EPISTEMOLOGY SEMINAR · MASTER ECONOMICS AND HUMANITIES

The Epistemology of Decision Theory

Bayesian Rationality and Empirical Methods, from Bernoulli to Savage and Randomized Experiments

Niels Boissonnet & Judith Favereau

Université Paris 1 Panthéon-Sorbonne · January 2021

Introduction: what is decision theory for?

The usual reading

- **Descriptive:** how decision makers (DMs) actually behave.
- **Normative:** how they ought to behave.

A further motivation (Gilboa): metatheoretical

- **Ontological:** give probability and statistics a subjective foundation.
- **Methodological:** justify certain empirical methods and produce identification strategies.

What does this imply for **experimental economics**?

Outline

Interpreting probabilities, founding statistics
Operationalising subjective probabilities
Decision theory and experimental economics

The thread

Decision theory as a theory of the **statistician**: first to found probabilities and statistics, then to model the **experimenter** who chooses an experimental design.

1

Interpreting probabilities, founding statistics

Frequencies, logical relations, degrees of belief: three readings of probability, and what each one implies for statistics.

Probabilities as frequencies

Foundations

Bernoulli, the origins of statistics, the law of large numbers, etc.

Idea

Number of actual occurrences divided by the number of possible occurrences.

Advantages

Natural, easy to operationalise, objective (?)

Criticisms

- **Natural but not very general:** meaningless when data are scarce. Not connected to uncertainty: “I do not know that p” $\neq$ “I find p more or less probable”.
- **Conceptually limited:** no Bayesian conditioning to explain how the statistician should revise her estimate.
- **Not that objective:** how should the statistician choose between different statistical tests?

Probabilities as logical relations

Foundations

Keynes (1921) and Wittgenstein's *Tractatus* (Wittgenstein, 1921).

Idea

Probability $=$ an **extension of propositional logic**. Instead of $p \to q$, write $P(q \mid p) = 2/3$: p implies q *to degree* $2/3$ (see also Wittgenstein on elementary propositions).

Advantages

Objective and general; the relations between propositions are well defined.

Criticisms

- **Metaphysical:** "Probability does not exist" (de Finetti and Ramsey).
- **Operational:** even if it is objectively true that $P(p) = 1/2$, how can we know it? Hence Keynes's **radical uncertainty**, and the uselessness of probability and statistics.
- **Holistic:** everyone should assign the same probabilities.

Subjective probability

Foundations

Ramsey (1931), de Finetti (1937) and Savage (1954).

Idea

Probability expresses a **degree of belief**.

Advantages

General; ontologically cheap; relations between propositions.

Criticism

- **Operational:** how can we determine this probability if it is idiosyncratic?
- And what if everyone holds a different one?

This operational question drives Section 2: how to **reveal** subjective probabilities from behaviour.

Consequences for statistics: Bayesians vs frequentists

In practice

Different methods which, in many circumstances, give “different” results.

Historically

Frequentist methods (maximum likelihood, hypothesis tests) dominate. But over the last thirty years or so, **Bayesian** methods have made a strong comeback, driven by the development of computers.

Example: tossing a coin

Let $p =$ "the coin lands tails".

Frequentist

- $P(p)$: over many **identical** tosses.
- The **data** (heads or tails) vary.
- The **parameters** are fixed $(1/2)$.
- After 10 tosses with 7 tails, the frequentist says $P(p) = 7/10$ (with a confidence interval).

Bayesian

- $P(p)$: over tosses that are **not identical** but uncertain (how the coin is tossed, etc.).
- The **data** are fixed.
- The **parameters** vary.
- A Bayesian first holds a **prior**, which she updates after the tosses to obtain a **posterior**.

Problem

How do we know which prior is the right one? What if researchers disagree about the prior? Is this equivalent to having several tests to choose from? We will come back to these questions.

2

Operationalising subjective probabilities

Beliefs matter through their effects on behaviour: from Ramsey and de Finetti to Savage, and the challenges of ambiguity and unawareness.

