Postwar game and decision theory: a historical perspective

Nicola Giocoli

University of Pisa

nicola.giocoli@unipi.it

Aim of the lecture

- Take two workhorses of modern neoclassical economics (NC econon):
  - ➢ Subjective expected utility theory (SEUT)
  - ➢ Bayesian game theory
- The question I will deal with today is simply: how, when & why did they affirm themselves at the core of NC econ?
- The lecture shows their rise was neither easy, nor immediate. Indeed, in their early days the odds were somehow *against* them.
- This story exemplifies that progress in economics is seldom obvious, nor necessarily grounded on theoretical superiority (whatever this may mean).

Bayesian decision theory

- *Bayesian decision theory* is the dominant theoretical model for both the descriptive and normative analysis of decisions.
- The standard assumption is that individuals follow *Bayesian rationality*, namely, reason probabilistically and act so as to maximize their own subjective expected utility.
- Three central features of Bayesian decision theory:
  1) Represent uncertainty in terms of <u>probabilistic beliefs</u>: uncertainty can always be probabilized → *subjectivist view of probability*.
  2) Beliefs are constrained by <u>consistency requirements</u> → *consistency view of rationality*.
     - ➢ Inconsistent beliefs entail irrationality, but *Bayes's rule* preserves consistency whenever new information arrive.
  3) No need of introspection for eliciting beliefs: just <u>observe agents' behavior</u> under uncertainty → *behaviorist approach*.

Why Bayesian rationality?

- Bayesian rationality is a solid approach, because:
  - … it may argued that any kind of reasoning that takes into account prior information along with basic logical rules and likelihood information is inherently Bayesian;
  - … it fits perfectly with the way 20^{th}-century NC econ has dealt with the problem of choice, i.e., by modeling economic agents as consistent decision-makers.
- Viewed in retrospect, it is fair to say that Bayesian rationality is the culmination of the choice-theoretic approach to rational behavior.
  - "The crowning glory of decision theory" (Kreps 1988).
- But: evidence has emerged in experimental econ, cognitive psychology & neurosciences showing that human beings do not always (to say the least…) use Bayesian reasoning to asses their environment and make predictions.
- Moreover, historically speaking, its success among economists is somehow puzzling. In the late 1950s few would have bet on it!

Puzzle #1

- Bayesian decision theory was born in *statistics*, not economics.
- Our first key question then is: *how, when and why did economic agents come to be modeled as Bayesian ~~decision-makers~~ statisticians*?
- This is seemingly trivial: the easy answer is that credit must be given to L.J. Savage's 1954 axiomatization of SEUT in *The Foundations of Statistics* (*FoS*).
- Yet, this outcome was far from warranted, given Savage's real goal and achievement in *FoS*.
- We need to tell the story of how a self-recognized *failure* eventually became an (unintended) triumph.
- By understanding this story, the "how, when and why?" of Bayesianism in economics becomes a meaningful issue.

Marschak’s remark

- In his 1946 *JPE* review of the *Theory of Games and Economic Behavior*, Jacob Marschak (director at the time of the Cowles Commission) underlined that…
- … John von Neumann had proposed in that book a new notion of rational behavior: rationality in a game required playing mixed strategies; therefore, Marschak noted, “to be an economic man, one must be a statistical man”, …
- … but Marschak also remarked *en passant* that…
- … according to Abraham Wald’s statistical decision analysis, “*being a statistical man implies being an economic man*”.
- As a matter of fact, Savage built SEUT in order to axiomatize *the Wald’s part* of Marschak’s remark, but he actually ended up re-affirming and strengthening the von Neumann’s part!

A new approach to rationality

- Post-WWII NC econ witnessed the emergence of a new meaning of rationality: from rationality as (mere) utility maximization to rationality as consistent behavior.
- The change was motivated by NC economists' embracing a combination of *mathematical formalism* (the axiomatic method) and *logical positivism* (behaviorism, operationalism).
- The focus was on *consistent choices* and on the analytical properties economic agents (or their preferences) need to satisfy to undertake them.
    - ➢ Examples: Samuelson's revealed preference theory; von Neumann's expected utility theory; rationality $\leftrightarrow$ transitive & complete preferences.
- Main consequence for NC econ was to formalize a unified notion of rational decision-makers: economic agents (regardless of their nature: individuals, firms, groups, governments, etc.) came all to be modeled as *consistent choice-algorithms*.
- Yet, this change was not unique to NC econ, but common to other social sciences too.

