Robust Policy Assessment

and

Intertemporal Valuation

Lars Peter Hansen (University of Chicago)

Northwestern Economics-Kellogg Summer School in
Economic Theory Workshop 2026

August 27, 2026

Collaborators: Barnett (ASU), Brock (University of Wisconsin),
Zhang (Argonne), Souganidis (University of Chicago),
Cerreia-Vioglio, Maccheroni, Marinacci (Bocconi),
Sargent (University of Pennsylvania)

Economists' wisdom

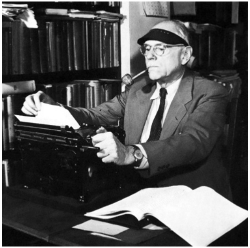

"uncertainty must be taken in a sense radically distinct from the familiar notion of risk, from which it has never been properly separated.... and there are far-reaching and crucial differences in the bearings of the phenomena ...."

Knight (1921)

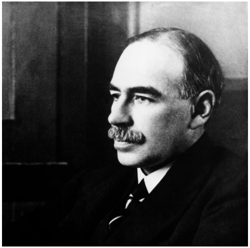

"We have, as a rule, only the vaguest idea of any but the most direct consequences of our acts... Our knowledge of the future is fluctuating, vague, and uncertain."

Keynes (1937)

Tension

- limited understanding of the mechanism by which policy influences economic outcomes
- demand for precise answers from the public and policymakers

Evidence-based policy?

- The challenges of economic policy are often dynamic in nature.
- Empirical researchers and applied statisticians gravitate toward problems with data richness.
- But the data tend not to be rich along all of the relevant directions for policy.
- Model-based approaches remain essential with important subjective inputs needed to fill knowledge gaps.

Statisticians’ wisdom

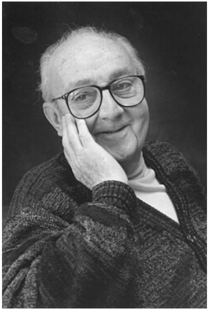

“Since all models are wrong, the scientist must be alert to what is importantly wrong. It is inappropriate to be concerned about mice when there are tigers abroad.” Box (1976)

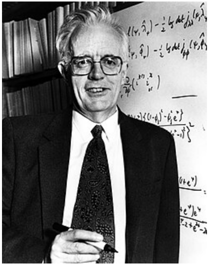

“The idea that complex physical, biological or sociological systems can be exactly described by a few formulae is patently absurd. The construction of idealized representations that capture important stable aspects of such systems is, however, a vital part of general scientific analysis and statistical models, especially substantive ones ...” Cox (1995)

Uncertainty tradeoffs

- How much weight do we assign to:
  - best guesses
  - potentially bad outcomes

  when making decisions?
- Do we act now, or do we wait until we learn more?

Outline

- decision theory support
- social and market valuation
- application: climate change and R&D investment

Part I: Advances in Decision Theory Under Uncertainty

Decision theory aims to develop and justify approaches that are "rational," or more appropriately, "prudent," offering a broad perspective on uncertainty.

Do not know:

- outcomes but probabilities are known - risk
- which among multiple probability models is best - ambiguity (prior uncertainty)
- ways in which a model might give flawed probabilistic predictions - misspecification (can include likelihood uncertainty)

We endow our policy maker with preferences that encode these uncertainty aversions. We then explore sensitivities to these aversions.

How Do We Use Decision Theory

Provides:

- entertains what is sometimes called deep uncertainty in the policy arena
- tractable ways to capture alternative uncertainty components in dynamic environments
- a vehicle to add more structure to uncertainty quantification

Approach

- draw on and develop modifications of Savage-style axiomatic formulations from decision theory to extend notions of uncertainty beyond risk in ways that make contact with applied challenges in economics and other disciplines
- distinguish concerns about potential misspecification of model-implied probabilities from concerns about the misspecification of prior probabilities over models

This opens the door to better ways to carry out uncertainty quantification for the dynamic, stochastic economic models used for private sector planning and governmental policy assessment.