Pragmatism as the foundation of subjective probability

Peirce, *How to Make Our Ideas Clear* (1878)

"Consider what effects, that might conceivably have practical bearings, we conceive the object of our conception to have. Then, our conception of these effects is the whole of our conception of the object."

Ramsey, *Philosophical Papers*

"The essence of pragmatism I take to be this, that the meaning of a sentence is to be defined by reference to the actions to which asserting it would lead, or, more vaguely still, by its possible causes and effects."

Beliefs have substance only because they have **effects on behaviour**. This effect must be revealed, and represented, through behaviour.

Representing subjective probabilities

Ramsey and de Finetti

- From axioms, reconstruct the probability that the agent (the statistician) assigns to events from **monetary bets** — so utility is assumed known.
- The reverse of von Neumann and Morgenstern (1944), who assume probabilities known and derive the agent's utility.

Problem with this approach

The utility of money.

Savage (1954)

From axioms, derives **both** utility **and** probability.

See also

Anscombe and Aumann (1963): a more restrictive theorem, but a much, much simpler one.

Savage's framework

What the experimenter observes

- S: the set of **states of the world**;
- C: the set of **consequences**;
- F: the set of **acts** f, i.e. functions assigning a consequence to each state of the world: $f(s) = c$ means that act f yields consequence c in state s;
- a **preference relation** $\succsim$ on acts ($f \succsim g$: act f is preferred to act g).

Theorem (Savage, 1954)

If seven axioms (P1–P7) on $(S, C, F, \succsim)$ hold, then there exist a utility function u, unique up to a positive affine transformation, and a unique probability p such that

$$f \succsim g \iff \int_S u(f(s))\,\mathrm{d}p(s) \geq \int_S u(g(s))\,\mathrm{d}p(s).$$

The key axiom: the sure-thing principle

If two acts give the **same consequence** in a state of the world (the "sure thing"), then that state does not explain why the agent chooses one act over the other. $\Longrightarrow$ **Separation of tastes and beliefs.**

Formally

With two states s_1, s_2, the axiom says that for all acts f, f', g, g':

$$\left.\begin{array}{l} f(s_1) = g(s_1) \text{ and } f'(s_1) = g'(s_1) \\ f(s_2) = f'(s_2) \text{ and } g(s_2) = g'(s_2) \end{array}\right\} \Longrightarrow (f \succsim g \iff f' \succsim g')$$

Interpretation

Suppose studying (act f) brings me an income of 100 ...

Criticism 1: the Ellsberg paradox (Ellsberg, 1961)

Urn: 90 balls: 30 red, the others green or blue (distribution unknown). Choose between:

	Red	Green	Blue
Bet A	€100	€0	€0
Bet B	€0	€100	€0
Bet C	€100	€0	€100
Bet D	€0	€100	€100

Observed in the lab

$$A \succ B \implies P(R) > P(G)$$
$$D \succ C \implies P(R)+P(B) < P(G)+P(B)$$
$$\implies P(R) < P(G)$$

Remark: this violates the **sure-thing principle**, with *blue* playing the role of the sure thing.

Response to criticism 1: theories of ambiguity

MaxMin (Gilboa and Schmeidler, 1989)

The agent evaluates each act with the probability that is *least* favourable to it, i.e. the one that minimises her expected utility (see also regret aversion). With $\mathcal{C} \subseteq \Delta(S)$ a set of priors:

$$f \succsim g \iff \min_{p \in \mathcal{C}} \int_S u(f(s)) \, \mathrm{d}p(s) \geq \min_{p \in \mathcal{C}} \int_S u(g(s)) \, \mathrm{d}p(s)$$

Non-additive probabilities (Schmeidler, 1989)

As in Savage, but the probability p is not additive: $p(s \text{ or } s') \neq p(s) + p(s')$. See also Cumulative Prospect Theory (Tversky and Kahneman, 1992).

