Uncertainty and all that

Massimo Marinacci

AXA-Bocconi Chair in Risk
Department of Decision Sciences and IGIER
Università Bocconi

EAERE 23rd Annual Conference
Athens, 1 July 2017

Probability of facts and of theories

- Decisions' consequences depend on external factors (contingencies)
- Probability of contingencies
- Probabilistic theories on contingencies (e.g., generative mechanisms, DGP)
- Thinking over such theories
- Two layers of uncertainty

Decision problems: the toolbox, I

A decision problem consists of

- a space A of actions
- a space C of material (e.g., monetary) consequences
- a space S of environment states
- a consequence function $\rho : A \times S \longrightarrow C$ that details the consequence
$$c = \rho(a, s)$$
of action a when state s obtains
- We abstract from state misspecification issues (e.g., unforeseen contingencies)

Example (i): natural hazards

Public officials have to decide whether or not to evacuate an area because of a possible earthquake

- A two actions a_0 (no evacuation) and a_1 (evacuation)
- C monetary consequences (damages to infrastructures and human casualties; Mercalli-type scale)
- S possible peak ground accelerations (Richter-type scale)
- $c = \rho\left(a, s\right)$ the monetary consequence of action a when state s obtains

Example (ii): monetary policy example

ECB or the FED have to decide some target level of inflation to control the economy unemployment and inflation

- Unemployment u and inflation π outcomes are connected to shocks $\varepsilon = (\varepsilon_u, \varepsilon_\pi)$ and the policy a according to

$$u = \theta_0 + \theta_{1\pi}\pi + \theta_{1a}a + \varepsilon_u$$
$$\pi = a + \varepsilon_\pi$$

$\theta = (\theta_0, \theta_{1\pi}, \theta_{1a})$ are three structural coefficients

- $\theta_{1\pi}$ and θ_{1a} are slope responses of unemployment to actual and planned inflation (e.g., Lucas-Sargent $\theta_{1a} = -\theta_{1\pi}$; Samuelson-Solow $\theta_{1a} = 0$)
- θ_0 is the rate of unemployment that would (systematically) prevail without policy interventions

Example (ii): monetary policy

Here:

A the target levels of inflation

C the pairs $c = (u, \pi)$

S has structural and random components

$$s = (\varepsilon, \theta)$$

The reduced form is

$$\begin{aligned} u &= \theta_0 + (\theta_{1\pi} + \theta_{1a})\, a + \theta_{1\pi}\varepsilon + \varepsilon_u \\ \pi &= a + \varepsilon_\pi \end{aligned}$$

and so ρ has the form

$$\rho\,(a, w, \varepsilon, \theta) = \begin{bmatrix} \theta_0 \\ 0 \end{bmatrix} + a \begin{bmatrix} \theta_{1\pi} + \theta_{1a} \\ 1 \end{bmatrix} + \begin{bmatrix} 1 & \theta_{1\pi} \\ 0 & 1 \end{bmatrix} \begin{bmatrix} \varepsilon_u \\ \varepsilon_\pi \end{bmatrix}$$

Example (iii): climate policy

- A policy maker has to decide some target greenhouse gas emissions level to control damages associated with global temperatures increases.
- Different sources of uncertainty are relevant (cf. Heal and Millner 2014)

Example (iii): climate policy

- *Scientific uncertainty*: how do emissions E translate in increases of temperatures T? Assume

$$T = \theta_T E + \varepsilon_T$$

  where θ_T is a structural CCR (carbon-climate response) parameter and ε_T is a random component (cf. Matthews et al. 2009)
- *Socioeconomic uncertainty*: how do increases of temperatures T translate in economic damages D? Assume a DICE quadratic

$$D = \theta_{1D} T + \theta_{2D} T^2 + \varepsilon_D$$

  where θ_{1D} and θ_{2D} are structural parameters and ε_D is a random component
- We abstract from issues about the objective functions

Example (iii): climate policy

- Random components: shocks (i.e., minor omitted explanatory variables which we are "unable and unwilling to specify") or measurement errors
- Cf. the works of Hurwicz, Koopmans and Marschak in the 1940s and 1950s

Example (iii): climate policy

Here:

A emission policies

C the economic damages (in GDP terms)

