℓ . Given the additive structure by which the jump intensities enter Eq. 27, we deduce the robust adjustment for each intensity separately by solving, for each $\ell = 1, 2, \dots, L$, the minimization problem:

$$\min_{g^\ell} \mathcal{J}^\ell \left[g^\ell (V^\ell - V) + \xi (1 - g^\ell + g^\ell \log g^\ell) \right]. \quad [29]$$

In Eq. 29, the term:

$$\mathcal{J}^\ell (1 - g^\ell + g^\ell \log g^\ell)$$

measures the conditional relative entropy of the jump intensity change induced by the scaling. The minimizing g^ℓ is

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

and the minimized objective given by

$$\mathcal{J}^\ell \xi \left(1 - \exp \left[-\frac{1}{\xi} (V^\ell - V) \right] \right). \quad [30]$$

The minimized objective is increasing and concave in the difference in the value function $V^\ell - V$.

In what follows, we are interested in the derivative of the minimized function in Eq. 30 with respect to the state vector given by

$$\mathcal{J}^\ell g^{\ell*} (DV^\ell - DV),$$

where $g^{\ell*}$ is the minimizer used to alter the jump intensity. With these computations, we are led to modify both the discounting process Δ and the stochastic flow process Φ . The modified Δ process, which now is

$$\Delta_t \stackrel{\text{def}}{=} \exp \left(-\int_0^t \left[\delta + \sum_{\ell=1}^L g^{*\ell}(X_u) \mathcal{J}^\ell(X_u) \right] du \right),$$

accounts for the altered jump probabilities. The second two contributions to the flow process become the following:

$$\Phi_t^2 \stackrel{\text{def}}{=} \xi \Lambda_t \cdot \sum_{\ell=1}^L D \mathcal{J}^\ell(X_t) \left[1 - \exp \left(-\frac{1}{\xi} [V^\ell(X_t) - V(X_t)] \right) \right]$$

$$\text{and } \Phi_t^3 \stackrel{\text{def}}{=} \sum_{\ell=1}^L \mathcal{J}^\ell(X_t) g^{*\ell}(X_t) \Lambda_t \cdot DV^\ell.$$

Using analogous reasoning as in our previous investigation of robustness, since the $g^{*\ell}$'s are solved by minimization, we ignore their partial derivatives when differentiating the intensity with respect to the states. With these changes to Δ and Φ^2 , and Φ^3 , we obtain the robust counterpart decomposition described in Subsection 4.2.

Remark 4.3: We may also include robustness considerations to the diffusion specification as described in Section 3.1 in the formulation here.

5. Summary

Our representation of intertemporal marginal valuations has four fundamental components:

- i) stochastic responses that are the marginal responses of future variables of interest to a change in an initial time period;
- ii) an uncertainty-adjusted probability measure that accounts for model ambiguity and misspecification concerns;
- iii) discounting that incorporates prospective probabilities of big events or large uncertain changes in the future;
- iv) marginal implications of big events in the future, contributed by marginal jump intensities and continuation values.

Asset-pricing-type representations support two decompositions of marginal valuation. The first one reflects the state interdependencies; and the second one reflects the three types of stochastic flows, two of which are induced by prospective Poisson jumps with state-dependent intensities.

6. Constructing Prudent Policies

We see multiple uses of the conceptual framework that we have delineated in this paper.

First, we have shown how to operationalize approaches to addressing “deep uncertainty” for computing marginal valuations relevant to alternative decisions. Our perspective, consistent with a variety of approaches in control theory and decision theory, is to use probabilistic formulations as an important starting point but to back off from a full commitment to this practice. Instead, we explore ways to relax this full confidence, consistent with the axiomatic formulations of decision theory. We show how this limited confidence is operationalized through the construction of uncertainty-adjusted probability distributions to use when

computing values. Some treatments for confronting deep uncertainties use multiple scenarios. Elsewhere, ref. 44 has discussed the limitations and challenges of this approach within the context of central bank climate policy. To be useful as a policy guide, scenarios need both a dynamic structure, as components of uncertainty will get resolved in the future, and some notion of plausibility when there are a vast number of alternative scenarios to consider.