Recent references distinguishing the two forms of misspecification

- Hansen, L. P., & Sargent, T. J. (2022). Structured ambiguity and model misspecification. Journal of Economic Theory.
- Hansen, L. P., & Sargent, T. J. (2024). Risk, ambiguity, and misspecification: Decision theory, robust control, and statistics. Journal of Applied Econometrics.

Aim

Develop computationally tractable methods for exploring subjective uncertainty including potential model misspecification and ambiguity across models.

Motivation:

- assess the impact of uncertainty on prudent policy outcomes
- isolate the forms of uncertainty that are most consequential for these outcomes.

Inputs

- tools from probability and statistics to limit the type and amount of uncertainty that is entertained
- aversion to uncertainty about probabilities over future events

Anscombe-Aumann (AA)

- preferences defined over acts
- an act maps states $\rightarrow$ probabilities over outcomes

AA refers to a horse race versus a lottery for motivation.

- the probability distribution over outcomes is the lottery
- the probability over states is the horse race

In a static setting, we assume:

- a state is a parameter vector that indexes a statistical model
- each statistical model induces a probability distribution over outcomes
- a probability distribution over "states" is a prior distribution

Horse race and lottery uncertainty

Figure: ambiguity

Figure: risk

Static decision theory

Consider a parameterized model of a random vector with realization w:

$$\ell(w \mid \theta) d\tau_o(w)$$

where:

$$\int_W \ell(w \mid \theta) d\tau_o(w) = 1,$$

$\theta \in \Theta$ and Θ is a parameter space, and W is the space of possible realizations of w. Place a baseline prior distribution π_o over Θ and consider a 'rule' $\gamma(w)$.

By analogy, think of

- θ as the outcome of "horse race"
- the distribution induced by $\gamma(w)$ conditioned on θ as the "lottery"

Observations

- parameter space Θ can be infinite dimensional
- a prize rule, γ, is within a restricted set of functions. For instance,
  - data $y = \zeta(w)$
  - prize rule:
$$\gamma_\delta(w) = \Psi[\delta \circ \zeta(w), w]$$

The paper on which this talk is based shows the similarities and differences with machine-learning methods, specifically with PAC (Probably Approximately Correct) Bayesian methods.

Subjective expected utility

Order preferences over γ by

$$\int_{\Theta}\left[\int_{W} u[\gamma(w)]\ell(w \mid \theta)d\tau_o(w)\right]d\pi_o(\theta).$$

Supported by Savage and AA axioms.

Ambiguity?

if I knew of a good way to make a mathematical model of those phenomena [vagueness and indecision], I would adopt it, but I despair of finding one. One of the consequences of vagueness is that, ... we are able to elicit precise probabilities by self-interrogation in some situations but not in others.

Personal communication from Leonard "Jimmie" Savage to Karl Popper in 1958

Divergences

Use a convex function ϕ_p for constructing divergence between probability measures. Restrict $\phi_p(1) = 0$ and $\phi_p''(1) = 1$ (normalization).

- Consider alternative priors of the form $d\pi(\theta) = n(\theta)d\pi_o(\theta)$ for $n \geq 0$ satisfying:

$$\int_\Theta n(\theta)d\pi_o(\theta) = 1.$$

  Call this collection $\mathcal{N}$.
- For priors indexed by $n \in \mathcal{N}$, form

$$\int_\Theta \phi_p[n(\theta)]d\pi_o(\theta) \geq 0.$$

We often use relative entropy or Kullback-Leibler divergence: $\phi_p(n) = n \log n$.

Ambiguity aversion preferences

- variational preferences (Maccheroni, Marinacci, and Rustichini, 2006)

$$\min_{n \in \mathcal{N}} \int_{\Theta} \left(\int_{W} u[\gamma(w)] \ell(w \mid \theta) d\tau_o(w) \right) n(\theta) d\pi_o(\theta) + \xi_p \int_{\Theta} \phi_p[n(\theta)] d\pi_o(\theta)$$

  for $\xi_p > 0$ and a convex function ϕ_p such that $\phi_p(1) = 0$

  or

- max-min utility (Gilboa and Schmeidler, 1989)

$$\min_{n \in \mathcal{N}^o} \int_{\Theta} \left(\int_{W} u[\gamma(w)] \ell(w \mid \theta) d\tau_o(w) \right) n(\theta) d\pi_o(\theta)$$

  for $\mathcal{N}^o \subset \mathcal{N}$

Comments about prior divergences

- Replace the independence axiom by weaker notions of constant independence.
- Interpret some previous contributions to decision theory literature as representing a prior ambiguity.
- With relative entropy divergence, the implied preference ordering agrees with smooth ambiguity preferences but is rationalized in a fundamentally different way.