S has structural and random components

$$s = (\varepsilon, \theta)$$

where

$$\varepsilon = (\varepsilon_T, \varepsilon_D)$$

are the random components affecting the climate and economic systems, and

$$\theta = (\theta_T, \theta_{1D}, \theta_{2D})$$

are their structural coefficients

Example (iii): climate policy

- Action a is an emission policy, with cost $c(a)$
- $d(a, \varepsilon, \theta)$ economic damage function
- $\rho(a, \varepsilon, \theta) = -d(a, \varepsilon, \theta) - c(a)$ is the overall consequence of policy a
- From

$$\begin{cases} T = \theta_T a + \varepsilon_T \\ D = \theta_{1D} T + \theta_{2D} T^2 + \varepsilon_D \end{cases}$$

it follows that

$$\begin{aligned} d(a, \varepsilon, \theta) &= -(\theta_{1D}\theta_T + 2\theta_{2D}\varepsilon_T) a - \theta_{2D}\theta_T^2 a^2 - \theta_{1D}\varepsilon_T \\ &\quad -\theta_{2D}\varepsilon_T^2 - \varepsilon_D \end{aligned}$$

Decision problems: the toolbox, II

- The quartet (A, S, C, ρ) is a *decision form under uncertainty*
- The decision maker (DM) has a preference $\succsim$ over actions
  - we write $a \succsim b$ if the DM (weakly) prefers action a to action b
- The quintet $(A, S, C, \rho, \succsim)$ is a *decision problem under uncertainty*
- DMs aim to select actions $\hat{a} \in A$ such that $\hat{a} \succsim a$ for all $a \in A$
- Static setting, we abstract from temporal/dynamic issues (cf. Gollier 2013)

Consequentialism

What matters about actions is not their label / name but the *consequences* that they determine when the different states obtain

- *Consequentialism*: two actions that are realization equivalent – i.e., that generate the same consequence in every state – are indifferent
- Formally,

$$\rho(a,s) = \rho(b,s) \quad \forall s \in S \Longrightarrow a \sim b$$

or, equivalently,

$$\rho_a = \rho_b \Longrightarrow a \sim b$$

- Here $\rho_a : S \to C$ is the section of ρ at a given by $\rho_a(s) = \rho(a,s)$
- The section ρ_a is a (Savage) *act*

Probability models

- Because of their ex-ante structural information, DMs know that states are generated by a probability model m that belongs to a given subset M of $\Delta(S)$
- Each m describes a possible *DGP*, so it represents *physical uncertainty* (risk)
- DMs thus posit a model space M in addition to the state space S, a central tenet of classical statistics a la Neyman-Pearson-Wald
- When the model space is based on experts' advice, its nonsingleton nature may reflect different advice

Models: a toy example

Consider an urn with 90 Red, or Green, or Yellow balls

- DMs bet on the color of a ball drawn from the urn
- State space is $S = \{R, G, Y\}$
- Without any further information, $M = \Delta(\{R, G, Y\})$
- If DMs are told that 30 balls are red, then

$$M = \{m \in \Delta(\{R, G, Y\}) : m(R) = 1/3\}$$

Models and experts: probability of heart attack

Two DMs: John and Lisa are 70 years old

- smoke
- no blood pressure problem
- total cholesterol level 310 mg/dL
- HDL-C (good cholesterol) 45 mg/dL
- systolic blood pressure 130

What's the probability of a heart attack in the next 10 years?

Models and experts: probability of heart attack

Based on their data and medical models, experts say

Experts	John's m	Lisa's m
Mayo Clinic	25%	11%
National Cholesterol Education Program	27%	21%
American Heart Association	25%	11%
Medical College of Wisconsin	53%	27%
University of Maryland Heart Center	50%	27%

Table from Gilboa and Marinacci (2013)

Uncertainty: a taxonomy

In this setup, we can decompose uncertainty in three distinct layers (cf. Hansen and Marinacci 2016):

- *Physical* or *aleatory uncertainty* (*risk*): uncertainty within a model m
- *Model ambiguity* or *uncertainty*: uncertainty across models in M
- *Model misspecification*: the true model does not belong to the posited set M

Models: a consistency condition

- Cerreia, Maccheroni, Marinacci, Montrucchio (*PNAS* 2013) take the "physical" information M as a primitive and thus enrich the standard framework
- DMs know that the true model m that generates observations belongs to the posited collection M
- In terms of preferences: betting behavior must be *consistent* with datum M, i.e.,

$$m(F) \geq m(E) \quad \forall m \in M \Longrightarrow \text{“bet on } F\text{''} \succsim \text{“bet on } E\text{''}$$

- The sextet $(A, S, C, M, \rho, \succsim)$ forms a *classical decision problem under uncertainty*
- We abstract from model misspecification issues

Classical subjective EU

We show that a preference $\succsim$ that satisfies Savage's axioms and the consistency condition is represented by the criterion

$$V(a) = \sum_{m \in M} \left(\sum_{s \in S} u(\rho(a, s))\, m(s) \right) \mu(m) \tag{1}$$