Second, empirically driven approaches to social valuation are significantly more coherent and credible when accompanied by a conceptual framework within which to interpret the evidence. This paper provides such a framework. Many of the valuation measurements from environmental economics and macroeconomics address policy challenges that are explicitly dynamic and arguably confront broadly conceived uncertainties. Some empiricists may lament our dependence on formal modeling with explicit interactions across states. Although modeling such interactions in a credible and tractable way can be very challenging, understanding how they might operate seems central to the design of prudent policies in an uncertain world. On this point, we are very sympathetic to the call from ref. 45 to capture the interactions among the sources of uncertainty when exploring challenges to climate change. Modular approaches like those discussed in ref. 46 simplify cross-disciplinary inputs into marginal valuations but may be problematic if they ignore interactions among state variables.

Third, our marginal value decompositions are designed to help researchers who build explicit dynamic stochastic models to deconstruct the implications when there are many “moving parts” to consider. By design, one of them features the quantitative impact of state variable interactions over time, and the other assesses the long-term impacts of particular courses of action, such as investment in scientific advances that support technologies of the future. (For instance, see ref. 47 for an application of these decompositions to large-project investments in addressing climate change.) The decompositions are applicable to the evaluation of both private sector behavior in decentralized economies and policy consequences. They are applicable to robustly optimal policy specifications and ad hoc alternatives.

Although we offer no simple solutions, we suggest a coherent framework for measuring and assessing policy impacts, embracing a broad notion of uncertainty. Our aim is to support critical examination of current approaches and to open the door to better quantifications in the future. Part of our motivation stems from government critiques of current methods (for a recent example, see ref. 48) and calls for better approaches (for a second example, see ref. 49). Although we do not offer a detailed commentary on these documents here, our framework provides tools to confront some of the challenges they point to. A principal example is our incorporation of broadly based uncertainty adjustments within the valuations themselves. A typical approach to uncertainty quantification is to perform an uncertainty-motivated sensitivity analysis outside or *external* to the formal realm of decision making. This external perspective, in which uncertainty is assessed outside formal measurement, is in contrast to the tools we discuss, and the analysis we provide in Section 3. Our approach is meant to *internalize* at least some aspects of the uncertainty quantification. The importance of this aspect to valuation will depend on how decision makers are or should be averse to the broadly conceived uncertainties. In our view, this perspective strengthens the conceptual foundation and quantitative rigor of marginal valuations and will provide more salient guides to prudent decision-making.

Data, Materials, and Software Availability. There are no data underlying this work.

ACKNOWLEDGMENTS. Fernando Alvarez, William Brock, Dennis Epple, Benjamin Hebert, Ken Miyahara Coello, Darrell Duffie, Greg Kaplan, Kevin Murphy, Aleksei Oskolkov, Diana Petrova, Richard Romano, Tom Sargent, Grace Tsiang, Xunyyu Zhou, Bowen Dong, Jessie Liao, and Zhaoyang Xu provided helpful comments.

Author affiliations: ^aDepartment of Economics and Booth School of Business, University of Chicago, Chicago, IL 60637; and ^bDepartment of Mathematics, University of Chicago, Chicago, IL 60637

