

Rationality under uncertainty: classic and current criticisms of the Bayesian viewpoint

This is the peer reviewed version of the following article:

Original:

Zappia, C. (2018). Rationality under uncertainty: classic and current criticisms of the Bayesian viewpoint. EUROPEAN JOURNAL OF THE HISTORY OF ECONOMIC THOUGHT, 25(6), 1387-1419 [10.1080/09672567.2018.1523210].

Availability:

This version is available <http://hdl.handle.net/11365/1071254> since 2021-04-03T11:48:19Z

Published:

DOI: <http://doi.org/10.1080/09672567.2018.1523210>

Terms of use:

Open Access

The terms and conditions for the reuse of this version of the manuscript are specified in the publishing policy. Works made available under a Creative Commons license can be used according to the terms and conditions of said license.

For all terms of use and more information see the publisher's website.

(Article begins on next page)

Rationality under uncertainty: classic and current criticisms of the Bayesian viewpoint*

Carlo Zappia (University of Siena)

1. Introduction

In economics the content of rational choice theory has not been disputed for decades. At least since the late 1940s the received view of rational choice has been Expected Utility [EU] theory. This view emphasizes the normative content of the theory, concentrating on what an ideal rational agent should do in every decision context, and enlarging the field of application of decision theory from certainty to risk and uncertainty (Sugden 1991). Such a theory is grounded on axioms that shape the rationality of agents as free of psychological assumptions or introspection, but insist on an operational view, providing continuity from Samuelson's (1938) revealed preference theory to Savage's (1954) Subjective Expected Utility [SEU] theory (Giocoli 2003).

Drawing on the works of Frank Ramsey (1964 [1931]) and Bruno de Finetti (1964 [1937]), the mathematical statistician Leonard Savage established the received view in his *Foundations of Statistics* (1954). Savage's contribution was meant to extend von Neumann and Morgenstern's (1947) approach, devised for taking into account decisions whose consequences depend on the actions of other agents in their *Theory of Games and Economic Behaviour*, but soon adapted to decisions concerning risky prospects such as lotteries. Savage used subjective probabilities to represent an individual's degrees of belief in the future realization of all kinds of external, pay-off relevant events—in principle, even singular events—and proved that the maximization of EU can be derived from a series of axioms on actions to be taken in uncertain contexts.

Savage's procedure to concentrate on the equivalence of maximization of EU with compliance with a set of axioms set a new standard. Indeed, in the years following von Neumann and Morgenstern's introduction of EU the axiomatic approach to decision-making flourished. It was on the basis of a series of postulates, explicitly intended to extend "rules of conventional logic ... into the realm of decision," that Marschak (1950) defended the maximization EU when making decisions.

* This is the preliminary draft of a paper presented at the Annual Meeting of the European Society for the History of Economic Thought - ESHET, University of Antwerp (May 2017) and at Annual Meeting of the Associazione Italiana per la Storia dell'Economia Politica - STOREP, University of Piacenza (June 2017). Comments by Andrea Salanti and Erik Schokkaert are gratefully acknowledged. A revised version of this paper is forthcoming in the *European Journal of the History of Economic Thought*.

Savage developed further this idea of extending logic, providing the “as if” justification for EU maximisation even under uncertainty.

After Savage’s axiomatization, at least among decision theorists, but increasingly in the entire economics profession, rationality has been equated with a theory prescribing what individuals should do in view of a series of foundational axioms. This normative turn in decision theory—emphasizing the prescriptive component of theoretical analysis and moving beyond the normative as an ethically motivated viewpoint (Hands 2015)—incorporates a specific notion of rationality. Rationality has been associated with the idea of behaving in a consistent way, regardless of the actual content of the decision to be made, and the decision theoretic approach built on it has been termed Bayesian.¹

The view that rationality implies Bayesianism has been put under scrutiny in recent years. A main line of research criticizing this view has concentrated on the need to incorporate behavioural assumptions, given the overwhelming evidence suggesting that the Bayesian approach is descriptively inadequate. The prototype behavioural theory, Kahneman and Tversky’s (1979) prospect theory, was presented as an alternative formal theory of behaviour intended to describe the choices that individuals actually make. Yet, it was not proposed as a substitute for the normative mainstream, and much of the following behavioural economics literature has not aimed at a new normative ideal either (Kahneman and Thaler 2006). As repeatedly noted in the literature, prospect theory did not intend to challenge the normative value of Bayesian decision theory (Heukelom 2014).

This paper concentrates instead on a second line of criticism, that addressing the normative significance of the axioms that define rationality in the Bayesian approach. In particular, the notion of consistency to be used by rational agents has been questioned in contemporary decision theory. As summarized in a recent methodological paper by decision theorist Itzhak Gilboa (2015), though orthodox economists may keep denying that violations of the standard model can be attributed to rational decision-makers, it has become increasingly evident that the significance of certain kinds of violations cannot be dismissed as only of descriptive relevance: violations have been so substantial as to make no longer obvious which notion of consistency should be used to define rationality. This is a highly relevant viewpoint in current research in decision theory since it can be attributed to a substantial group of scholars working on decision-making who aim to establish a new normative

¹ The term Bayesian relates to the use of Bayes Theorem in the updating of probability priors, but the reference here mostly is to the assumption that individuals are supposed to be endowed with a single, definite probability prior in any possible instance.

standard for choice under uncertainty.² And although such a literature of non-Bayesian models, also referred to as “ambiguity aversion models,” has been criticized in its turn—see in particular Al-Najjar and Weinstein (2009) and Hammond and Zank (2014)—it is of special prominence in the contemporary debate on rational choice theory.