That is, acts a and b are ranked as follows:

$$a \succsim b \Longleftrightarrow V(a) \geq V(b)$$

Here

- u is a von Neumann-Morgenstern utility function that captures *risk attitudes* (i.e., attitudes toward physical uncertainty)
- μ is a *subjective prior probability* that quantifies the epistemic uncertainty about models; its support is included in M
- If M is based on the advice of different experts, the prior may reflect the *different confidence* that DMs have in each of them

Classical subjective EU

We call this representation *Classical Subjective Expected Utility* because of the classical statistics tenet on which it relies

- If we set
$$R(a, m) = \sum_{s \in S} u(\rho(a, s))\, m(s)$$
we can write the criterion as
$$V(a) = \sum_{m \in M} R(a, m)\, \mu(m)$$

- In words, the criterion considers the expected utility $R(a, m)$ of each possible model m, and averages them out according to the prior μ

Classical subjective EU

- Each prior μ induces a *predictive probability* $\bar{\mu} \in \Delta(S)$ through reduction

$$\bar{\mu}(E) = \sum_{m \in M} m(E)\, \mu(m)$$

In turn, the predictive probability enables to rewrite the representation as

$$V(a) = R(a, \bar{\mu}) = \sum_{s \in S} u(\rho(a, s))\, \bar{\mu}(s)$$

- This reduced form of V is the original Savage subjective EU representation

Classical subjective EU: some special cases

- If the support of μ is a singleton $\{m\}$, DMs subjectively (and so possibly wrongly) believe that m is the true model
  The criterion thus reduces to a Savage EU criterion $R(a, m)$
- If M is a singleton $\{m\}$, DMs know that m is the true model (a rational expectations tenet)
  (i) There is no epistemic uncertainty, but only physical uncertainty (quantified by m)
  (ii) The criterion again reduces to the EU representation $R(a, m)$, but now interpreted as a *von Neumann-Morgenstern criterion*

Classical subjective EU: some special cases

- Classical subjective EU thus encompasses both the Savage and the von Neumann-Morgenstern representations
- If $M \subseteq \{\delta_s : s \in S\}$, there is no physical uncertainty, but only epistemic uncertainty (quantified by μ). By identifying s with δ_s, wlog we can write $\mu(s)$ and so the criterion takes the form

$$V(a) = \sum u(\rho(a, s))\, \mu(s)$$

where it is u that matters

Classical subjective EU: some special cases

- Singleton M have been pervasive in economics
- Since the 70s, economics has emphasized the study of agents' reactions to "opponents" actions (from the Lucas critique in macroeconomics to the study of incentives in game theoretic settings)
- Rational expectations literature had to depart from the "particle" view of agents of the Keynesian macroeconomics of the 50s and 60s

Factorization

- In applications, states often have structural and random components

$$s = (\varepsilon, \theta)$$

- So, here $m(\varepsilon, \theta)$ is a joint probability
- Adopt the factorization $m = q \times \delta_\theta$, that is,

$$m(\varepsilon, \theta') = \begin{cases} q(\varepsilon) & \text{if } \theta' = \theta \\ 0 & \text{else} \end{cases}$$

where $q(\varepsilon)$ is the probability of ε and δ_θ is the (degenerate) probability distribution concentrated on θ

- Each model corresponds to
  1. a distribution q of the random component ε
  2. a model climate system/economy θ

Factorization

- In the factorization $m = q \times \delta_\theta$, two kinds of model uncertainties emerge
- *Theoretical model uncertainty* about the economic and physical theories that underpin the models: different θ correspond to different theories
- *Stochastic model uncertainty* about the statistical performance of such theories, due to shocks and to measurement errors: different q correspond to different performances

Factorization

- Theoretical model uncertainty is the more "fundamental"
- To focus on it, assume that the distribution q is known and common across models
- Different models m thus correspond to different structural components θ
- Formally, we can parametrize models via their structural components:

$$m_\theta = q \times \delta_\theta \qquad \forall \theta \in \Theta$$

Factorization

- Physical uncertainty is quantified by q
- Epistemic uncertainty is about the structural coefficient θ
- To address it, the DM has a prior probability $\mu(\theta)$ that quantifies DM's degree of belief that θ is the true parameter

Classical subjective EU under factorization

Under factorization, the classical subjective expected utility criterion becomes

$$V(a) = \sum_{\theta \in \Theta} \left(\sum_{\varepsilon \in E} u(\rho(a, \varepsilon, \theta))\, q(\varepsilon) \right) \mu(\theta)$$

or, equivalently,

$$V(a) = \sum_{\theta \in \Theta} R(a, \theta)\, \mu(\theta)$$

where $R(a, \theta) = \sum_{\varepsilon \in E} u(\rho(a, \varepsilon, \theta))\, q(\varepsilon)$

