Risk, Ambiguity, and Misspecification

Decision Theory, Robust Control, and Statistics

Lars Peter Hansen (University of Chicago)

JAE lecture at the IAAE Conference, Oslo

Collaborator: Thomas J. Sargent (NYU)

June 28, 2023

Link to Paper

The 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, 2022

Haunted by Hayek's forewarning

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

From Hayek's Nobel address (1974)

Tension

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

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?

Talk outline

- decision theory
- uncertainty quantification
- market and social valuation
- application to financial markets
- application to climate change policy

Decision theory

Provides:

- ▷ axiomatic formulations of "rationality" in the presence of uncertainty, broadly conceived
- ▷ tractable ways to capture alternative uncertainty components in dynamic environments

Statisticians' wisdom

"In what circumstances is a minimax solution reasonable? I suggest that it is reasonable if and only if the least favorable initial distribution is reasonable according to your body of beliefs." Irving J. Good (1952)

"Now it would be very remarkable if any system existing in the real world could be exactly represented by any simple model. However, cunningly chosen parsimonious models often do provide remarkably useful approximations." George Box (1979)

Aims

- allow for a broad perspective on uncertainty
  - risk - unknown outcomes with known probabilities
  - ambiguity - unknown weights to assign to alternative probability models
  - misspecification - unknown ways in which a model might give flawed probabilistic predictions
- include formulations that are dynamic and recursive

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 misspecifications of likelihoods from concerns about robustness of alternative priors

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

Targets

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

Goals:

- ▷ 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 or dislike of uncertainty about probabilities over future events

Anscombe-Aumann (AA)

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

AA refer 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

ambiguity

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
- prize rules are within a restricted set of functions. For instance, $\gamma(w) = \Gamma(d, w)$ for $d \in D$.
- a direct extension that merely complicates notation is to let the prize rule depend on information about the future shock

The paper that this talk is based on shows the similarities and differences to machine-learning methods, specifically to PAC (probably almost 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 any good way to make a mathematical model of these phenomenon [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 L. J. Savage to Karl Popper in 1957

Divergences

Use a convex function ϕ_p for constructing divergence between probability measures. Each ϕ is a convex function with $\phi(1) = 0$ and $\phi''(1) = 1$ (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 $\phi_m(1) = 1$.

Observations:

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

Robust Bayes with model misspecification, I

Represent preferences over γ using:

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

Robust Bayes with model misspecification, II

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).

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

Hansen and Sargent (JET, 2022) confront a tension between dynamic consistency and admissibility.

Uncertainty quantification

Two questions:

- How much uncertainty aversion should we impose?
  - trace through sensitivity to the choice of penalty parameters or constraints
  - inspect the impact on the implied worst-case distributions from the min-max problem
- Which source of uncertainty matters the most?
  - activate the robustness concerns one source at a time
  - compare the decision outcomes to those from a decision problem with all concerns activated simultaneously

Valuation

What does asset valuation provide?

Asset pricing theory: how do markets assess the investment opportunities in the face of uncertain future net payoffs?

- ▷ "assets" include financial, physical, human, organizational and environmental "capital"
- ▷ associated with each asset is a prospective sequence of net payoffs to investments (payoffs can be negative)
- ▷ apply these tools to both social and market valuation!!

The social cost of climate change and social value of research and development are asset prices with uncertain social "cash flows."

Valuation under ambiguity and misspecification aversion

- use discounted expected values of social cash flows as is typical in cost-benefit analyses
- but expectations are constructed using the minimizing probabilities in order to capture the full uncertainty adjustments

Apply stochastic discounting under a probability measure inferred from a decision problem.

Uncertainty in financial markets

Ambition:

- a mechanism for inducing fluctuations in asset values
- investors fear persistence in bad times and fear the lack of persistence in good times

Link to Paper 1

Link to Paper 2

Link to a perspective piece

Market valuation: uncertain growth-rate dynamics

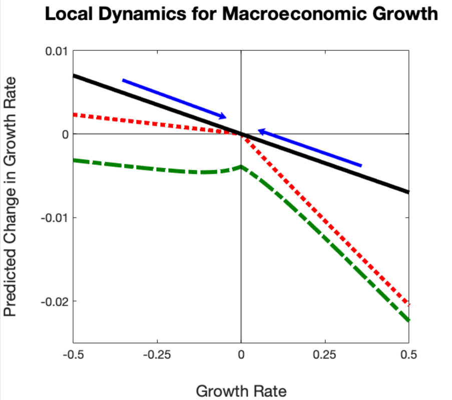

black line: model approximation; red dotted curve: adjusts for ambiguity aversion; green dashed curve: adjusts for ambiguity and misspecification aversion

Compounding the uncertainty

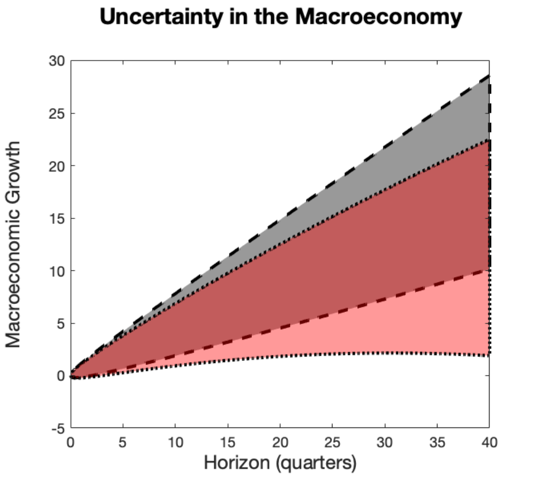

The grey shaded area gives the region between the .1 and .9 deciles. The red shaded area gives the region within the .1 and .9 deciles that includes both model ambiguity and model misspecification concerns.

Broader perspective

- difficult to disentangle risk aversion from belief distortions
- belief distortions are more compelling in environments in which uncertainty is complex
- statistical and decision-theoretic tools provide valuable ways to assess environmental complexity
- ambiguity and model misspecification aversion induce nonlinearity dynamics into valuation

Social valuation: climate policy under uncertainty

- There are many calls for immediate climate policy implementation
- The existing limits to our understanding of the timing and magnitude of climate change impacts have led to apprehension by some
- We study how a decision-maker confronts uncertainty in a setting where:
    - there will be future information about damage severity
    - but the value of further empiricism in the near term is limited

Social assets

that are pertinent to devising robust climate-economic policies

- social cost of global temperature change
- social benefit to research and development that supports the discovery of economically viable new green technologies

Both can be viewed as assets or liabilities that generate intertemporal social cash flows.

What is the challenge?

Four sources 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: research and development invested today may eventually lead to economically viable technologies

Research question: Which of the four sources is of most concern for designing policy?

Our initial research shows that: the unknown outcome of R&D investment is the most potent contributor to uncertainty for climate-economics policy. This source of uncertainty leads to doing more green R &D investment.

Concluding remarks

- Uncertainty matters for policy tools like the social cost of global warming and social investment in green research and development.
- Understanding the sources of uncertainty, broadly conceived, used by the private sector and by governments will make economic policy more effective.