This perspective, criticizing the normative value of the Bayesian viewpoint, has emerged from an analysis of the so-called Ellsberg Paradox. As we shall see, urn examples of the kind suggested by Daniel Ellsberg (1961) evidenced the inability of a significant number of “reasonable” decision-makers to commit to a sharp probability function, specifically violating the additive property of the probability function that Bayesians assume can always be elicited from choices. Ellsberg suggested that these violations are “deliberate,” and that they challenge the prescriptive, rather than the descriptive, content of the underlying choice theory. But while the Ellsberg Paradox is amply referred to in the current literature—indeed it is the starting point of most studies addressing decision under uncertainty—there usually is no analysis of the kind of criticism Ellsberg made to Savage, in what can be characterized as the first, *classic* criticism of the Bayesian viewpoint.

This paper provides a comparison of Gilboa’s *current* critique of Bayesianism with Ellsberg’s 1960s critique of Savage’s understanding of rationality. After a summary of Gilboa’s viewpoint about the arbitrariness of the Bayesian approach and his suggested alternative to it, the paper concentrates on the Ellsberg Paradox. The paper illustrates Savage’s own view of the methodological issues intertwined with his approach and Ellsberg’s criticism of it, with specific regard to the notion of rationality as consistency promoted by Savage’s Bayesian approach. The paper shows that Ellsberg’s view—partly hidden in his doctoral thesis and therefore largely unknown—consists of a generalization of the Bayesian viewpoint, aimed at presenting a normative theory for decision contexts that are too ambiguous to be interpreted through sharp probability priors.

A comparative assessment of the classic and current criticisms reveals that Ellsberg’s analysis of Savage’s axioms thoroughly anticipated today’s criticism of Bayesian rationality. In particular, we shall see that both Gilboa’s suggested definition of rationality and the related procedure to test it against empirical evidence were already formulated in Savage’s proposal and dissected in Ellsberg’s critique. Indeed, the early 1960s dispute among the two decision theorists already hinged on the two main aspects that constitute the methodological bases of the current critique of Bayesian rationality: first, a specific interpretation of rationality—that a decision is rational for a decision-maker if, when exposed to the analysis of her choices, she does not regret having made it—and, second, a procedure

² Notable contributors include Eichberger and Kelsey (1999), Gilboa and Schmeidler (2001), Gilboa, Postlewaite and Schmeidler (2008, 2009, 2012), Gilboa and Marinacci (2016), Hansen (2014), Machina and Siniscalchi (2014), and Mukerji (2009).

as to how the validity of normative assumptions should be tested—that is, by taking into account only decision-makers’ reflective behaviour.

2. Current criticism of the Bayesian viewpoint

As is well-known, Bayesianism holds that every kind of uncertainty—and not only risky situations such as lottery tickets for which an objective probability distribution is available—should be quantified with probability over the state space. In the absence of objective probabilities, the decision-maker should use her own, subjective probabilities to guide her decision. Rather than rely on introspection or logical consideration of symmetry—as in “intuitive” approaches to probability such as those identified by Keynes (1973 [1921]) and Jeffreys (1948)—the probability function measuring the individuals’ subjective beliefs in the realization of payoff relevant events is elicited from an analysis of “actual” choices.³ In actual decision settings where the choices made by individuals may be observed, including those in laboratory environments, a decision over actions whose consequences are not certain, but conditioned on the realizations of external events, reveals a degree of belief in the event.⁴

On this view, the laws of probability provide the constraints on the way the individual should form beliefs for making decision when information is incomplete and update them when new pieces of information are made available. Mainly, attributing probabilities to any kind of event guarantees the kind of consistency necessary both to avoid accepting losing prospects—as in the Dutch Book example—and to satisfy dynamic consistency—by means of Bayes’s rule of updating of beliefs in view of information gathering. On these grounds, rational choice theory under uncertainty simply

³ In Ramsey’s words (1964 [1931], p. 72), giving birth to the operational perspective characterizing the Bayesian approach, a degree of belief can be expressed as “the extent to which we are prepared to act on it” (see also de Finetti 1964 [1937], p. 102). Savage’s claim that preferences are defined by choices was formulated in term of acts, i.e. decisions whose consequences depend on the realization of exogenous events. That a decision-maker prefers act *f* to *g*, simply means that “if he were required to decide between *f* and *g*, no other act being available, he would decide on *f*” (Savage 1954, p. 17).

⁴ This is a form of radical behaviourism, in which the psychology of choice is summarized by the idea that mental states are distinguishable only in terms of behaviour. As in Samuelson’s (1938) revealed preference approach, Bayesian decision-making can thus be given an entirely behavioural foundation, freed from any process of introspection.

implies adhering to the tenets of Bayesianism.⁵ Following on Ramsey and de Finetti, Leonard Savage (1954) was responsible for providing the axiomatic foundations of this approach.

In recent years, an increasing number of scholars have questioned this approach. The attack has been multifaceted, but we shall concentrate here on the issue of consistency, following on the authoritative viewpoint expressed by Itzhak Gilboa, a main contributor to the revival of decision theory under uncertainty as different from risk. Elaborating on a previous analysis with his mentor and co-author David Schmeidler (Gilboa and Schmeidler 2001), Gilboa (2015) offers a methodological justification for the dismissal of the consistency view proposed by Savage and endorsed by the whole Bayesian approach. Gilboa's unusual readiness to discuss the methodological foundations of contemporary decision theory allows for a focus on a number of relevant normative issues usually taken for granted, but rarely properly examined among contemporary decision theorists.