Model misspecification concerns

- replace $\ell(w \mid \theta)d\tau_o(w)$ with $m(w \mid \theta)\ell(w \mid \theta)d\tau_o(w)$ for m satisfying: $\int_W m(w \mid \theta)\ell(w \mid \theta)d\tau_o(w) = 1$, and denote the set of all such m's as $\mathcal{M}$.
- rank alternative γ's conditioned on θ by solving:

$$\min_{m \in \mathcal{M}} \int_W \left(u[\gamma(w)]m(w \mid \theta) + \xi_m \phi_m[m(w \mid \theta)]\right)\ell(w \mid \theta)d\tau_o(w)$$

  for $\xi_m > 0$ and ϕ_m is a convex function satisfying $\phi_m(1) = 0$.

Observations:

- do not impose a prior distribution over $\mathcal{M}$ (conditioned on θ).
- links to parts of robust control theory
- allow for ambiguity in the AA lotteries, and hence we are outside the standard decision theory framework

Decision theory references address model misspecification

- Cerreia-Vioglio, S., Hansen, L. P., Maccheroni, F., & Marinacci, M. (2026). Making decisions under model misspecification. The Review of Economic Studies.
- Gilboa, I., Maccheroni, F., Marinacci, M., Schmeidler, D. (2010). Objective and subjective rationality in a multiple prior model. Econometrica.

Robust Bayes with model misspecification

Represent preferences over γ with:

$$\min_{n \in \mathcal{N}} \min_{m \in \mathcal{M}} \int_{\Theta} \left(\int_{W} u[\gamma(w)] m(w \mid \theta) \ell(w \mid \theta) d\tau_o(w) \right) n(\theta) d\pi_o(\theta) \\ + \xi_m \int_{\Theta} \left(\int_{W} \phi_m[m(w \mid \theta)] \ell(w \mid \theta) d\tau_o(w) \right) n(\theta) d\pi_o(\theta) \\ + \xi_p \int_{\Theta} \phi_p[n(\theta)] d\pi_o(\theta)$$

Joint divergence over (m, n). The penalty parameters ξ_m and ξ_p could be multipliers on constraints.

Dynamics

Use conditional counterparts to the previous analysis

- explore the consequences of misspecifying Markov transition dynamics by representing potential changes in probabilities as nonnegative martingales
- explore consequences of misspecifying priors/posteriors over alternative parameters
- address dynamic consistency
  - recursive construction of possible conditional probabilities over parameterized models
  - recursive construction of statistical divergences and their set counterparts

Part II: Intertemporal marginal valuations

- formulation
- representation
- deconstruction

Why intertemporal marginal valuations?

They:

- are used in a variety of settings: public economics, environmental economics, and macroeconomics for policy assessment;
- help to navigate movements from suboptimal allocations to more efficient alternatives;
- provide interpretations of first-order conditions for (robustly) optimal policies, including dynamic investment choice;
- support the implementation of a (robustly) optimal policy;
- are proportional to derivatives of values with respect to endogenous state variables.

I provide formulas for intertemporal marginal valuation that deconstruct the valuations into alternative interpretable components.

Examples

In addition to the marginal valuation of physical, human, and intangible capital:

- social cost of carbon
- social value of research and development
- natural capital
- social capital

Confront broad-based uncertainties arising from imperfect measurement and quantification.