The "economic" approach to statistics

- The rise of behaviorist statistics: Wald's works in *statistical decision theory* (1939-1950).
- Wald viewed statistical problems as *decision problems*: when faced with the data, statisticians have to decide *what to do* (i.e., how to act), not just *what to say* (i.e., list admissible actions).
  - ➢ Statistics is a guide to action, not just a language to read/interpret reality.
- In general, a *statistical decision problem* (SDP) arises whenever there is a set of alternative choices and the statistician's preference over them depends on an unknown prob distr.
- Basic SDP: the statistician has to choose what "experiment" to perform (say, sampling) in order to extract information from the data & what action to take following the "experimental outcome" (say, making a final decision or continuing "experimentation"). Both choices are "*economic choices*" (= based on costs and benefits).
  - ➢ Wald's basic ideas date back to 1939, but the decisive step came in 1943, out of a very urgent problem, i.e., how to test the quality of war materials (bullet-proof tank armors!).
- He thought that following this *economic* (indeed, *choice-theoretic*) approach would greatly generalize statistics: inference is not anymore restricted to fixed "experiments" and to point/interval testing as in frequentist approach; it becomes a general approach to decision-making under less than complete information about the environment.

von Neumann's minimax rationality

- In a 1928 paper and then, more famously, in the *Theory of Games* (1944) John von Neumann had developed the theory of two-person zero-sum games and proposed <u>the minimax/maximin rule</u> as *the* characterization of rational behavior in a strategic situation.
- The intuition behind the minimax approach is that when you have uncertainty about what your game rival will do and there is no way to get more information (including assigning probabilities), to be rational means to always consider *the worst case scenario* (i.e. your rival is willing and capable to inflict you the maximum loss) and behave so that to minimize the loss you will suffer in such a case (i.e., "play safe"). This *guarantees* you the minimax payoff.
- This approach characterizes strategic rationality *without* any specific assumption on your rival's actions or beliefs, including any use of probability concerning your rival's behavior.
  - ➢ To von Neumann, to be rational meant to be in *full control* of one's own outcome in a game, i.e., to nullify the impact of strategic uncertainty.
  - ➢ It is a kind of "rationality of actions" (Montesano 1998): rationality as "choosing according to a criterion", which simply means following consistently a criterion and being able to justify it.

Wald’s brilliant idea

- Wald’s puzzle: how to solve statistical decision problems (SDPs)?
- The solution: extend von Neumann’s reasoning to SDPs!
- Any SDP can be reconfigured as a game played against a hostile opponent called Nature → games against Nature.
- Nature “plays the game” by randomly selecting the *true value* of a prob distr.
- Winning the game (= solving the SDP) entails figuring out what this true value is.
- The statistician’s rational way to play is, as in von Neumann, to minimize the cost of the mistake she will inevitably make by guessing that value wrong.
- The bottom line: *statisticians must behave as rational agents in a game*; hence, they should solve their SDPs by choosing the minimax action.

Bayesianism kicks in

- Wald defined a risk function $r(F, \delta)$ measuring the cost of following decision rule δ when the true distribution is *F*. This cost is made of two parts:

  i) the *expected loss* of wrongly acting on the basis of δ when *F* holds,

  ii) the *expected cost of experimentation*, i.e., of collecting information about *F*.

- Choosing the minimax solution of a SDP means choosing δ such as to minimize the maximum value of the risk function (= minimize the maximum cost).
- Wald demonstrated that…
- … every SDP is equivalent to a two-person zero-sum game à la von Neumann, where Nature is the statistician's opponent;
- … the solution to a SDP can be found via the minimax and Bayes's rules:
  - all minimax solutions of SDPs are also Bayes's solutions;
  - given any a priori distribution, any solution of SDPs must be among Bayes's ones.
- However, Wald was *not* a true Bayesian: existence of a priori distributions, Bayes's rule and minimax solutions were *mere math tools* $\rightarrow$ no philosophical commitment to them.