Gilboa's (2015) analysis starts from a historical reconstruction that hinges on the turn towards an understanding of rationality as consistency experienced in economic theory. While "philosophers of the eighteenth and nineteenth centuries did not shy away from making statements about the substantive meaning of rationality" and from expressing "views about what 'Rational Man' should think and do, on issues that are often a matter of value judgment," the rise of neoclassical mathematical economics in the early twentieth century "could be viewed as taking a step back, and reducing the concept of rationality to consistency." He follows on the now widely agreed view (Giocoli 2003) that rationality "started to be defined as behaving in a way that is sufficiently coherent to allow certain formal representation," remarking that, as a result, "rationality ceased to be a matter of content, and became a matter of form" (Gilboa 2015, p. 315). The spreading of this viewpoint from decision theory to economics in general was pervasive and long-lasting, and is still textbook reference.

It is well-known, though, that the Bayesian paradigm came under attack since the early days of its systematization. The status of the crucial independence axiom was questioned by Maurice Allais (1953) with a series of choice examples while the new mainstream was still in the making. But Savage (1954, pp. 102-103) objected to the critique that the impossibility to represent the actual choices of

⁵ The tenets of Bayesianism are usually listed as follows: (1) a representation of the environment in terms of state space (all information is summarized in states of nature), (2) prior probabilities (degrees of belief are given by a single probability measure defined over the state space), (3) Bayesian updating (in view of new information posterior probabilities are derived from Bayes's law), and (4) maximization of EU (individuals act as if they are maximizing a utility function). See Gilboa *et al.* (2012, pp. 14-15).

individuals through a utility function—something that characterizes the modal choices presented by Allais—did not concern the deliberate choices of a rational individual, who would certainly change her own “irrational” choices after closer examination, that is, in view of the axioms of the theory and their normative value.⁶ It took more than 25 years for the issue to come back to the fore, starting from the work of Kahneman and Tversky (1979). But Kahneman and Tversky’s prospect theory—motivated by the Allais paradox and similar failures of EU theory—was intended as a descriptive theory of behaviour: the normative value of Savage’s theory was not questioned, and apparently it is not yet in a significant part of behavioural economics (Kahneman and Thaler 2006).

Gilboa concentrates instead on a second line of attack that expands the criticism from Allais’s risky situations to more general decision contexts and uncertainty proper, a line of research originated by Ellsberg (1961). Gilboa (2015, p. 328) reminds us that “in the second half of the twentieth century, the Bayesian paradigm ... was criticized for its descriptive validity, with experimental work that followed Ellsberg (1961),” and that relevant criticism was also raised on theoretical grounds by statisticians such as Shafer (1976) who objected to the way Bayesians describe information available to the decision-maker. However, Savage’s axiomatization of decision-making remained a “cogent argument that rationality implies Bayesianism” even after the counterexamples proposed by Ellsberg, since works aimed at offering generalizations of the Bayesian approach “were not axiomatically related to decision-making” (Gilboa 2015, p. 329). Indeed, the influence of the Ellsberg Paradox on theoretical research had to wait decades before becoming significant (Camerer and Weber 1992).

It was only toward the end of the twentieth century, Gilboa argues, that economists started questioning the axioms that define rationality, with the aim to propose alternative axiomatic justifications. And while most of theoretical, experimental, and empirical studies were still concerned with the descriptive validity of the Bayesian model, “normative issues inevitably surfaced ... [and] the essence of rationality was questioned.” Thus, by the end of the century, “it was no longer obvious which notion of consistency should be used to define rationality” (Gilboa 2015, p. 315).

Before moving on to his proposed way forward, it must be noted that Gilboa’s assessment is a conventional one among scholars interested in decision theory under ambiguity/uncertainty.⁷ Gilboa (2009, p. 132) makes a distinction between counter-examples addressing behavioural issues in view of the descriptive failures of the theory, and an analysis motivated by a sense of “cognitive unease”

⁶ On the history of the independence axiom see Fishburn and Wakker (1995).

⁷ See for instance Blume and Easley (2008), Binmore (2009), Eichberger and Kelsey (2009), and Machina and Siniscalchi (2014). Another prominent contributor to the literature, Peter Wakker, is less adamant about the actual normative significance of recent developments (Wakker 2008, p. 431).

with the manner that the Bayesian paradigm deals with information. It is this cognitive unease, pointing to the normative failure of the Bayesian approach, that characterizes David Schmeidler's (1989) point of attack and the recent advances in non-Bayesian theories of decision-making. Hence, for Gilboa, the idea that one should move over urn examples and concentrate on uncertainty proper is due to contemporary decision theorists such as Schmeidler and himself, while Ellsberg only addressed descriptive failures of Savage's axioms.⁸

This is not intended to suggest that Gilboa does not give prominent role to Ellsberg's urn examples, as is common in the literature of decision-making under uncertainty (Machina *et al.* 2011). For instance, Gilboa is keen to distinguish between framing effects à la Kahneman and Tversky and the ambiguity issues highlighted by the Ellsberg Paradox, and clarifies that while both were intended as descriptive analyses, Ellsberg's has stronger normative implications. People who initially do not recognize the equivalence of two representations of the same problem, as in Kahneman and Tversky's experiments, are usually embarrassed to discover that, while people refusing to bet on Ellsberg's ambiguous urns do not change their minds even after Savage's sure-thing principle is explained to them. However, while taking the Ellsberg Paradox as a crucial starting point, Gilboa refers to its crucial descriptive relevance, but never examines Ellsberg's normative aims and seems to downplay Ellsberg's role in fostering an alternative, normatively motivated, approach to decision-making. This may be correct in terms of the actual influence of Ellsberg on the following developments but, as we shall see, it is inaccurate to suggest that Ellsberg was silent on normative issues.

3. Gilboa's alternative to the consistency view of rationality. Or, does rationality imply Bayesianism?

Gilboa notes that, since the Bayesian approach to rationality can be reduced to an issue of consistency to be satisfied by individuals' choices, the standard definition of rationality used in decision theory and economic applications is purely behavioural. Rationality is made equivalent to "behaving in compliance with Savage's axioms or some other set of axioms" (Gilboa 2009, p. 136). In order to put forward his criticism of the Bayesian approach, then, for Gilboa it is mandatory to identify a different notion of rationality.