Ingredients

- start with a value function, V,
    - impose the controls (investments and other actions): do not have to be optimal
    - characterize explicitly the state variable interactions
- form stochastic (nonlinear) impulse responses
- deduce an asset pricing-type formula with marginal utility adjustments
- construct an altered probability measure that adjusts for uncertainty aversion
- model infrequent big changes as Poisson jumps
- deduce additive decompositions of the contributions to the marginal valuations

Stochastic State Evolution

Initially, consider diffusion dynamics:

$$dX_t = \mu(X_t)dt + \sigma(X_t)dW_t$$

where X is a Markov diffusion and W is a multivariate standard Brownian motion.

We construct a stochastic impulse response to measure the impact of a marginal change in a state.

Stochastic Responses

Use the theory of stochastic flows (Kunita) to:

- form the vector stochastic response process, Λ, that gives the marginal impact on future X's of a marginal change in one of the initial states;
- initialize the process at one of the coordinate vectors to specify the initial state of interest.

The stochastic evolution of Λ depends on X.

Characterizing Λ

- The process Λ^i evolves as:

$$d\Lambda_t^i = (\Lambda_t)' \frac{\partial \mu_i}{\partial x}(X_t)dt + (\Lambda_t)' \frac{\partial \sigma_i}{\partial x}(X_t)dW_t.$$

  where Λ^i is the i^{th} component of Λ.
- Stack all of the Λ^i's and study the joint dynamics (X, Λ).
- Linear VAR (vector autoregression) counterpart is obtained with a drift $\mu(x)$ that is linear in x and a Brownian exposure matrix σ that is constant. Λ is not stochastic in this special case.
- Λ is stochastic in general.

An Initial Pricing Formula

Represent partial derivatives as asset prices:

$$
\begin{aligned}
&\frac{\partial V}{\partial x}(X_0) \cdot \Lambda_0 \\
&= \delta \int_0^\infty \exp(-\delta t) E\left[\frac{\partial U}{\partial x}(X_t) \cdot \Lambda_t \mid X_0, \Lambda_0\right] dt.
\end{aligned}
$$

where δ is the subjective rate of discount and U is the utility contribution to the value function.

- Initializing Λ_0 at alternative coordinate vectors gives the derivatives of interest.
- $\frac{\partial U}{\partial x}(X_t)$ is marginal utility contribution in the future.
- Λ is the vector of stochastic impulse responses.

Observation: this approach is different from solving a vector system of co-states forward.

Partitioning Valuations into Alternative Components I

Recall:

$$\frac{\partial V}{\partial x}(X_0) \cdot \Lambda_0 \\ = \delta \int_0^\infty \exp(-\delta t) E \left[\frac{\partial U}{\partial x}(X_t) \cdot \Lambda_t \mid X_0, \Lambda_0 \right] dt.$$

- decomposition by horizon: changing a state today has different implications for each future marginal payoff horizon;
- decomposition by state: changing a state today impacts other states in the future - state feedback effects;

(i) decomposing the marginal cash flow:

$$\frac{\partial U}{\partial x}(X_t) \cdot \Lambda_t$$

into state-specific components; and ii) computing discounted expected values for each component.

Adding in jumps

Analyze the problem from a "pre-jump" perspective.

- Let $\mathcal{J}^\ell(x)$ denote the jump intensity to jump type ℓ, for $\ell = 1, 2, ..., L$.
- Let V^ℓ denote the state dependent continuation value for type ℓ.

Jump Contributions

- Alter the discount rate to adjust for a damage curve revelation jump or a technology discovery jump:

$$\delta + \sum_{\ell=1}^{L} \mathcal{J}^{\ell}(X_t).$$

- Additional stochastic flow contributions:
  - marginal impact of a jump:

$$\Lambda_t \cdot \sum_{\ell=1}^{L} \left[\frac{\partial \mathcal{J}^{\ell}}{\partial x}(X_t) \right] \left[V^{\ell}(X_t) - V(X_t) \right];$$

  - marginal impact when you jump:

$$\Lambda_t \cdot \sum_{\ell=1}^{L} \mathcal{J}^{\ell}(X_t) \left[\frac{\partial V^{\ell}}{\partial x}(X_t) \right].$$

Partitioning Valuations into Alternative Components II

Decompose the two additional flow contributions:

- marginal impact of a jump:

$$\Lambda_t \cdot \sum_{\ell=1}^{L} \left[\frac{\partial \mathcal{J}^\ell}{\partial x}(X_t) \right] \left[V^\ell(X_t) - V(X_t) \right];$$

- marginal impact when you jump:

$$\Lambda_t \cdot \sum_{\ell=1}^{L} \mathcal{J}^\ell(X_t) \left[\frac{\partial V^\ell}{\partial x}(X_t) \right].$$

by jump types and compute the corresponding discounted expected values.