Factorized classical subjective EU: monetary policy example

- Back to the monetary example

$$u = \theta_0 + \theta_{1\pi}\pi + \theta_{1a}a + \varepsilon_u$$
$$\pi = a + \varepsilon_\pi$$

- Distribution q of shock ε is known
- Model economy θ is unknown
- So, belief μ is directly on θ
- The monetary policy problem is then

$$\max_{a\in A} V(a) = \max_{a\in A} \sum_{\theta\in\Theta} \left(\sum_{\varepsilon\in E} u(\rho(\varepsilon,\theta))\, q(\varepsilon) \right) \mu(\theta)$$

Factorized classical subjective EU: climate policy example

- Back to the climate policy example

$$\begin{cases} T = \theta_T a + \varepsilon_T \\ D = \theta_{1D} T + \theta_{2D} T^2 + \varepsilon_D \end{cases}$$

- Distribution q of shocks ε is known
- Model climate system and model economy θ is unknown
- So, belief μ is directly on θ
- The climate policy problem is then

$$\max_{a \in A} V(a) = \max_{a \in A} \sum_{\theta \in \Theta} \left(\sum_{\varepsilon \in E} u(\rho(a, \varepsilon, \theta)) \, q(\varepsilon) \right) \mu(\theta)$$

Factorized classical subjective EU: climate policy example

- Under risk neutrality and zero actions' cost, we have

$$R(a, \theta) = -\theta_{1D}\theta_T a - \theta_{2D}\theta_T^2 a^2 - \theta_{2D}$$

provided the random components have zero mean and unit variance

- The climate policy problem is then

$$\max_{a \in A} -a\mathbb{E}_{\mu}(\theta_{1D}\theta_T) - a^2\mathbb{E}_{\mu}(\theta_{2D}\theta_T^2) - \mathbb{E}_{\mu}(\theta_{2D})$$

with optimal policy

$$\hat{a} = -\frac{\mathbb{E}_{\mu}(\theta_{1D}\theta_T)}{\mathbb{E}_{\mu}(\theta_{2D}\theta_T^2)}$$

Road map

- Decision problems
  - toolbox
  - Savage setup
  - classical subjective expected utility
- **Model uncertainty: ambiguity / robustness models**
- Issues
  - ambiguity / robustness makes optimal actions more prudent?
  - ambiguity / robustness favors diversification?
  - ambiguity / robustness affects valuation?
  - model uncertainty resolves in the long run through learning?
  - sources of uncertainty: a Pandora's box?

Ambiguity / Robustness: the problem

- Physical and epistemic uncertainties need to be treated differently
- The standard expected utility model does not
- Since the 1990s, a strand of economic literature has been studying *ambiguity* / *Knightian uncertainty* / *robustness* / *deep uncertainty*
- We consider two approaches
  - non-Bayesian (Gilboa and Schmeidler 1989; Schmeidler 1989)
  - Bayesian (Klibanoff, Marinacci, Mukerji 2005)
- Both approaches broaden the scope of traditional EU analysis
- Normative focus (no behavioral biases or "mistakes"; see Gilboa and Marinacci 2013)

Ambiguity / Robustness: the problem

- Intuition: betting on coins is greatly affected by whether or not coins are well tested
- Models correspond to possible biases of the coin
- By symmetry (uniform reduction), heads and tails are judged to be equally likely when betting on an untested coin, never flipped before
- The same probabilistic judgement holds for a well tested coin, flipped a number of times with an approximately equal proportion of heads to tails
- The evidence behind such judgements, and so the confidence in them, is dramatically different: ceteris paribus, DMs may well prefer to bet on tested (phys. unc.) rather than on untested coins (phys. & epist. unc.)
- Experimental evidence: Ellsberg paradox

Ambiguity / Robustness: relevance

- A more robust rational behavior toward uncertainty emerges
- A more accurate / realistic account of how uncertainty affects valuation (e.g., uncertainty premia in market prices)
- Better understanding of exchange mechanics
  - a dark side of uncertainty: no-trade or small-trade results because of cumulative effects of physical and epistemic uncertainty; See the recent financial crisis
- Better calibration and quantitative exercises
  - applications in Finance, Macroeconomics, and Environmental Economics
- Better modelling of decision / policy making
  - applications in Risk Management; e.g., the otherwise elusive precautionary principle may fit within this framework

Ambiguity / Robustness: relevance

- Caveat: risk and model uncertainty can work in the same direction (magnification effects), as well as in different directions
- Magnification effects: large "uncertainty prices" with reasonable degrees of risk aversion
- Combination of sophisticated formal reasoning and empirical relevance