Wald’s “heir”

- Wald’s approach was very successful at the time: it promised to place statistics on firmer operational grounds as the new science of decision-making under uncertainty → Arrow (1951).
- Unfortunately, Wald died in an airplane crash in December 1950.
- Savage set out to accomplish his unfinished project.
- His main idea was: let’s take Wald seriously!
  - If statisticians must act like rational economic men, one has first to define exactly *what is* rational behavior → von Neumann’s view is too limited.
- Historically speaking, it would be as a mere by-product of this idea that rational economic men would end up being modeled as Bayesian statisticians.

Savage's *Foundations of Statistics*

- Savage's 1954 book was the declared continuation of what Wald had started.
- *FoS* targeted statistics, *not economics*: its goal was to fulfill Wald's project by improving the behaviorist foundations of orthodox statistical techniques.
- Note that Savage aimed neither at revolutionizing statistics (he, like Wald, just wanted to generalize usual inference tools), nor at describing how economic agents really behave.
- His decision theory was explicitly *normative*: he just wanted to teach statisticians how they *should* behave to solve their SDPs.
- Savage's two-part program in *FoS*:
  - ➢ First, develop a rigorous decision theory under uncertainty as a normative benchmark for statisticians.
  - ➢ Second, reinterpret/extend orthodox inference methods according to the new theory.
- SEUT was the outcome of the first part.

What SEUT was for

- The *axiomatic method* was at center stage of *FoS*'s first part.
  - ➢ Like in von Neumann's EUT, the task was to select a list of axioms and impose that agent's choices under uncertainty respect them.
- Savage's task was to develop a *normative guide* for the "good statistician": *a rational statistician* is one who forms consistent beliefs, i.e., whose beliefs satisfy the axioms and are revised according to Bayes's rule.
- *SEU theorem*: rational choice under uncertainty can be explained as if it has been made according to the expected utility rule with the attribution of a numerical subjective probability to agent's beliefs.
  - ➢ To derive the theorem Savage had to solve several technical problems (mainly, to transform qualitative beliefs into numerical ones). His ingenious solutions would stimulate a huge literature afterwards.

Subjectivist, consistent, behaviorist… and axiomatic

- SEUT satisfies all three characteristics of Bayesian rationality:
  1) SEUT represents uncertainty in terms of *probabilistic beliefs*.
  2) Beliefs need be constrained by *consistency requirements*.
  3) Beliefs are elicited *operationally*.
- Generally, modern decision theory characterizes rationality as “rationality of preferences” (Montesano 1998), i.e., as “a preference system satisfying a set of formal properties (axioms)”.
- Theorists look for rules to pass from the preferences and expectations over the outcomes of the agent’s actions to the preferences over the actions themselves.
- Any such rule is said to be rational if it satisfies a pre-assigned list of axioms.
- Savage’s theory is one of those rules.
  - ➢ So it embodies a different kind of rationality than von Neumann’s (and Wald’s) “rationality of actions” → Savage’s “silent shift” from one kind of rationality to another.

The achievement

- Savage's achievement in developing SEUT is undisputable.
- Economists like Samuelson, Arrow, Friedman, etc., praised it immediately.
- Even before the publication of *FoS* (e.g. Paris 1952 conference), SEUT became a benchmark for theoretical and empirical discussion *among economists* about what it means to be rational.
- So Kreps's "crowning glory" acclaim is well deserved.
- But…

Savage’s fiasco

- … but recall Savage’s real goal in *FoS*.
- He aimed at fulfilling Wald’s program, i.e., transform orthodox inferential statistics into a behaviorist, choice-based discipline.
- With SEUT he had accomplished the first part of his project, i.e., develop a complete characterization of rational choice under uncertainty.
- Unfortunately, in the rest of the book Savage *failed* to reinterpret orthodox statistical techniques in subjectivist/behaviorist terms.
- Indeed, *he proved that this could not be done* and admitted that such a surprising outcome destroyed the whole *FoS*’s program!
- A brand new kind of statistics was called for if the discipline was to be grounded upon rigorous subjectivist/behaviorist basis: this opened the door for his *later* (= 1960s) championing of Bayesian statistics.
- The failure explains why today the second part of *FoS* is totally neglected. But even *at the time* Savage’s fiasco was already clear and explicit. He himself openly recognized it!
- How could a part of such a botched project become a pillar of NC econ?