Robust Valuation

Adjust the valuation for model ambiguity and misspecification concerns:

- Minimize over possible probability distributions subject to a penalization.
- Use the minimizing probability to provide an uncertainty adjustment pertinent for representing marginal valuations.

Corresponding robust decision theory under deep uncertainty: formulate a recursive max-min game instead of a single-agent maximization problem.

Uncertainty-adjusted Probabilities

Draw on insights from robust Bayesian theory and derivative claims pricing to:

- derive a "worst-case" probability distribution isolating where potential misspecifications are most concerning.
- this "uncertainty-adjusted probability" does not represent the policy maker's beliefs.
- nevertheless, we justify the use of the adjusted probability to represent marginal valuations.

Uncertainty Quantification and Model Deconstruction I

We use the policy problem to add structure to uncertainty quantification by:

- depicting prudent decisions as dependent on marginal valuations;
- representing marginal valuations as asset prices with uncertain social and economic payoffs;
- partitioning marginal valuations into distinct components based on alternative contributions to the payoffs.

For our climate change example, two marginal valuations will be uniquely relevant:

- social cost of climate change
- social value of research and development.

Uncertainty quantification and Model Deconstruction II

Two questions:

- How much uncertainty aversion should we impose?
  - trace through sensitivity to the choice of penalty parameters;
  - inspect the implied worst-case distributions;
  - explore the impact on the implied uncertainty-adjusted valuations and their decompositions.
- Which source of uncertainty matters the most?
  - activate the robustness concerns one source at a time;
  - compare the results to those from a decision problem with all concerns activated simultaneously;
  - compute a Shapley value decomposition.

Reference for robust marginal valuation

- Hansen, L. P., & Souganidis, P. (2025). Stochastic responses and marginal valuation. *PNAS*.

Part III

- Describe an illustrative application that includes "big project" research and development relevant to climate change.
- Extend an AK model with adjustment costs to include R&D investment and technology discovery.

Climate Change Challenge

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

Rising, Tedesco, Piontek, Stainforth, Nature, 2022

Economic Damage Challenge

> *"Experts being clearer about what economics can and cannot tell us would not resolve disagreements about climate policy. ... The full effects of climate change are unknowable, and a more constructive public discussion about climate policy will require getting more comfortable with that."*

Noah Kaufman, Atlantic, 2026

Production Inputs

- capital
- fossil-fuel-based energy

This energy source also contributes to global warming.

Production Outputs

- consumption contributes directly to the welfare
- investment in capital increases future output
- investment in R&D increases the stock of knowledge

Prudent choices for investments and energy input depend on marginal valuations (partial derivatives of a value function of a robustly optimal policy problem).

Two particular marginal valuations are relevant:

- social cost of climate change
- social value of research and development.

Climate Externality

- include a damage function that depends on temperature;
- resulting damages captured by multiplicative reduction in consumption

What are the uncertainties?

Four channels of uncertainty:

- productivity: capital investment today alters future output
- geosciences: CO_2 emissions today impact the future climate
- economics: climate change in the future alters economic opportunities and social well-being
- technology: R&D investments today may eventually lead to economically viable technologies

Two Types of Big Events

- reveal damage uncertainty for larger values of the temperature anomaly represented by alternative damage function curves - the intensity depends on the temperature anomaly;
- discover a new technology that eliminates the need for the dirty energy input - the intensity depends on a knowledge stock.

I include the first of these to explore a tradeoff between acting now or waiting until more is known about the damages for more extreme temperature changes.