Ambiguity / Robustness: a Bayesian approach

- A first distinction: DMs do not have attitudes toward uncertainty per se, but rather toward physical uncertainty and toward epistemic uncertainty
- Such attitudes may differ: typically DMs are more averse to epistemic than to physical uncertainty
- Berger and Bosetti (2016) provide experimental evidence

Bayesian approach: a tacit assumption

Suppose acts are monetary

- Classical subjective EU representation can be written as

$$\begin{aligned} V(a) &= \sum_{m \in M} R(a,m)\,\mu(m) \\ &= \sum_{m \in M} \left(u \circ u^{-1}\right)(R(a,m))\,\mu(m) \\ &= \sum_{m \in M} u(c(a,m))\,\mu(m) \end{aligned}$$

  where $c(a,m)$ is the certainty equivalent

$$c(a,m) = u^{-1}(R(a,m))$$

  of action a under model m

- Recall that $R(a,m) = \sum_{s \in S} u(\rho(a,s))\,m(s)$

Bayesian approach: a tacit assumption

- The profile

$$\{c(a,m) : m \in \operatorname{supp} \mu\}$$

  is the scope of the model uncertainty that is relevant for the decision
- In particular, DMs use the decision criterion

$$V(a) = \sum_{m \in M} u(c(a,m))\, \mu(m)$$

  to address model uncertainty, while

$$R(a,m) = \sum_{s \in S} u(\rho(a,s))\, m(s)$$

  is how DMs address the physical uncertainty that each model m features
- Identical attitudes toward physical and epistemic uncertainties, both modeled by the same function u

Bayesian approach: representation

- The smooth ambiguity model generalizes the representation by distinguishing such attitudes
- Acts are ranked according to the smooth (ambiguity) criterion

$$
\begin{aligned}
V(a) &= \sum_{m \in M} \left(v \circ u^{-1}\right)(R(a,m))\,\mu(m) \\
&= \sum_{m \in M} v(c(a,m))\,\mu(m)
\end{aligned}
$$

- The function $v : C \longrightarrow \mathbb{R}$ represents attitudes toward model uncertainty
- A negative attitude toward model uncertainty is modelled by a concave v, interpreted as aversion to (mean preserving) spreads in certainty equivalents $c(a,m)$
- Ambiguity aversion amounts to a higher degree of aversion toward epistemic than toward physical uncertainty, i.e., a v more concave than u

Bayesian approach: representation

- Setting $\phi = v \circ u^{-1}$, the smooth criterion can be written as

$$V(a) = \sum_{m \in M} \phi(R(a, m))\, \mu(m)$$

- This formulation holds for any kind of acts (not just monetary)
- Ambiguity aversion corresponds to the concavity of ϕ
- If $\phi(x) = -e^{-\lambda x}$, it is a Bayesian version of the multiplier preferences of Hansen and Sargent (2001, 2008)
- Sources of uncertainty now matter (no longer "uncertainty is reduced to risk")

Bayesian approach: example

- Call I the tested coin and II the untested one
- Actions a_I and a_{II} are, respectively, bets of one euro on coin I and on coin II
- $S = \{H, T\} \times \{H, T\} = \{HH, HT, TH, TT\}$
- The next table summarizes the decision problem

	HH	HT	TH	TT
$\mathbf{a}_I$	1	1	0	0
$\mathbf{a}_{II}$	1	0	1	0

Bayesian approach: example

- Given the available information, it is natural to set

$$M = \left\{ m \in \Delta(S) : m(HH \cup HT) = m(TH \cup TT) = \frac{1}{2} \right\}$$

- M consists of all models that give probability $1/2$ to either outcome for the tested coin; no specific probability is, instead, assigned to the outcome of the untested coin

Bayesian approach: example

- Normalize $u(1) = 1$ and $u(0) = 0$, so that

$$V(\mathbf{a}_I) = \sum_{m \in M} \phi(m(HH \cup HT))\, d\mu(m) = \phi\left(\frac{1}{2}\right)$$

  and

$$V(\mathbf{a}_{II}) = \sum_{m \in M} \phi(m(HH \cup TH))\, d\mu(m)$$

- If μ is uniform, $V(\mathbf{a}_{II}) = \int_0^1 \phi(x)\, dx$. If ϕ is strictly concave, by the Jensen inequality we then have

$$V(\mathbf{a}_{II}) = \int_0^1 \phi(x)\, dx < \phi\left(\int_0^1 x dx\right) = \phi\left(\frac{1}{2}\right) = V(\mathbf{a}_I)$$