- P. Dasgupta, K. G. Mäler, Net national product, wealth, and social well-being. *Environ. Dev. Econ.* **5**, 69–93 (2001).
- Y. Cai, K. L. Judd, T. S. Lontzek, “The social cost of carbon with climate risk. Economic working paper (2017). <https://www.hoover.org/sites/default/files/research/docs/18113-judd1.pdf>.
- W. D. Nordhaus, Revisiting the social cost of carbon. *Proc. Natl. Acad. Sci. U.S.A.* **114**, 1518–1523 (2017).
- K. Rennett *et al.*, Comprehensive evidence implies a higher social cost of CO₂. *Nature* **610**, 687–692 (2022).
- L. P. Hansen, T. J. Sargent, T. D. Tallarini, Robust permanent income and pricing. *Rev. Econ. Stud.* **66**, 873–907 (1999).
- F. Alvarez, U. J. Jermann, Using asset prices to measure the welfare cost of business cycles. *J. Polit. Econ.* **112**, 1223–1256 (2004).
- R. E. Lucas, *Models of Business Cycles* (Basil Blackwell, New York, 1987).
- S. Hallegatte, A. Shah, C. Brown, R. Lempert, S. Gill, “Investment decision making under deep uncertainty-application to climate change” (Tech. Rep. 6193, World Bank Policy Research Working Paper, 2012).
- H. R. Maier *et al.*, An uncertain future, deep uncertainty, scenarios, robustness and adaptation: How do they fit together? *Environ. Model. Softw.* **81**, 154–164 (2016).
- V. A. Marchau, W. E. Walker, P. J. Bloemen, S. W. Popper, *Decision Making Under Deep Uncertainty: From Theory to Practice* (Springer Nature, 2019).
- J. Rising, M. Tedesco, F. Piontek, D. A. Stainforth, The missing risks of climate change. *Nature* **610**, 643–651 (2022).
- L. Berger, M. Marinacci, Model uncertainty in climate change economics: A review and proposed framework for future research. *Environ. Resour. Econ.* **77**, 475–501 (2020).
- L. P. Hansen, T. J. Sargent, Structured ambiguity and model misspecification. *J. Econ. Theory* **199**, 105165 (2022).
- S. Cerreia-Vioglio, L. P. Hansen, F. Maccheroni, M. Marinacci, Making decisions under model misspecification. *Rev. Econ. Stud.* forthcoming **2**, rda046 (2025).
- J. Yang, H. J. Kushner, A Monte Carlo method for sensitivity analysis and parametric optimization of nonlinear stochastic systems. *SIAM J. Control. Optim.* **29**, 1216–1249 (1991).
- E. Gobet, R. Munos, Sensitivity analysis using Itô-malliavin calculus and martingales, and application to stochastic optimal control. *SIAM J. Control. Optim.* **43**, 1676–1713 (2005).
- R. Frisch, “Propagation problems and impulse problems in dynamic economics” in *Economic Essays in Honour of Gustav Cassel*, J. Åkerman, Ed. (Allen and Unwin, 1933), pp. 171–205.
- E. Fournie, J. M. Lasry, J. Lebuchoux, P. L. Lions, N. Touzi, Applications of Malliavin calculus to Monte Carlo methods in finance. *Finance Stochast.* **3**, 391–413 (1999).
- J. Borovička, L. P. Hansen, J. A. Scheinkman, Shock elasticities and impulse responses. *Math. Financial Econ.* **8**, w20104 (2014).
- C. A. Sims, Macroeconomics and reality. *Econometrica* **48**, 1–48 (1980).
- R. Dorfman, An economic interpretation of optimal control theory. *Am. Econ. Rev.* **59**, 817–831 (1969).
- J. M. Bismut, Growth and optimal intertemporal allocation of risks. *J. Econ. Theory* **10**, 239–257 (1975).
- S. Peng, Y. Shi, Infinite horizon forward-backward stochastic differential equations. *Stoch. Process. Appl.* **85**, 75–92 (2000).
- H. E. Ryder, G. M. Heal, Optimal growth with intertemporally dependent preferences. *Rev. Econ. Stud.* **40**, 1–31 (1973).
- M. Rubinstein, The valuation of uncertain income streams and the pricing of options. *Bell J. Econ.* **7**, 407–425 (1976).
- S. Ross, A simple approach to valuation of risky streams. *J. Bus.* **51**, 453–475 (1978).
- L. P. Hansen, S. F. Richard, The role of conditioning information in deducing testable restrictions implied by dynamic asset pricing models. *Econometrica* **50**, 587–614 (1987).
- F. H. Knight, *Risk, Uncertainty, and Profit* (Houghton Mifflin, 1921).
- F. Maccheroni, M. Marinacci, A. Rustichini, Dynamic variational preferences. *J. Econ. Theory* **128**, 4–44 (2006).
- D. H. Jacobson, Optimal stochastic linear systems with exponential performance criteria and their relation to deterministic differential games. *IEEE Trans. Autom. Control.* **AC-18**, 124–131 (1973).
- M. R. James, Asymptotic analysis of nonlinear stochastic risk-sensitive control and differential games. *Math. Control. Signals Syst.* **5**, 401–417 (1992).
- L. P. Hansen, T. J. Sargent, Robust control and model uncertainty. *Am. Econ. Rev.* **91**, 60–66 (2001).