Risk Components

- Brownian (normally distributed) shocks
- Poisson (big) events

- Two types of Poisson events:
  - reveal damage uncertainty represented by alternative damage function curves for larger values of warming - the intensity depends on the temperature anomaly
  - discover a new technology that eliminates the need for the dirty energy input - the intensity depends on a knowledge stock

Uncertain Probabilities

Relax full confidence (explore sensitivity to ambiguity and misspecification)

- Brownian uncertainty - replace the standard Brownian motion with the Brownian motion altered to have a state-dependent drift or local mean
- jump uncertainty - scale each of the state dependent jump intensities

Uncertainty-adjusted Probabilities

We use them to:

- isolate where potential misspecification is most concerning;
- represent marginal valuations in the presence of uncertainty aversion.

Climate Uncertainty Plot

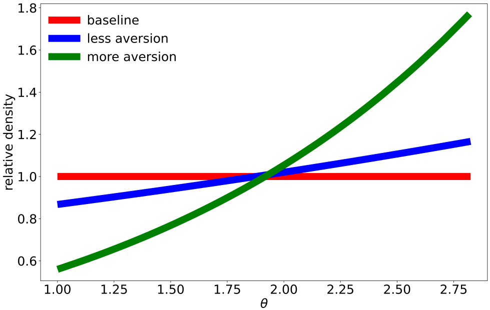

Figure: Uncertainty-adjusted densities for temperature responses to cumulative emissions. The baseline and two uncertainty-adjusted densities are relative to a uniform baseline density.

Jump Time Density

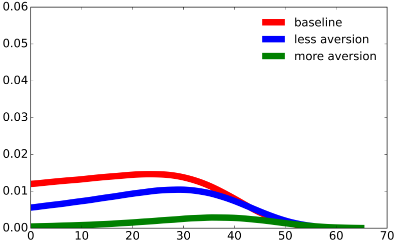

(a) Contribution of technology jump.

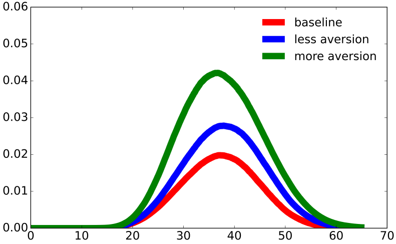

(b) Contribution of damage jumps.

Figure: Jump-time densities under the baseline probability and uncertainty-adjusted probabilities for two specifications of uncertainty aversion.

Damage Curvature Relative Densities

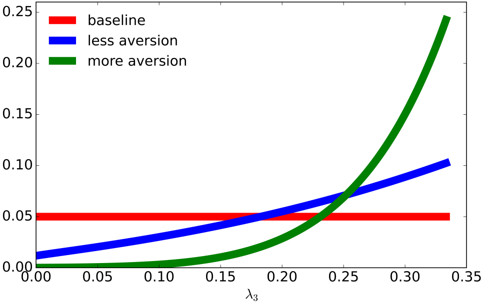

Figure: Uncertainty-adjusted probabilities for the alternative damage curvatures.

R&D investment before and after a damage realization event

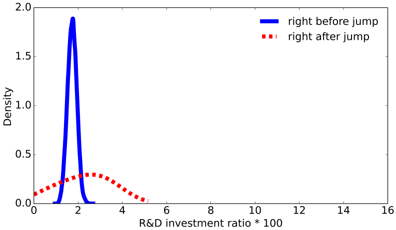

(a) baseline

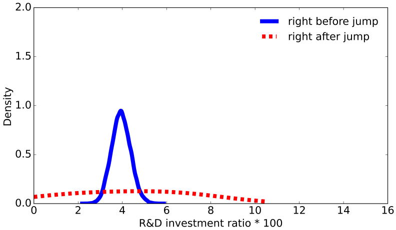

(b) less aversion

Figure: The blue solid line shows the distribution of outcomes just before a damage jump realization, and the red dashed line shows the distribution of outcomes just after a damage jump realization.