Bayesian approach: extreme attitudes and maxmin

- Under extreme ambiguity aversion (e.g., as $\lambda \uparrow \infty$ when $\phi(x) = -e^{-\lambda x}$), the smooth ambiguity criterion in the limit reduces to the maxmin criterion

$$V(a) = \min_{m \in \operatorname{supp} \mu} \sum_{s \in S} u(\rho(a, s)) \, m(s)$$

- Pessimistic criterion: DMs maxminimize over all possible probability models in the support of μ
- The prior μ just selects which models in M are relevant
- Waldean version of Gilboa and Schmeidler (1989) seminal maxmin decision model

Bayesian approach: extreme attitudes and maxmin

- If $\operatorname{supp} \mu = M$, the prior is actually irrelevant and we get back to the Wald (1950) maxmin criterion

$$V(a) = \min_{m \in M} \sum_{s \in S} u(\rho(a,s))\, m(s)$$

- When M consists of all possible models, it reduces to the statewise maxmin criterion

$$V(a) = \min_{s \in S} u(\rho(a,s))$$

  A very pessimistic (paranoid?) criterion: probabilities, of any sort, do not play any role (Arrow-Hurwicz decision under ignorance)

- Precautionary principle

Bayesian approach: extreme attitudes and no trade

In a frictionless market a primary asset y that pays $y(s)$ if state s obtains, is traded

- Its market price is p
- Investors may trade x units of the asset (buy if $x > 0$, sell if $x < 0$, no trade if $x = 0$)
- State contingent payoff is $\mathbf{x}(s) = y(s)x - px$
- Trade occurs only if $V(\mathbf{x}) \geq V(\mathbf{0}) = 0$

Bayesian approach: extreme attitudes and no trade

- Dow and Werlang (1992): under maxmin behavior, there is no trade on asset y whenever

$$\min_{m \in \operatorname{supp} \mu} \mathrm{E}_m(y) < p < \max_{m \in \operatorname{supp} \mu} \mathrm{E}_m(y) \qquad (2)$$

- High ambiguity aversion may freeze markets
- Inequality (2) requires $\operatorname{supp} \mu$ to be nonsingleton: the result requires ambiguity
- More generally: a lower trade volume on asset y corresponds to a higher ambiguity aversion (e.g., higher λ when $\phi(x) = -e^{-\lambda x}$) if (2) holds
- Bottom line: it reinforces the idea that uncertainty can be an impediment to trade

Bayesian approach: quadratic approximation

- The smooth ambiguity criterion admits a simple quadratic approximation that leads to a generalization of the classic mean-variance model (Maccheroni, Marinacci, Ruffino 2013)
- The random variable

$$\rho(a, \cdot) : S \longrightarrow C$$

induced by action a is denoted by $\mathbf{a}$ and called (Savage) *act*

Bayesian approach: quadratic approximation

- The robust mean-variance rule ranks acts **a** by

$$\mathrm{E}_{\bar{\mu}}(\mathbf{a}) - \frac{\lambda}{2}\sigma^2_{\bar{\mu}}(\mathbf{a}) - \frac{\theta}{2}\sigma^2_{\mu}(\mathrm{E}(\mathbf{a}))$$

  where λ and θ are positive coefficients
- Here $\mathrm{E}(\mathbf{a}) : M \to \mathbb{R}$ is the random variable

$$m \mapsto \mathrm{E}_m(\mathbf{a}) = \sum_{s \in S} \mathbf{a}(s)\, m(s)$$

  that associates the EV of act **a** under each possible model m
- $\sigma^2_{\mu}(\mathrm{E}(\mathbf{a}))$ is its variance

Bayesian approach: quadratic approximation

- The robust mean-variance rule

$$\mathrm{E}_{\bar{\mu}}(\mathbf{a}) - \frac{\lambda}{2}\sigma^2_{\bar{\mu}}(\mathbf{a}) - \frac{\theta}{2}\sigma^2_{\mu}(\mathrm{E}(\mathbf{a}))$$

  is determined by the three parameters λ, θ, and μ. When $\theta = 0$ we return to the usual mean-variance rule
- The taste parameters λ and θ model DMs' attitudes toward physical and epistemic uncertainty, resp.
- Higher values of these parameters correspond to stronger negative attitudes