⁸ Gilboa (2009, p. 133) claims: "Schmeidler's intuition ... has a behavioural manifestation in Ellsberg's paradox. But Schmeidler did not start out by attempting to explain the experimental evidence." See also Gilboa *et al.* (2012, pp. 18-19), and Gilboa and Marinacci (2016, p. 398).

Gilboa's proposal for an alternative notion of rationality starts from observed behaviours, but introduces the idea of "reasoning about behaviour" as a necessary component for the valuation of the axioms' relevance to decision-makers. Following on his elaboration with Schmeidler (Gilboa and Schmeidler 2001, pp. 17-18), Gilboa suggests to use a definition of rationality that has to do not only with behaviour, but also with "cognition." He observes that a mode of behaviour must be considered "irrational" for a decision-maker if, when "exposed to the analysis of her choice, she would have liked to change her decision." The cognitive aspect is relevant here insofar as this definition "is based on a sense of embarrassment about one's decisions, and not only on observed behaviour." Whatever the reasonableness of the axioms suggesting rational actions, a mode of behaviour is "subjectively rational for a decision-maker if she cannot be convinced that this decision is wrong" Gilboa (2015, p. 316).⁹

For instance, as we shall see in section 6, in the case of Ellsberg's urns it is typical to observe choices that do not imply degrees of belief compatible with a single probability function, contrary to what is requested by Savage's representation theorem for the maximization of SEU. But this cannot be considered a descriptive failure of the theory since these modal choices are usually upheld even after consideration of the axioms that are violated. In fact, "if people [exposed to the Ellsberg Paradox] shrug their shoulders when we explain the logic of our axioms, or admit that these are nice axioms but argue that they are impractical, we should better refine the theories" (Gilboa 2009, p. 139).

Gilboa's suggestion is then is to concentrate on the normative component of a theory—which in principle is not subject to falsification through counterexamples—but also to allow for a text connected to empirical relevance. He reminds us that, usually, normative science aims to be prescriptive in view of certain general principles regarded as evident, at least by the theoretician. Yet, Gilboa (2015, p. 313) concludes, this does not seem to be enough when adverse empirical evidence is overwhelming: in such a case it must be admitted that "being different from reality is hardly a sufficient criterion for a good normative theory." Hence, logical or probabilistic reasoning can be viewed as normative theories only insofar as this is "how people would like to reason." To be useful, normative theories "cannot derive their authority from the academic fame of the professors who preach them." Rather, Gilboa (2015, p.313) concludes, "logic and probability should be put to a test, in which people express their preference to adopt these tools in order to reason better (in their own eyes)" (see also Gilboa *et al.* 2009, p. 292).

⁹ In a recent interview Gilboa claims: "My definition of rationality started with this: asking what do people have in mind when they refer to something as 'rational.' But the best definition I came up with is in terms of what most people would be willing to accept as their decision making models, as opposed to what they would like and could change" (Herfeld 2015, p.128).

Coming back to the issue of whether rationality implies Bayesianism, then, the question addressed by Gilboa is that when tested along the lines just sketched the Bayesian approach does not pass the test in relevant decision contexts, specifically when people act in an “uncertain world” (Gilboa 2015, p. 330). Gilboa admits that, even among economists who are willing to consider alternative descriptive models, the majority still contends that “people may indeed be irrational in a variety of ways, including violating Savage’s axioms, but that rational people should not behave so wildly, and that rationality implies Bayesianism.” The mainstream argument still is that the Bayesian approach offers a normatively convincing answer to decision-making under uncertainty when it suggests that the way to obtain subjective rationality is to select a prior probability, use it for belief formation and choice, and then update it in the face of new information.

However, this is but one approach, which may not be the best choice for all people in all problems.¹⁰ It is evident in any problem that has to be faced before any information has been gathered—typically when decisions are not routine, but of crucial relevance—that the selection of a sharp prior is “highly arbitrary.” Therefore, Gilboa (2015, p. 330) concludes, “the Bayesian approach may not be the most subjectively rational for many people”: while decisions will be coherent given the selection of the prior, the arbitrariness of the latter or the unwillingness of the decision-maker to commit to a single one, exposes the entire decision process to criticism.

Gilboa’s point is that decision theory has made progress exactly on this front. That is, it has suggested alternatives to the Bayesian approach, and that these alternatives constitute axiomatically based, general purpose models of decision-making. Stimulated by what he perceived as a normative failure of Bayesianism, Schmeidler (1989) has initiated research on models in which individuals’ priors are represented through a non-additive probability measure, and choice alternatives are ordered by means of a particular integral operator for non-additive weights called Choquet integral, giving birth to so-called Choquet EU. Under uncertainty, Schmeidler showed that it is rational for decision-makers not to commit on a sharp additive prior, but on a fuzzy one, so to act as if they maximize Choquet EU instead of SEU. In a different, but related, investigation Gilboa and Schmeidler (1989) concentrated on a representation in which the decision-maker entertains a set of priors rather than a single one, and axiomatized the possibility to use a Maxmin of EUs with respect to different priors as an alternative decision rule: in this second instance, every act is evaluated by the worst possible

¹⁰ Gilboa’s claim is enlightening when he makes reference to the other sciences using the Bayesian approach: only economics – and neither statistics nor computer science nor philosophy – accepts Bayesianism as “the sole claimant to the throne of rationality” for the representation of uncertainty (Gilboa *et al.* 2009, p. 287).

EU it could obtain, ranging over the various priors in the set, and this minimal EU is maximized by the decision-maker.