Summing up doesn’t make the total right

- *Consider* that SEUT was developed as a *normative* theory to teach *statisticians* how to solve their *statistical* decision problems, not to describe how generic decision-makers, like me and you, actually behave.
- *Add* that in terms of its most urgent and relevant goal, i.e., making orthodox inference techniques more robust, *FoS* was a total failure.
- *Add* that, even as the trigger of the Bayesian revolution in statistics, *FoS*’s success was neither immediate (not even for Savage himself!) nor complete: Bayesian statisticians remained a (significant) minority with respect to their orthodox colleagues.
- *Add* that SEUT was hardly the only option available for 1950s NC econ. Several alternatives existed, including a few decision criteria capable of also handling cases of so-called *complete ignorance* (i.e., the Holy Grail of “true” uncertainty, à la Knight & Keynes).
- *The end result is*… our question once again! Why the Bayesian turn in economics? Why did SEUT became the orthodox way of *describing* economic agents’ behavior? Why, to reword Marschak’s remark, “to be an economic agent came to imply being a *Bayesian* statistician”?

The moral of the story (part I)

- What this story tells us is that the history of economic ideas seldom proceeds linearly: an established argument today may not have been considered such when it was first proposed.
- How does economics pick winners? Is it just a matter of sheer theoretical superiority (whatever this may mean)? Or is it a matter of superior descriptive power?
- In the case of SEUT, it can be argued that the answer is negative under both respects: Savage's approach was *theoretically and descriptively weak*.
  - ➢ On the former weakness, why should economics attribute such a prominent role to a (admittedly, solid) building block of a theoretical edifice that collapsed from day one?
  - ➢ On the latter weakness, just think of Ellsberg's paradox.
- I have no sure answer to the "how, when & why" question. What really matters is to raise it!
- Still, I may suggest three, not mutually exclusive conjectures. Each of them corresponds to one of the three possible motivations for progress in economics.
  - ➢ Note: empirical validation is *not* among them → "No mere fact ever was a match in economics for a consistent theory" (Milgrom & Roberts 1987).

Three paths to progress

- Follow the money
- Defend your bastion
- A quest for elegance & integrity

Follow the money

Conjecture #1. The spread of Bayesian decision theory in major US business schools.

- The spread of SEUT has been helped by the new 1950s-1960s approach to business education as *a formal social science*, rooted in quantitative analysis (e.g. operations research) and behaviorist disciplines. This approach was first proposed in the business schools of Harvard and Carnegie, and, later, Stanford, Northwestern and Chicago among others.
- The key was operationalism: “You can’t teach frequentist stuff to businessmen! They *have to* make decisions under uncertainty” (Howard Raiffa, interview in Fienberg 2008).
  - ➢ Same issue as Wald’s tank armor testing: he *had* to decide whether it was bullet-proof.
- SEUT is very operational: it gives green light to decisions made following managers’ informed/educated guess. So it quickly became very popular in business schools.
- This conjecture captures the first path to success in economics.
- Winning econ ideas are those that find direct (and often remunerative) applicability in the real world, regardless of their intrinsic validity: marketable ideas!

Defend your bastion

Conjecture #2. SEUT as a much-awaited-for support.

- Expected utility theory (EUT) has a long history. It dates back to 17th century (Blaise Pascal's wager); economists have regularly used it since mid-18th century (e.g. Ferdinando Galiani).
- When von Neumann's axiomatized EUT, many economists viewed this as his most important achievement in the *Theory of Games*.
  - ➢ Which of course it was *not* → another instance of economists' cherry-picking attitude...
- But EUT was criticized for being based on unavailable *objective* probs & for reviving a non-operational notion of "cardinal utility" (which NC econ had abandoned in the 1930s).
- While the latter critique was quickly defeated, the unavailability of probs was a serious problem that undermined a pillar of mainstream micro → "no sound criteria exist for behavior under uncertainty, hence it is basically random" (Alchian, *JPE* 1950).
- Though part of a failed project, Savage's SEUT filled the gap in a convenient way: it showed that *subjective* probs could be used, provided they were policed by consistency rules.
- This exemplifies the second path to success: winning ideas are those useful to protect the mainstream, i.e., to preserve its theoretical core from dangerous critiques: Lakatos *docet*!

A quest for elegance & integrity

Conjecture #3. *The consistency view of rationality*.