Bayesian approach: quadratic approximation

- The information parameter μ determines the variances $\sigma^2_{\bar{\mu}}(\mathbf{a})$ and $\sigma^2_{\mu}(\mathrm{E}(\mathbf{a}))$ that measure the physical and epistemic uncertainty that DMs perceive in the evaluation of act **a**
- Higher values of these variances correspond to a DM's poorer information regarding such uncertainties
- As usual, the risk premium is

$$\frac{\lambda}{2}\sigma^2_{\bar{\mu}}(\mathbf{a})$$

- Novelty: the ambiguity premium is

$$\frac{\theta}{2}\sigma^2_{\mu}(\mathrm{E}(\mathbf{a}))$$

Ambiguity / Robustness: a non Bayesian approach

- Need to relax the requirement that a single number quantifies beliefs: the multiple (prior) probabilities model
- DMs may not have enough information to quantify their beliefs through a single probability, but need a set of them
- Expected utility is computed with respect to each probability and DMs act according to the minimum among such expected utilities

Non Bayesian approach: representation

- Epistemic uncertainty quantified by a set C of priors
- DMs use the criterion

$$V(\mathbf{a}) = \min_{\mu \in C} \sum_{m \in M} \left(\sum_{s \in S} u(\rho(a,s))\, m(s) \right) \mu(m) = \min_{\mu \in C} \sum_{s \in S} u(\rho(a,s))\, \bar{\mu}(s) \tag{3}$$

- DMs consider the least among all the EU determined by each prior in C
- The predictive form (3) is the original version axiomatized by Gilboa and Schmeidler (1989)

Non Bayesian approach: comments

- This criterion is less extreme than it may appear at a first glance
- The set C incorporates
  - the attitude toward ambiguity, a taste component
  - its perception, an information component
- A smaller set C may reflect both better information – i.e., a lower perception of ambiguity – and / or a less averse uncertainty attitude
- In sum, the size of C does not reflect just information, but taste as well

Non Bayesian approach: comments

- With singletons $C = \{\mu\}$ we return to the classical subjective EU criterion
- When C consists of all possible priors on M, we return to the Wald maxmin criterion

$$\min_{m \in M} \sum_{s \in S} u\left(\rho\left(a, s\right)\right) m\left(s\right)$$

- No trade results (kinks)

Non Bayesian approach: comments

A more general α-maxmin model has been axiomatized by Ghirardato, Maccheroni, and Marinacci (2004):

$$
\begin{aligned}
V(a) \quad = \quad & \alpha \min_{\mu \in C} \sum_{m \in M} \left(\sum_{s \in S} u\left(\rho\left(a, s\right)\right) m\left(s\right) \right) \mu(m) \\
& + (1 - \alpha) \max_{\mu \in C} \sum_{m \in M} \left(\sum_{s \in S} u\left(\rho\left(a, s\right)\right) m\left(s\right) \right) \mu(m)
\end{aligned}
$$

Non Bayesian approach: variational model

- In the maxmin model, a prior μ is either "in" or "out" of the set C
- Maccheroni, Marinacci, Rustichini (2006): general variational representation

$$V(\mathbf{a}) = \inf_{\mu \in \Delta(M)} \left(\sum_{m \in M} \left(\sum_{s \in S} u(\rho(a,s))\, m(s) \right) \mu(m) + c(\mu) \right)$$

where $c(\mu)$ is a convex function that weights each prior μ

- If c is the dichotomic function given by

$$\delta_{\mathrm{C}}(\mu) = \begin{cases} 0 & \text{if } \mu \in \mathrm{C} \\ +\infty & \text{else} \end{cases}$$

we get back to the maxmin model with set of priors C

Non Bayesian approach: multiplier model

- If c is given by the relative entropy $R(\mu||\nu)$, where ν is a reference prior, we get the multiplier model

$$V(a) = \inf_{\mu \in \Delta(M)} \left(\sum_{m \in M} \left(\sum_{s \in S} u(\rho(a,s))\, m(s) \right) \mu(m) + \alpha R(\mu||\nu) \right)$$

popularized by Hansen and Sargent in their studies on robustness in Macroeconomics

- Also the mean-variance model is variational, with c given by a Gini index

Illustration: climate policy example under ambiguity

- Recent related works: Athanassoglou and Xepapadeas (2012), Millner, Dietz, and Heal (2013); Heal and Millner (2014, 2015), Drouet, Bosetti and Tavoni (2015), Berger (2015), Berger, Emmerling, and Tavoni (2017), Lemoine and Traeger (2016), Chambers and Melkonyan (2017), Koundouri, Pittis, Samartzis, Englezos, and Papandreou (2017), Rudik (2017), Xepapadeas and Yannacopoulos (2017)