Both the non-additive Choquet EU model and the Maxmin EU model can be used to explain experimental evidence such as the one related to the Ellsberg Paradox. Their main relevance has been to provide an a-priori argument for the contention that, when uncertainty is pervasive, the Bayesian approach is too restrictive to represent information satisfactorily. Furthermore, both models can be thought of as having a behavioural justification since they rely on observable preferences and derive a representation thereof, which need not describe any actual mental process: in the main, “both representations allow a naïve cognitive interpretation, but they are also compatible with other cognitive processes that may result in behavior that satisfies the relevant axioms” (Gilboa *et al.* 2012, p. 11).¹¹ And since the aim to set out normative alternatives to the Bayesian approach has been reached, Gilboa concludes, it is notable that we have positive formalized contributions to substantiate the view that consistency of the Bayesian kind no longer is a necessary condition for rationality.¹²

4. The turn to normativism in decision theory

It is a main aim of this paper to trace the origins of Gilboa’s argument that economic theory should advance beyond the conventional assumption that rationality implies Bayesianism, and that this step entails a reorientation of normative analysis. So before concentrating on why the Ellsberg Paradox is crucial for the demise of a strict Bayesian viewpoint—and on significant aspects of its assessment Gilboa seems to have overlooked—this and the following section re-examine the rationale underlying

¹¹ The influence of these axiomatically based alternatives to strict Bayesianism has been huge, both in applications (e.g., Mukerji and Tallon 2004) and in refinements concerning operational decision rules (e.g., Klibanoff *et al.* 2005). However, it is important to note that discordant views about the actual significance of the approach exist as well, even among main contributors to the topic. For instance, criticism of the normative achievements of the approach—with specific regard to updating and the failure of ambiguity aversion model to satisfy consequentialism, a theme not addressed in this paper—is presented in Al-Najjar and Weistein (2009).

¹² Similarly, in their survey of rationality issues, Blume and Easley (2008, p. 891) conclude: “Rationality does not mean expected utility. Expected utility is one small class of decision models for choice under uncertainty. Its dominance in application was understandable 30 years ago when few alternatives were on the table. Since then decision theorists have been creative in developing better-behaved alternatives, and equilibrium and game theorists have been clever in applying them.”

normative arguments in view of the historical episode that definitely turned decision theory into a normative science.

As already noted, the mid-twentieth century developments established a consensus among economists on the basis of which rational choice theory has been increasingly viewed as a normative rather than positive theory of the behavior of individual agents, a “theory that describes what one ought to do in order to be rational” (Hands 2015, p. 6). This consensus was mainly related to the reinterpretation of von Neumann and Morgenstern’s EU theory proposed by Marschak (1950), who made explicit the axiomatic bases justifying the maximization of EU. On this ground, the usual axioms for choice under uncertainty—mainly transitivity and completeness, plus technical conditions such as continuity—were supplemented by the so-called independence axiom, thus accounting for risky prospects under the same theoretical construct.¹³

Following on the probabilistic approach of Bruno de Finetti (1964 [1937]), this viewpoint consolidated through the work of Savage (1954), who extended it to choice under uncertainty. Savage argued that even when objective probabilities of events conditioning the payoffs of individual agents are not available—typically for singular events such as “the outcome of the next presidential election” (Savage 1954, p. 21)—the subjective probabilities that can be elicited from choices can be used as the probabilities for the maximization of (subjective) EU by an ideal, rational agent. Due to the endorsement of Savage’s subjective perspective by decision theorists, the traditional Knightian distinction between risk and uncertainty lost relevance: probabilistic reasoning, typically admitted to deal only with risk, could be used also to address uncertainty. The new mainstream in decision theory then came to be characterized as Subjective EU (Luce and Raiffa 1957, Raiffa and Schlaifer 1961).¹⁴

The conventional assessment is that Savage’s *Foundations* signal a substantial retreat of decision theory from descriptive ambitions. Savage’s reaction to Allais’s (1953) example, as documented in his 1954 volume, is seen as a decisive step in the shaping of rational choice theory. Savage used a

¹³ Herfeld (2017) provides evidence from archival material about the Cowles Commission in the late 1940s that Marschak and Tjalling Koopmans were responsible, as directors of the Commission, of a commitment at Cowles to the axiomatization of economics. Their aim was to spread this viewpoint even in branches, such as decision-making, which were traditionally seen as not easily separable from other behavioural sciences. Marschak (1946, p. 114) evidenced as a major methodological step taken by von Neumann and Morgenstern the separation of formal axiomatic structure and empirical interpretation, and promoted the development of economics by means of the “tools of modern logic” in a series of Cowles Discussion Papers culminating in Marschak (1950) and (1951).

¹⁴ The term “Bayesian” started being used first among statisticians interested in applications of Savage’s approach to statistical practice and then became usual even among decision theorists (Fienberg 2006).

chapter devoted to “historical and critical comments on utility” to contend, in reply to Allais, that decision theory has a normative status that cannot be criticized on the grounds of descriptive failures. Savage’s reaction became—and still mostly is—the conventional one of what Allais called the “American School” when challenged by counterexamples (Hirshleifer and Riley 1992). As a matter of fact, at least until the consolidation of experimental economics, the status attributed to Allais’s counterexample was that of an empirical violation, merely consisting of a logical mistake.¹⁵

It is worth noting though that Savage’s determinacy to put forward a normative interpretation of his theory, and to downplay the significance of descriptive violations, had already emerged as the outcome of the debate on empirical evidence that was of crucial relevance soon after von Neumann and Morgenstern (1947). In particular, Savage’s collaboration with Milton Friedman shows that the turn away from the search for empirical evidence validating EU theory—something Friedman and Savage placed at the core of their classic 1948 paper on the theory of decision under risk—was already evident before the meeting with Allais. While in their 1948 paper Friedman and Savage (1948, p. 282) had aimed to provide a “crude empirical test” of the hypothesis of maximization of EU, and found that “empirical observations are entirely consistent with the hypothesis,” a follow-up paper of 1952 concentrates instead on the significance of the axioms being proposed, and investigates their potential for normative economics, endorsing Marschak’s (1951) position that the axioms justifying the maximization of EU under risk should be considered like elements of logic and mathematics.¹⁶

Friedman and Savage present their 1952 argument as a reply to William Baumol’s (1951) critique of their previous paper about the status of the EU theory. Apart from the issue raised by Baumol about the cardinal measurability of the EU function—which they acknowledge to be a semantic question with “unfortunate connotations” they inadequately addressed in 1948—the 1952 paper concentrates on the “grounds for accepting or rejecting the [EU] hypothesis.”¹⁷ Friedman and Savage (1952, p.