- SEUT was welcomed by NC economists as *the accomplishment of the consistency approach* to rationality, itself a key ingredient of the so-called formalist turn in economics.
- In his revealed pref theory Samuelson had talked of "agents as consistent guinea pigs". From there it was just a small step to Savage's "agents as consistent gamblers".
  - ➢ Small, at least, once Savage had demonstrated that his subjectivism did not entail a return to unobservable mental entities, but was rather based on *observable choices*.
- Of course, choices were *not really* observable (experimental econ was yet to come!). The real motivation behind the whole consistency approach to rationality was the pursuit in economics of the same formal rigor typical of the axiomatic approach to mathematics.
  - ➢ The *mathematician*, rather than the physicist, was the economists' new role model.
  - ➢ Savage's axiomatic SEUT was perfectly attuned with the new formalist paradigm.
- This is the third path to success: winning ideas are those that fit perfectly with (or bring to completion) an underlying general view of what econ is about and/or what it means to be an economist: <u>the seductive power of the "spirit of system"</u> (cf. Adam Smith, *WN,* Bk IV).

Puzzle #2

- Nash equilibrium is “the embodiment of neoclassical rationality” (Aumann 1985).
- If so, why were noncoop game theory in general, and Nash eq in particular, almost completely ignored (when not expressly *rejected*) by NC economists in the 1950s & 1960s?
  - Note: I refer here to game theory as the analysis of rational strategic behavior, not as a useful tool box of math techniques for other areas of economic analysis (most notably, Arrow/Debreu’s 1954 use of Nash eq in proving existence of GEE).
- And, as a related question, why the late 1970s-early 1980s boom in game theory?
- How, when and why did game theory conquer the core of NC econ?
  - Two landmark textbooks certifying the conquest: Tirole 1988; Kreps 1990.
- Here I will only focus on the “boom” part of the puzzle.

Cherchez Bayes… once again

- The turning point was John Harsanyi's 1967-68 three-part paper "Games with incomplete information played by Bayesian players" published in… *Management Science* (remember the business schools?).
  - Possibly, the single most important game-theoretic work after von Neumann & Nash.
- The paper started Bayesian game theory, i.e., the theory of (realistic) games where agents have private information and form consistent beliefs about other players' information.
  - Other prominent game theorists were involved from the beginning in Harsanyi's project: Robert Aumann, Michael Maschler, Reinhard Selten.
- The issue was: how to model players' beliefs about each others' beliefs?
  - Assume a parameter of the game is unknown; player 1 formulates a belief about it; player 2 then formulates a belief over player 1's belief about the parameter; player 1 then formulates a belief over player 2's belief over his own belief; etc. → an infinite hierarchy of beliefs.
  - Harsanyi's answer: uncertainty is not about parameters' values, but *about the information other players have about them*; this information can be represented by a random variable; the model must then include probability distributions describing what players subjectively believe about these random variables; these probabilistic beliefs are policed by Bayes's rule → Bayesian game theory could begin!
- This answer helped the development of the so-called *Nash refinement literature*.
- But the importance of Harsanyi's work goes much beyond that literature…

Harsanyi’s real breakthrough

- Harsanyi’s paper brought a fully-fledged revolution in economics.
  - Triggering a theoretical revolution is a *fourth*, different path to success.
- What is a “revolution”? A big methodological question.
  - A paradigm change à la Khun? In this sense, Harsanyi’s was *not* a revolution: the old pillars of NC econ (rationality, equilibrium, maximization, etc.) were still there.
  - A significant transformation/expansion in the set of problems your theory can deal with? In this sense, Harsanyi’s was a true revolution because…
- …by developing a general methodology to model games where players have private information, Harsanyi effectively started *modern information economics*.
  - “Before 1967, the lack of a general framework for informational problems had inhibited economic inquiry about markets where people do not share the same information. The unity and scope of modern information economics were found in Harsanyi’s framework” (Myerson 2008).

“Two sides of the same coin”

- Bayesian game theory made economists realize that in strategic situations “uncertainty and private information are essentially two sides of the same coin” (Marinacci 2015):
  - ➢ uncertainty generates private info (when agents have different info about a phenomenon);
  - ➢ private info generates uncertainty (think of adverse selection or moral hazard problems).
- So, explaining game theory’s boom with the Nash refinements literature is only a partial answer: it downplays the extent of the Harsanyi revolution.
- The true cause of the boom was the tremendous breadth and salience of information econ.
- Almost everything in modern applied game theory descends from that 1967-68 paper.
  - ➢ Think e.g. of mechanism design theory (MDT).