Smooth ambiguity (Klibanoff et al., 2005)

Represents the agent's attitude towards ambiguity by a function ϕ (with μ a belief over priors):

$$f \succsim g \iff \int_{\Delta(S)} \phi\Big(\int_S u(f(s)) \, \mathrm{d}p(s) \Big) \, \mathrm{d}\mu(p) \geq \int_{\Delta(S)} \phi\Big(\int_S u(g(s)) \, \mathrm{d}p(s) \Big) \, \mathrm{d}\mu(p)$$

Criticism 2: small worlds, imagination and radical uncertainty

With subjective probabilities, the agent can **imagine any state of the world**. But she may be unable to ask herself certain questions.

Examples

- “Do cats grow on trees?”
- “From 1314 to 1677, what was one twenty-ninth of the average temperature on the planets of Proxima Centauri?”
- “Where did I put my keys?” — “Did you look in the car?” — “Oh yes, you’re right!”
- **Radical translation** in Quine (Quine, 1960); see also Davidson’s principle of charity, or the Sapir–Whorf hypothesis.
- What if time were not absolute, but relative to space?

Response to criticism 2: unawareness, or the lack of imagination

- Here **negative introspection** fails: the agent may not know that she does not know — impossible in Savage's model.
- The answer: **unawareness**. Typically, the **state space itself varies**.
- This allows us to think of a counterpart, **reverse Bayesianism**: learning *extends* the set of possibilities rather than shrinking it.

Assessment

A very technical literature, which struggles to establish itself because it is not easy enough (or too easy) to apply or to observe.

A very clear link with Keynes's **radical uncertainty**.

3

Decision theory and experimental economics

Applying decision theory to the experimenter: the framework of Banerjee, Chassang and Snowberg (2017), Bayesian versus randomized experiments.

A missing link between decision theory and experimental economics

Decision theory and experimental economics are regarded as two of the most important fields in economics today (Banerjee et al., 2017, p. 1).

Both cut across many fields

- Decision theory and psychology, decision theory and finance, . . .
- Experimental economics and development economics, labour economics, political science, . . .

BUT no link between the two

Exceptions: Chassang et al. (2012); Kasy (2013); Banerjee et al. (2014).

- An experiment founded on decision theory would be rejected.
- Experiments do not respect the axioms of subjective rationality.

Banerjee et al. (2017) • Decision theory for the experimenter

Decision theory as a framework for the experimenter

- Codifies experimenters' practice.
- In return, the issues raised by experimental economics can shed light on decision theory.

→ **Win–win.**

Main difficulty: an inconsistency between the experimenter's **individual** preferences and her preferences **as a group**.

- Individually, experimenters are **Bayesians**.
- As a group, they are **ambiguity-averse** DMs → they find it imprudent to choose a prior to evaluate new information.

A parallel with the **Ellsberg paradox** (Ellsberg, 1961), see Section 2.

[Editors' note: here decision theory is applied to the *experimenter*; we do not ask whether it applies to the participants.]

Banerjee et al. (2017) • What the framework does

The framework of Banerjee et al. (2017)

- Addresses the decision problem of **choosing an experimental design**.
- Offers **decision rules**.

→ Maximises the experimenter's utility *and* seeks to satisfy an **audience** that can veto the choice.

Reminder

For Banerjee et al. (2017), the goal of an experiment is to produce **evidence on the effectiveness of public policies**. We return below to what is at stake here, and to the different experimental designs that can be considered (their pros and cons).

Banerjee et al. (2017) • Internal and external stakes

The experimenter in this framework = a researcher who must inform about the **impact of a public policy** through an experiment.

Stakes of the decision

- **Internal:** identifying the effects on the experimental population.
- **External:** extrapolating the effects to another population.

$\Rightarrow$ The main methodological issues of experiments in the social sciences, **translated into decision terms**:

- lab experiments and direct control $\rightarrow$ the problem of **artificiality**;
- field experiments and indirect control $\rightarrow$ the problem of **generalisability**.

Banerjee et al. (2017) • Formalising the framework

1. **Actions and preferences.** The decision maker (the experimenter) must decide whether to implement a specific policy. To inform her judgement, she can run an experiment $\rightarrow$ the probability that the policy succeeds is conditional on the treatment and the covariates:
$$p_x^t \equiv \text{prob}\left(Y_i^t = 1 \mid x_i = x\right)$$
2. **Environments and decision problems.** Environments z are defined to distinguish internal from external problems in decision making.
3. **Experiments and decision rules.** The experimenter's decision is twofold: choose an **experimental design** and a **decision rule**, which together translate experimental data into a policy decision.