¹⁵ It took decades before Savage’s position was questioned extensively. Before Allais and Hagen (1979), only Slovic and Tversky (1974) cogently made the point that the intuitive rationale underlying the violation was not abandoned by individuals, even after careful consideration of EU theory. Machina (1987) reviews the wave on non-EU theories originated by the Allais Paradox in the 1980s.

¹⁶ Friedman and Savage’s (1948) early work assessed the historical developments of the economic theory of risk in order to make apparent the ability of von Neumann and Morgenstern’s (1947) approach to provide a long awaited axiomatic bases for a utility function justifying both buying of insurance and gambling. Together with the 1952 follo-up, the two papers are usually considered as part of a coherent effort at systematization, though a fundamental step towards normativism that is put forward in the second did not feature the previous one. Starmer (2005) is a notable exception.

¹⁷ Baumol (1951, p. 65) presented an hypothetical example suggesting behavior that would contradict the EU hypothesis and yet would not be clearly “pathological,” arguing that he would not consider

463) consider the possibility to regard the maximization of EU as a “maxim of behavior,” whose significance depends “not on its empirical verification for the economic behavior of men at large, but on its acceptability, to individuals who are particularly concerned with such decisions, as a rule guiding ‘wise’ behavior in the face of uncertainty”—a viewpoint already hinted at by Savage (1951) when presenting Wald’s (1945) approach to statistics¹⁸. On these bases, they present a set of postulates “more plausible than the [EU] hypothesis to which they are logical equivalent.” The set now includes a postulate termed P3, that is, Savage’s version of the independence axiom for uncertain events, that in the *Foundations of Statistics* Savage calls P2, or the “sure-thing” principle.¹⁹ The 1952 paper, then, contains a preliminary presentation of the step forward Savage have been working on in those years, deemed to be in agreement with Marschak’s (1951) presentation of the axiomatic system for rational choice under risk (Friedman and Savage 1952, p. 467).²⁰

The emphasis on normative justification is apparent in the defense against Baumol’s allegation that the EU hypothesis is descriptively false. Friedman and Savage’s (1952, p. 466) argument consists of stating what are the features that makes EU not only a promising conjecture about behavior under uncertainty, but a truly “scientific hypothesis”. By “scientific” they mean, in Popperian fashion, a proposition which lends itself to be contradicted, and this entails making it possible to specify behavior that contradicts it, a crucial issue that will re-emerge in the criticism put forward by Ellsberg

some of the results of the application of EU “introspectively obvious.” Friedman and Savage’s (1952, p. 465) reply intended “to repair in some measure our earlier failure” by making clear which are the grounds for acceptance or rejection of the hypothesis, and to contrast Baumol’s claim that “casual observation and introspection suggest that the hypothesis is false.” Although not explicitly referred to, Samuelson’s (1950) criticism was also in the background.

¹⁸ Giocoli (2013) analyses how their reading of Wald, who proposed to see statistical practice as a decision problem under uncertainty from the statistician’s viewpoint, influenced Marschak and Savage in designing Bayesian rationality for economic agents.

¹⁹ The 1952 P3 postulate states that for subjective probability distribution f , g , and h defined over a finite set of outcomes X , and a parameter a such that $0 < a < 1$, $af + (1-a)h \leq ag + (1-a)h$ if and only if $f \leq g$, where \leq means f is not preferred to g . On the introduction of the independence principle and Savage’s sure-thing version see Fishburn and Wakker (1995).

²⁰ In the paper, Savage is explicitly reported to be doing further work on the postulate system. Apart from the forthcoming book on the foundations of statistics, reference is probably made to Savage’s 1952 notes for the Paris Conference, clearly endorsing the normative viewpoint (Savage 1953). That Savage had been working on this issue since 1949, when a preliminary report was presented at a meeting of the Econometric Society in Boulder (Savage 1950) is seldom noticed, despite Arrow’s (1951) remark (but see Giocoli 2013). On the evolution of Savage’s thought towards the *Foundations*, see also Moscati (2016), who clarifies that Savage mentioned the sure-thing principle, and defended its normative appeal, already in his correspondence with Samuelson, in mid-1950s.

later. Given that “direct” evidence to contrast Baumol’s claim that behavior is conceivable that would contradict it is not yet available—although “some recent experiments by Mosteller and Nogee add to the direct observational evidence and fail to contradict the hypothesis” (Friedman and Savage 1952, p. 466)—the “real appeal” of the hypothesis consists of “indirect” evidence, namely, the “plausibility” of the set of postulates which makes the EU hypothesis implied as a (representation) theorem.