What if…?

- Still in the late 1960s very few economists cared about game theory. Once again, the odds were strongly *against* Harsanyi's success. However…
- The game-theoretic boom came when Bayesian decision theory was *already* the established approach to NC decision-making under uncertainty.
- Had Bayesian decision theory *not yet* been judged a proper approach to modeling rational behavior, it is indeed likely that Harsanyi's paper would never enjoy much success.
- It may thus be argued that the rise of the new view of rationality ("the economic man as a Bayesian statistician") also created the conditions for the subsequent boom of Bayesian game theory.
  - ➢ So even truly theoretical revolutions may in the end be just an indirect offspring of the other three (more mundane…) explanations of econ progress.
- We are back to square one: what made Bayesian rationality so special?

The moral of the story (part II)

- In their influential *Games and Decisions* (1957), Duncan Luce and Howard Raiffa concluded:
  1) that Savage's SEUT was only good for situations of *partial ignorance* (= when the agent has *some* beliefs about the state of the world), which are intermediate between those of pure risk and those of complete ignorance (i.e., "true" uncertainty);
  2) that ways to axiomatize decision theory under complete ignorance already existed, providing a *more general* characterization of rational behavior.
- Still, NC economists endorsed SEUT wholeheartedly and never turned back (so far…).
- What would modern micro look like if economists had embraced one of these more general (though less operational) alternative approaches to rationality rather than the Bayesian one?
- I don't know, but the moral for econ students should be clear: never forget the *roads not taken* and the *many different motivations* behind winning ideas in economics.

Thank you!

Read my research at:

ssrn.com/author=92886

pisa.academia.edu/NicolaGiocoli

www.researchgate.net/profile/Nicola_Giocoli

References

- Alchian A. 1950, «Uncertainty, evolution and economic theory», Journal of Political Economy, 58:3, 211-221.
- Arrow K. 1951, «Alternative approaches to the theory of choice in risk-taking situations», *Econometrica*, 19:4, 404-437.
- Basili M. & Zappia C. 2009, "Shackle and modern decision theory", *Metroeconomica*, 60:2, 245-282.
- Chernoff H. 1954, "Rational selection of decision functions", *Econometrica*, 22, 422-443.
- Ellsberg D. 1961, "Risk, ambiguity, and the Savage axioms", *Quarterly Journal of Economics*, 75:4, 643-669.
- Fienberg, S.E. 2008, The Early Statistical Years: 1947-1967. A Conversation with Howard Raiffa, *Statistical Science*, 23:1, 136-149.
- Harsanyi, J.C. 1967-68, "Games with incomplete information played by 'Bayesian' Players, Parts I-III", *Management Science*, 8, 159-182, 320-34, 486-502
- Kreps D. 1988, *Notes on the Theory of Choice*, Westview Press.
- Kreps D. 1990, *A Course in Microeconomic Theory*, Princeton: Princeton UP.
- Luce R.D. & Raiffa H. 1957, *Games and Decisions*, New York: Wiley.
- Marinacci M. 2015, «Model uncertainty», *Journal of the European Economic Association*, 13:6, 998-1076.
- Marschak J. 1946, "Neumann's and Morgenstern's new approach to static economics", *Journal of Political Economy*, 54:2, 97-115.
- Milgrom P. & Roberts J. 1987, "Informational asymmetries, strategic behavior, and industrial organization", *American Economics Review*, 77:2, 184-193.
- Montesano A. 1988, "Rationalty in economics: a general framework", *Proto Sociology*, 12, 290-296.
- Myerson R. 2008, "John C. Harsanyi (1920-2000)", in: Durlauf S.N. e L.E. Blume (eds.), *The New Palgrave Dictionary of Economics*, second edition.
- Savage L.J. 1954, *The Foundations of Statistics*, New York: Wiley.
- Schmeidler D. 1989, "Subjective Probability and Expected Utility without Additivity", *Econometrica*, 57:3, pp. 571–578
- Smith A. 1904 [1776], *An Inquiry into the Nature and Causes of the Wealth of Nations*, http://www.econlib.org/library/Smith/smWN.html
- Tirole J. 1988, *The Theory of Industrial Organization*, Cambridge, Ma: MIT Press.
- von Neumann J. & Morgenstern O. 1953 [1944], *Theory of Games and Economic Behavior* (3rd ed.), Princeton: Princeton UP.
- Wald A. 1950, *Statistical Decision Functions*, New York: Wiley.