Banerjee et al. (2017) • First decision: which kind of experiment?

Option 1

Bayesian experiments?

Option 2: randomized experiments?

- They rest on Fisher's statistical framework, famous for its frequentist approach.
- Illustration: *the lady tasting tea*.
- Further reading: Teira (2013).

Warning

The range of possible experimental designs is in fact much wider; the choice considered here is, in the end, quite narrow.

1. Bayesian experiments (1/3) • An illustration

An experiment on education vouchers, meant to guide a school principal's decision

- **Policy decision:** whether or not to introduce such vouchers in a predetermined neighbourhood (an internal problem).
- Earlier, the principal **rejected** the vouchers, stressing that what really mattered for education was the children's socio-economic background.
- His **beliefs have changed**: he is now convinced that schooling is the sole determinant of educational success.
- He runs an experiment in a private school: **a single pupil** takes part.

1. Bayesian experiments (2/3) • Learning from one pupil

→ How can one learn from this experiment? With no control group, how can one draw conclusions?

The Bayesian answer

A Bayesian DM always has a **prior** that she can compare with the result of the single participant.

Prior: a probability distribution expressing someone's belief (here, the principal's) before there is any evidence (see Section 1).

- Suppose the principal believed that a child from a modest social background could not score above the **70th percentile**, and the pupil in the experiment scores in the **90th percentile**.
- The principal compares his belief with the result of the experiment.
- The principal **updates** his belief.
- The principal becomes convinced that vouchers should be given to poor children so that they can attend private schools.

1. Bayesian experiments (3/3) • Limits

- What if there are **several decision makers**, with different priors and beliefs?
- How can other decision makers be **convinced** (e.g. to extend the experiment to other neighbourhoods)?

Two answers

→ The DM may not fully believe her own prior: a problem of **ambiguity** (see Section 2).
→ The audience the DM must convince may reflect the DM's own internal doubts and criticisms.

⇒ The choice of experimental design depends on the amount of available data, the number of participants and, above all, the **weight** the DM attaches to her own beliefs and to those of the audience.

2. Randomized experiments

A brief definition

- Compare two groups: **control** and **treatment**.
- Secure the comparison by **randomly assigning** individuals to each group.

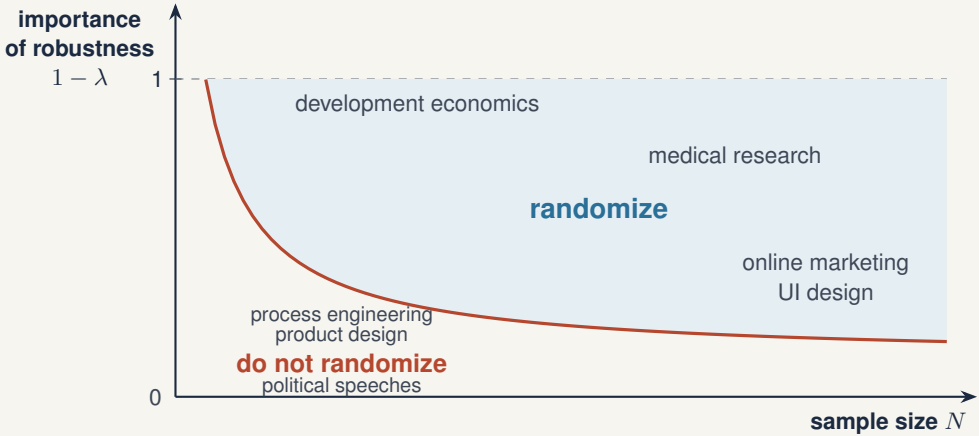
Illustration

Development economics vs medicine.

A methodological gold standard?

- Strong **internal** validity vs weak **external** validity.
- The problem of generalisability.

How to choose between the two kinds of experiment?