The historical interest of Friedman and Savage’s 1952 paper is that this is arguably the first addressing the issue of how it can be claimed that a proposition that cannot be empirically proved may still hold true, the question we saw Gilboa dealing with in the previous section. Although the paper also refers to the methodological position that an “as if” interpretation is in order—excluding that people actually compute EU values, and arguing that the value of the EU hypothesis is to be assessed in its ability to enable correct predictions (Friedman 1953)—its main proposal is to focus on the “intuitive appeal” of the postulates, particularly the independence axiom P3—soon to be renamed as the sure-thing principle P2. Friedman and Savage (1952, p. 469) argue that, as regards possible objections to the significance of this crucial postulate, the fundamental point is that if a “reader” considers the postulate in the light of the illustration they provide in the paper “he will concede that the principle is not one he would deliberately violate.”²¹ And this is why there is “reason for supposing that people do actually tend to avoid flagrant violations of the principle.” The emphasis then is already on deliberate choices: deliberate violations by a rational decision-maker are excluded due to the plausibility of axioms.

To sum up, Friedman and Savage (1952, p. 474) admit as “unexceptionable” Baumol’s objection that the EU hypothesis is not “a useful or valid interpretation of actual behavior.” But “the hypothesis is fruitful and justifies acceptance,” they conclude, considering “its coherence with the body of economic theory and, more important, ... the plausibility of postulates with which it can be shown to be equivalent rather than from repeated success in prediction.” A violator of the fundamental axioms of the theory can be considered as such not after “casual observation” of her choices but only after she has been informed of the highly logical appeal of axioms, and only if she—now deliberately—insists on her choices.²²

²¹ After Moscati’s (2016) assessment of the private exchange between Savage and Samuelson on the issue of the independence principle, it is conceivable that the reader Friedman and Savage refer to may have been Samuelson himself.

²² Even after inspection of the material available in Leonard Jimmie Savage Papers at Yale University Library, it is unclear whether Savage was already aware of Allais’s argument when the 1952 paper was submitted for publication. Moreover, the English version of Savage’s notes presented at the Paris International Colloquium on Risk in which Savage was informally tested by

5. The consistency view of rationality in Leonard Jimmie Savage

As seen in the previous section, all the elements of the methodology put forward by Savage in his *Foundations* were already surfacing in its preparatory works, including the fundamental point that, even though counterexamples may emerge—as claimed by Baumol even before Allais—this would not imply that the EU hypothesis is wrong. The *Foundations* presented the derivation of EU theory under uncertainty in an all-comprehensive axiomatic set-up, clarifying to which extent the use of subjective probabilities could be instrumental to decision-making, and using them for a substantial reorientation of statistical practice. As is well-known, the main formal result provided by Savage was the identification of a set of axioms on preferences over “acts”—choices whose consequences depend on the realization of external events—making possible to prove as a theorem that obeying to them is equivalent to the maximization of Subjective EU.²³

Establishing what became the Bayesian standard, individuals’ degrees of belief are assumed to be the “personal probabilities” attributed by the individual to payoff relevant events, in principle unrelated to anything objective.²⁴ The numerical values of these personal probabilities can be determined observing actual behavior in any instance, while introspection, or intuitive derivation of probabilities, is denied relevance. Through this operational perspective in terms of economic behavior every individual can then be endowed with whatever degree of belief, provided her behavior is consistent—as in Samuelson’s (1938) derivation of utility from preferences revealed by choices. Consistency is the only admitted constraint to an individual’s degrees of belief, and this holds true

Allais—published in French as Savage (1953)—is undated. Both Friedman and Savage (1952) and Savage (1953) do not refer to Allais.

²³ Savage’s “representation theorem” states that if an individual’s preferences satisfy his axioms, then these preferences can be represented by a utility function assigning a utility index to every consequence, and a probability function assigning a probability to every event. The crucial axiom to prove this result turns out to be the “sure-thing” principle: a preference between two acts is independent of the states in which the two yield identical consequences. For acts f and g , this requires $f \succeq g \leftrightarrow f' \succeq g'$ whenever the state space can be partitioned into two parts, the irrelevant event and the relevant event, such that, on the irrelevant event, $f \sim g$ and $f' \sim g'$, and on relevant event, $f \sim f'$ and $g \sim g'$. When probabilities are assigned to states, expected-utility maximization implies the sure-thing principle because the parts of the expected utilities of f and g or of f' and g' related to the irrelevant event cancel (Fishburn and Wakker 1995, p. 1136).

²⁴ Savage was keen to make a distinction which is no longer usual in current decision theory. He placed emphasis on the fact that the subjective probabilities of an individual are her own “personal” probabilities. Since his approach was intended to replace the objective probabilities of frequentism, he feared that the adjective subjective could be used to downplay the scientific value of his contribution. In Savage’s view, nothing is subjective, but the personal probabilities of the decision-maker.

because otherwise behavior would be such that a dominated choice would be chosen.²⁵ The axioms of subjective probabilities are then those which, when applied to consistent degrees of belief, guarantee that the decision-makers act as if they were EU maximizers, even under uncertainty.

But Savage's volume aimed also to set out a clarification of a number of methodological issues tangled with its axiomatic orientation, in the direction anticipated in the 1952 paper. In an attempt to define a coherent framework for decision-making, Savage presents his derivation of Subjective EU theory explicitly as an attempt to extend logical reasoning to situations of uncertainty. His declared aim is not to provide an empirical theory for predicting human behavior, but to develop logical tools for deciding between alternative courses of action. Decisions made in the face of uncertainty, Savage claims, require "formal reasoning": his purpose, then, is "to discuss the implications of reasoning for the making of decisions." Savage uses the analogy with logic already referred to by Marschak (1950): since reasoning is usually associated with logic, but the implication of "ordinary logic" are limited in the face of uncertainty, the subject matter of his investigation is to ask "whether logic cannot be extended, by principle as acceptable as logic itself, to bear more fully on uncertainty" (Savage 1954, p. 6). As a result, Savage concentrates on building up "a highly idealized theory of the behavior of a 'rational' person with respect to decisions," and although rational means "logical," decision-makers "have to make up his mind in situations in which criteria beyond the ordinary ones of logic will be necessary."