Emissions Before and after a damage realization event

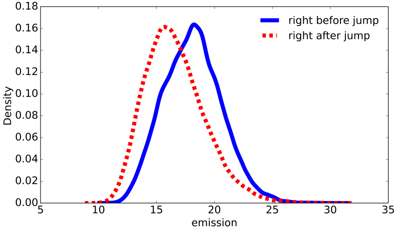

(a) baseline

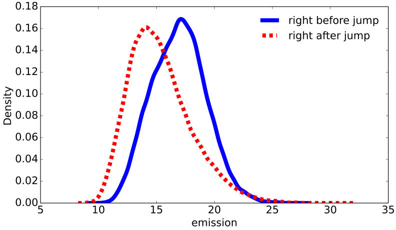

(b) less aversion

Figure: The blue solid line shows the distribution of outcomes just before a damage jump realization, and the red dashed line shows the distribution of outcomes just after a damage jump realization.

Channel-based uncertainty decompositions

Uncertainty channel	SVRD	SCGW
climate	0.49 (8%)	14.9 (8%)
damage	0.97 (16%)	38.0 (20%)
productivity	-0.25 (-4%)	-6.1 (-3%)
technology	4.94 (80%)	147.2 (76%)

Table: Shapley decompositions of the changes in the two social valuations, SVRD and SCGW, between the baseline and all-channel specifications. Entries are channel-specific contributions, and the numbers in the parentheses report percentages.

Uncertainty-adjusted Probabilities for the technology jump only specification

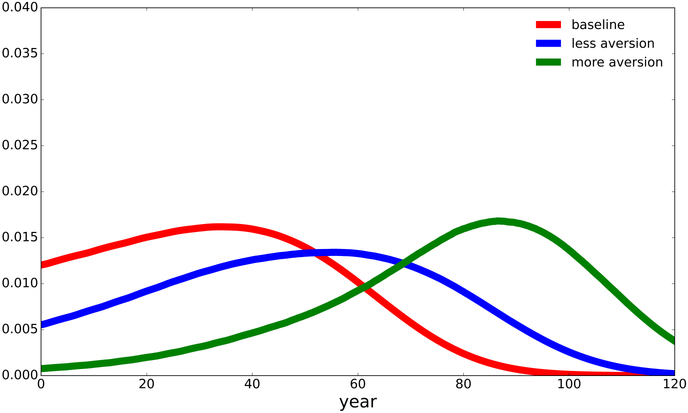

Jump-time densities for the technology jump

Social values of R&D and the corresponding investments for alternative aversions

ξ_m	∞	.10	.05	.01	.009	.008	.007
SVRD	4.27	6.42	8.70	8.02	7.68	7.30	6.92
R&D inv.	.0079	.0179	.0329	.0281	.0258	.0233	.0209

Table: Social value of R&D and the R&D investment/output ratios as functions of the robustness penalty parameter, ξ. Smaller values of ξ imply larger aversions to uncertainty.

R&D investment increases over a significant range of uncertainty aversions

Why are the R&D investment impacts not monotone in the uncertainty aversion of the planner?

Two counteracting forces:

- adopt a more cautious view of when the technological success will be realized;
- achieve a greater benefit to a successful discovery because of the substantial reduction in the exposure to uncertainty.

The second force is captured by the changes in marginal impact of the technology measured by:

$$\left[\frac{\partial \mathcal{J}^L}{\partial x}(X_t)\right]\left[V^L(X_t) - V(X_t)\right].$$

where V is the value function in the absence of technology success and V^L is the value function if a technology discovery occurs.

Reference for the application

- Barnett, M., Brock, W., Hansen, L. P., & Zhang, H. (2026). Uncertainty, social valuation, and climate change policy. *SSRN working paper*.

Work in Progress

Extending the model using mean-field control theory to include:

- cross-sectional heterogeneity in exposures to climate change and in the damage consequences;
- large population of individuals formulated as mean-field control with aggregate shocks, generalized to include robustness concerns.

Conclusions

- Sometimes, the best response to uncertainty is to be more proactive.
- More transparent communication of scientific uncertainty can lead to better policy.
- This research is part of a larger agenda to explore the impacts of uncertainty concerns for both private and public sector decision-making.