Bayes's rule

$$P(A|B) = \frac{P(B|A)P(A)}{P(B)}$$

where:

- *A* and *B* are events and $P(B) \neq 0$;
- *P(A)* (the "prior") and *P(B)* are the independent probabilities of observing A and B;
- *P(B|A)* is the probability of observing event B given that A is true (the likelihood);
- *P(A|B)* is the probability of observing event A given that B is true (the "posterior").
- Example: one of the tablets produced by company X is found defective; what is the prob that it has been produced by factory Z rather than by any other of X's factories?
  - Prob of event A: pct. of tablets being produced by factory Z (the prior): 66%
  - Prob of event B: pct of defective tablets in company X's total production: 3,2%
  - Likelihood: pct. of defective tablets out of all those produced by factory Z: 1,8%
  - The posterior is: 0,018 * 0,66 / 0,032 = 0,371 → there is a 37% posterior prob that the defective tablet comes from factory Z (against an a priori prob of 66%).

John von Neumann
(1903 – 1957)

Jacob Marschak
(1898 – 1977)

Abraham Wald
(1902-1950)

Leonard "Jimmie" Savage
(1917 – 1971)

An early 1950s description of a SDP

- "The statistician typically finds himself in situations more similar to that of the businessman. The problem of statistics can be formulated roughly as follows:
- It is known that one out of a number of hypotheses about a given situation is true.
- The statistician has the choice of one of a number of different experiments […] (drawing a sample of any size can be included in this scheme), the outcome of any one of which is a random variable with a probability distribution depending on which of the unknown hypotheses is correct.
- On the basis of that outcome, the statistician must take some action (accept or reject a hypothesis, estimate the mean of a distribution to be some particular value, *accept or reject a lot of goods*, recommend a change in production methods, etc.) the consequences of which depend on the action taken and on the hypothesis which is actually true. […]
- The various theories which have been proposed from time to time as foundations for statistical inference are therefore closely related to theories of economic behavior under uncertainty" (Kenneth Arrow, *Econometrica*, 1951)

How to elicit beliefs – the intuition

- By observing how the decision-maker chooses between two gambles we can infer her beliefs about the relative likelihood of two alternative states.
  - This intuition dates back to Frank Ramsey and Bruno de Finetti in the late 1920s - early 1930s.
- Take two mutually exclusive states *s1* and *s2* on whose occurrence the decision-maker has no available objective probabilities.
- Take two simple gambles, one paying 1$ if state *s1* occurs and zero if *s*2 occurs and one paying 1$ if *s2* occurs and zero if *s1* occurs.
- If the agent chooses the first gamble in preference to the second, then she reveals that she considers state *s1* subjectively more probable than state *s2* (= *revealed probabilistic beliefs*).
- Subjective beliefs could in general be inconsistent, but irrationality is avoided thanks to Bayes's rule, which guarantees that consistency is always preserved whenever new information arrives.

Some of the alternatives

- The postwar literature on decision-making under uncertainty is huge. Many criteria were suggested. For example…
- …Wald's maximin criterion
- …Savage's minimax risk (regret) criterion
- …Leonid Hurwicz's Pessimism-Optimism index
- …the Principle of Insufficient Reason
  - ➢ Note: all these could be given a *rigorous axiomatic foundation* (Chernoff 1954).
- …more heterodox approaches, like G.L.S. Shackle's non-probabilistic approach
  - ➢ Shackle's critiques have been formally answered only very recently, by Gilboa and Schmeidler's non-additive probability approach (Schmeidler 1989; Basili & Zappia 2009)