33. L. Ljungqvist, T. J. Sargent, *Recursive Macroeconomic Theory* (MIT Press, Cambridge, MA, ed. 4, 2018).
34. W. H. Fleming, P. E. Souganidis, On the existence of value function of two-player, zero-sum stochastic differential games. *Indiana Univ. Math. J.* **38**, 293–314 (1989).
35. L. P. Hansen, T. J. Sargent, G. A. Turmuhambetova, N. Williams, Robust control and model misspecification. *J. Econ. Theory* **128**, 45–90 (2006).
36. D. M. Kreps, E. L. Porteus, Temporal resolution of uncertainty and dynamic choice. *Econometrica* **46**, 185–200 (1978).
37. L. G. Epstein, S. E. Zin, Substitution, risk aversion and the temporal behavior of consumption and asset returns: A theoretical framework. *Econometrica* **57**, 937–969 (1989).
38. Z. Chen, L. Epstein, Ambiguity, risk, and asset returns in continuous time. *Econometrica* **70**, 1403–1443 (2002).
39. P. Klibanoff, M. Marinacci, S. Mukerji, Recursive smooth ambiguity preferences. *J. Econ. Theory* **144**, 930–976 (2009).
40. L. P. Hansen, T. J. Sargent, Robustness and ambiguity in continuous time. *J. Econ. Theory* **146**, 1195–1223 (2011).
41. L. P. Hansen, J. Miao, Aversion to ambiguity and model misspecification in dynamic stochastic environments. *Proc. Natl. Acad. Sci. U.S.A.* **115**, 9163–9168 (2018).
42. S. Cerreia-Vioglio, F. Maccheroni, M. Marinacci, L. Montrucchio, Ambiguity and robust statistics. *J. Econ. Theory* **148**, 974–1049 (2013).
43. E. W. Anderson, L. P. Hansen, T. J. Sargent, A quartet of semigroups for model specification, robustness, prices of risk, and model detection. *J. Eur. Econ. Assoc.* **1**, 68–123 (2003).
44. L. P. Hansen, Central banking challenges posed by uncertain climate change and natural disasters. *J. Monet. Econ.* **125**, 1–15 (2022).
45. N. P. Simpson *et al.*, A framework for complex climate change risk assessment. *One Earth* **4**, 489–501 (2021).
46. National Academies of Sciences, *Engineering and Medicine, Valuing Climate Damages: Updating Estimation of the Social Cost of Carbon Dioxide* (The National Academies Press, Washington, DC, 2017).
47. M. Barnett, W. A. Brock, L. P. Hansen, H. Zhang, Uncertainty, social valuation, and climate change policy (University of Chicago, Becker Friedman Institute for Economics Working Paper number 2024-75, July 3, 2025). <http://dx.doi.org/10.2139/ssrn.4872679>.
48. J. Christy, J. Curry, S. Koonin, R. McKittrick, R. Spencer, *A Critical Review of Impacts of Greenhouse Gas Emissions on the US Climate* (U.S. Department of Energy, Washington DC, Report, 2025).
49. Office of Management and Budget, "OMB circular no. a-4, regulatory analysis" (Tech. Rep. A-4, Executive Office of the President, 2023).