Adapted from Banerjee et al. (2017).

Persistent methodological limits

Randomized experiments to the rescue of Bayesian ones

- A way around the problem of **ambiguity**.
- BUT no answer at all to **criticism 2** (unawareness, see Section 2).

Bayesian experiments to the rescue of randomized ones: the hidden message of Banerjee et al. (2017)?

- External validity: the Bayesian approach reintroduces priors.
- Allows **reflexivity** about beliefs.
- **Structured speculation**:
  - who, or what, guides and informs this structure? Only beliefs?
  - how general can this framework be?

Towards **combining**, rather than separating, experimental designs?

Conclusion

Statistics serve plural functions: *for* the researcher and *about* the researcher

- Statistics allow researchers, by creating and collecting data, to say something about the **world**.
- Statistics allow us to analyse the **behaviour of researchers**.

We have tried to approach statistics from both angles

- by combining methods for collecting and creating data with statistical inference (**experiments**);
- AND through **decision theory**, which shows a particular use of statistics through an equally particular theory of probability.

Take-aways

1. Probability can be read as a **frequency**, a **logical relation** or a **degree of belief**; each reading founds statistics differently.
2. Pragmatism makes subjective probability **operational**: beliefs are revealed by choices, and Savage derives utility and probability together.
3. The **Ellsberg paradox** (ambiguity) and **unawareness** (lack of imagination) mark the limits of Savage's model.
4. Banerjee et al. (2017) apply decision theory to the **experimenter**: Bayesian as an individual, ambiguity-averse as a group.
5. Neither Bayesian nor randomized experiments solve every problem: towards **combining** designs.

References I

Anscombe, F. J. and Aumann, R. J. (1963). A definition of subjective probability. *Annals of Mathematical Statistics*, 34(1):199–205.

Banerjee, A., Chassang, S., and Snowberg, E. (2017). Decision theoretic approaches to experiment design and external validity. In Banerjee, A. and Duflo, E., editors, *Handbook of Economic Field Experiments*, volume 1, pages 141–174. North-Holland.

Chassang, S., Padró i Miquel, G., and Snowberg, E. (2012). Selective trials: A principal-agent approach to randomized controlled experiments. *American Economic Review*, 102(4):1279–1309.

de Finetti, B. (1937). La prévision : ses lois logiques, ses sources subjectives. *Annales de l'Institut Henri Poincaré*, 7(1):1–68.

Ellsberg, D. (1961). Risk, ambiguity, and the Savage axioms. *Quarterly Journal of Economics*, 75(4):643–669.

Gilboa, I. and Schmeidler, D. (1989). Maxmin expected utility with non-unique prior. *Journal of Mathematical Economics*, 18(2):141–153.

Keynes, J. M. (1921). *A Treatise on Probability*. Macmillan, London.

Klibanoff, P., Marinacci, M., and Mukerji, S. (2005). A smooth model of decision making under ambiguity. *Econometrica*, 73(6):1849–1892.

Peirce, C. S. (1878). How to make our ideas clear. *Popular Science Monthly*, 12:286–302.

Quine, W. V. O. (1960). *Word and Object*. MIT Press, Cambridge, MA.

Ramsey, F. P. (1931). Truth and probability. In Braithwaite, R. B., editor, *The Foundations of Mathematics and Other Logical Essays*, pages 156–198. Routledge & Kegan Paul, London.

Savage, L. J. (1954). *The Foundations of Statistics*. Wiley, New York.

References II

Schmeidler, D. (1989). Subjective probability and expected utility without additivity. *Econometrica*, 57(3):571–587.

Tversky, A. and Kahneman, D. (1992). Advances in prospect theory: Cumulative representation of uncertainty. *Journal of Risk and Uncertainty*, 5(4):297–323.

von Neumann, J. and Morgenstern, O. (1944). *Theory of Games and Economic Behavior*. Princeton University Press, Princeton.

Wittgenstein, L. (1921). Logisch-philosophische Abhandlung. *Annalen der Naturphilosophie*, 14:185–262. English translation: *Tractatus Logico-Philosophicus*, 1922.