Illustration: climate policy example under ambiguity

- Berger and Marinacci (2017), following Drouet, Bosetti and Tavoni (2015), consider the damage function

$$D = \theta_{1D} T + \theta_{2D} T^2 + \theta_{3D} T^6 + \theta_{4D} (\mathrm{e}^{-\theta_{5D} T^2} - 1) + \varepsilon_D$$

  with 3 possible specifications (quadratic, exponential and sextic)
- Consider 11 possible values of the CCR parameter θ_T
- So, there are 33 models to consider
- The next figures give the certainty equivalents

$$c(a, \theta) = u^{-1}(R(a, \theta))$$

  of the 33 models, with a power u

Illustration: climate policy example under ambiguity

Quadratic damage:

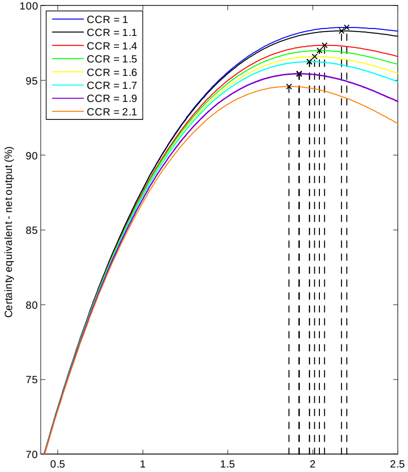

Illustration: climate policy example under ambiguity

Sextic damage:

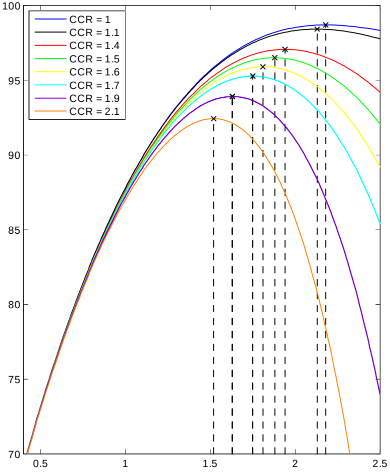

Illustration: climate policy example under ambiguity

Exponential damage:

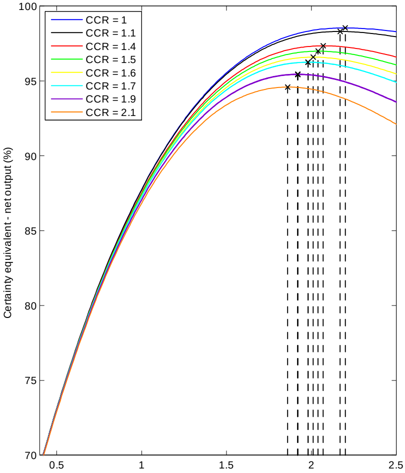

Illustration: climate policy example under ambiguity

- Assume a uniform μ, that is, $\mu(\theta) = 1/33$
- Ambiguity aversion makes optimal policies more prudent
- The next figure illustrates

Illustration: climate policy example under ambiguity

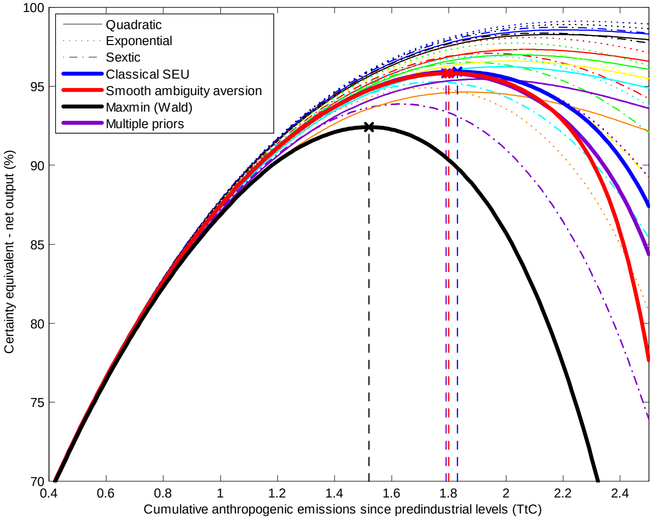

Sources of uncertainty

- We made a distinction between attitudes toward physical and epistemic uncertainty
- A more general issue: do attitudes toward different uncertainties differ?
- Source contingent outcomes: Do DMs regard outcomes (even monetary) that depend on different sources as different economic objects?
- Ongoing research on this subtle topic

Epilogue

- In decision problems with data, it is important to distinguish physical and epistemic uncertainty
- Traditional EU reduces epistemic uncertainty to physical uncertainty, and so it ignores the distinction
- Experimental and empirical evidence suggest that the distinction is relevant and may affect valuation
- We presented two approaches, one Bayesian and one not
- For different applications, different approaches may be most appropriate
- Traditional EU is the benchmark
- Yet, adding ambiguity broadens the scope (empirical and theoretical) and the robustness of results