Ellsberg paradox

- Take two urns. Urn I contains 100 red and black balls, in an unknown ratio. Urn II has 50 red and 50 black balls.
- Agents must choose an urn to draw from, and bet on the color that will be drawn. They will receive $100 if that color is drawn, and $0 if the other color is drawn.
- Subjects must decide whether they would rather bet on…
    1. A red draw from Urn I, or a black draw from Urn I
    2. A red draw from Urn II, or a black draw from Urn II
    3. A red draw from Urn I, or a red draw from Urn II
    4. A black draw from Urn I, or a black draw from Urn II
- One would expect subjects to be indifferent in cases 1 and 2, and indeed experiments show that they are.
- However, it turns out that agents uniformly prefer a draw from Urn II in cases 3 *and* 4.
- But these choices make it impossible to infer the agents' probabilistic beliefs → Savage's SEUT axioms are violated!
- Do people regard a draw of a particular color from Urn I as more likely? Certainly not, otherwise they would not choose Urn II in both cases 3 and 4. If agents choose Urn II in case 3, this implies that they believe (rightly or wrongly) that Urn I has more black balls than red. However, if that is their belief, then they ought to choose Urn I in case 4, but they don't.
- Urn I *may* have more black balls than red, but it *may* also have more red balls than black, so agents simply tend to avoid it.
- The primary conclusion one draws from this experiment (first proposed by Daniel Ellsberg in 1961) is therefore that people prefer *definite* to indefinite information, i.e., they hate ambiguity more than pure risk.

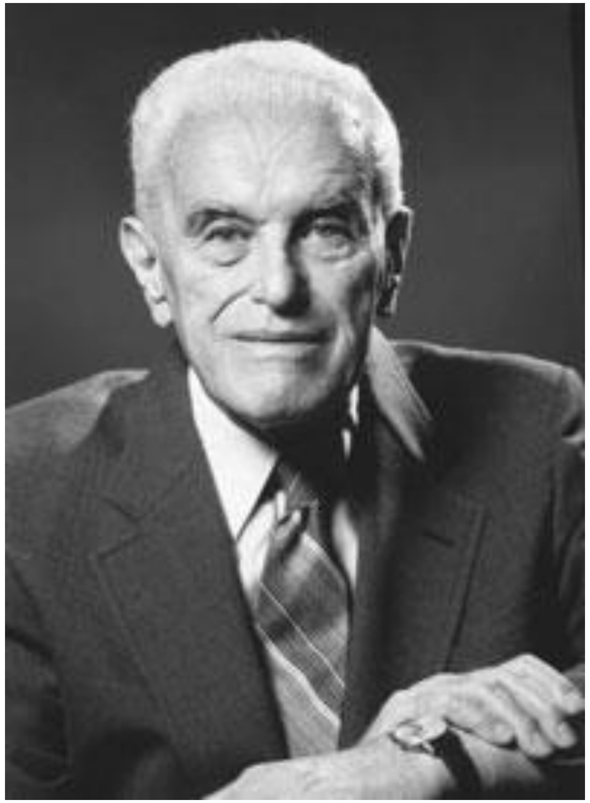

John C. Harsanyi
(1920-2000)

The refinements literature

- Nash eq is just a *necessary* condition of rational strategic behavior, not a sufficient one.
- Two main problems:

1. Multiplicity of Nash equilibria (with some "irrational" equilibria).
   - Hence the goal of eliminating unreasonable equilibria.
2. Games in extensive form (EF) descending from the same normal form (NF) may have different equilibria.
   - Hence the goal of ensuring consistency between NF and EF solutions.

- A refinement of Nash eq is a more accurate characterization of rational strategic behavior that (tries to) solve both problems.
- Harsanyi's 1967-68 paper was a milestone for this literature because it allowed the interpretation of EF games as Bayesian games.
- The Nash refinement program attracted to game theory a bunch of brilliant math economists at the same time they began to abandon general equilibrium theory.
- Lots of new results followed: the game theory "industry" was born.
- So, is this the reason for the game-theoretic boom in the late 1970s – early 1980s?

Mechanism design theory

- MDT is at center stage of modern micro: it is a general approach to the engineering of socially desirable economic institutions and/or incentive systems in an environment populated by interacting rational agents.
- Bayesian game theory has been crucial in developing MDT.
  - MDT's key results (like the revelation principle) were derived in the early 1980s within the framework of Bayesian games and using Bayesian Nash equilibrium as the key solution concept.
  - Since then, MDT problems have been formulated in Bayesian terms, i.e., following Harsanyi's modeling strategy.
- Starting from the early 1980s any economist wishing to contribute to MDT – which means more or less *everywhere* in today's micro – has first to learn how to model and solve a Bayesian game.
- This further contributed to place game theory firmly at the core of modern micro.