Savage notes that logic admits both an empirical and a normative interpretation. But since he regards the axioms of the theory of decision as "logic-like criteri[a] of consistency in decision situations," he concentrates on the latter one: "The principal value of logic is in connection with its normative interpretation, that is, as a set of criteria by which to detect ... any inconsistencies there may be about our beliefs, and to derive from the beliefs we already hold such new ones as consistency demands" (Savage 1954, p. 20). Pursuing the analogy with logic, he concludes, "the main use I would

²⁵ Starting from de Finetti's (1931) idea that an individual's degree of belief in a proposition (or event) corresponds with the lowest odds at which she would bet on the truth of that proposition, the condition of coherence can be summarized as follows: for the measurement of the degree of belief as a probability, "the subjective theory considers the conditions of a bet which this individual would accept" under circumstances such that "the restrictions necessary and sufficient for avoiding inconsistencies among conditions of betting" be satisfied (De Finetti 1951, p. 223). Ramsey (1964 [1931], p. 83) made the same point arguing: "if anyone's mental condition violated these laws [the laws of probability], his choice would depend on the precise form in which the options were offered him, which would be absurd. He could have a book made against him by a cunning better and would then stand to lose in any event." In the literature this is usually referred as the Dutch Book principle.

make of P1 and its successors [the axioms of subjective probability] is normative, to police my own decision for consistency and, where possible, to make complicated decisions to depend on simpler ones” (p. 20).²⁶

Crucially, being already involved in the defence of his approach even before the actual publication of his volume—and having excluded the descriptive path to a rejection of “scientific” hypotheses—Savage could not avoid facing the cumbersome issue of a critical method for the testing of a normative proposition. The idealized theory of the behaviour of a rational individual when making decision under uncertainty cannot be deemed to be scientific simply by equating rational with logical, as suggested by Marschak. As just seen, Savage admits that under uncertainty “criteria beyond the ordinary ones of logic will be necessary,” and indeed his “sure-thing” axiom is introduced exactly to do so. As a result, Savage (1954, p. 7) suggests a step further: “... when certain maxims are presented for your consideration, you must ask yourself whether you try to behave in accordance with them, or, to put it differently, how would you react if you noticed yourself violating them.”

In Savage’s view, as we already know from the previous section, this would mean analyzing “deliberate” behavior, something Marschak did not contemplate. Now in response to Allais’s (1953) critique of Subjective EU and his own failure to adhere to it, Savage (1954, pp. 102-104) argues that theories of rational behaviour have a normative status that is unquestionable even in the light of adverse evidence, and that contrary evidence can only indicate “irrational” choices.²⁷ He admits that when confronting counter-examples, it may be difficult to distinguish between making a mistake and not obeying to Subjective EU axioms, and offers his own methodological way out: “In general, a person who has tentatively accepted a normative theory must conscientiously study situations in which the theory seems to lead him astray; he must decide for each by reflection ... whether to retain

²⁶ That this justification for consistency has an ex-cathedra flavour is plainly admitted by Savage (1954, pp. 20-21) when providing an intuitive rationale for “why and when consistency is a desideratum.” He states: “Suppose someone says to me, ‘I am a rational person, that is to say, I seldom, if ever, make mistakes in logic. But I behave in flagrant disagreement with your postulates, because they violate my personal taste, and it seems to me more sensible to cater to my taste than to a theory arbitrarily concocted by you’. I do not see how I could really convert him, but I would be inclined to match his introspection with some of my own.”

²⁷ As is well-known, when he was presented Allais’s example at an informal meeting during the 1952 conference in Paris, Savage expressed preferences contradicting his axioms. Upon reflection, though, he claimed: “it seems to me that in reversing my preference ... I have corrected an error” (Savage 1954, p. 103). For a historical reconstruction of the Paris conference, mostly from Allais’s viewpoint, see Jallais and Pradier (2005). Heukelom (2014) illustrates how Savage answered to Allais in the private exchange they had after the conference. Mongin (2014) explores the methodological features of the debate.

his initial impression of the situation or to accept the implications of the theory for it” (Savage 1954, p. 102). As a matter of fact, his own experience as a violator is that a “reasonable” decision-maker who instinctively violates the theory when confronting Allais’s example will reverse her choice after “thorough deliberation.” But he is also let to admit that if the individual maintains preferences that are in conflict with an axiom even after revision, “he must abandon or modify the principle: for that kind of discrepancy seems intolerable in a normative theory.”²⁸

It can then be concluded that an inspection of the methodological bases of Savage’s *Foundations of Statistics* shows that the original place where Gilboa’s suggested definition of rationality was first articulated is precisely the foundational volume of Bayesianism. While he was putting forward what is today termed a Bayesian view of rationality, Savage also suggested that one could consider irrational only decision-makers whose behaviour is inconsistent with their own theory, and that empirical evidence should be scrutinized with respect to decision-makers readiness to revise their choices in the light of their theories of rationality. And indeed the way to deal with empirical violations of a normative theory—termed a “quasi-empirical” test of rationality by Guala (2000)—was debated since the early days of decision theory. But while Allais’s argument was deemed to be inconclusive for long, a further round of discussion about similar methodological issues came to the fore early, as a new paradox was put forward by Daniel Ellsberg.

6. Classic criticism of the Bayesian viewpoint: the Ellsberg Paradox

Despite its widespread acceptance among decision theorists, further fundamental criticism of Savage’s axiomatization of decision-making appeared early in the literature. In a similar, but distinct fashion, Daniel Ellsberg (1961) and William Fellner (1961) argued that the neglect of the risk/uncertainty distinction suggested by Savage’s axiomatization was inappropriate for what they viewed as a much needed development of criteria for basing action upon vague knowledge. In order to emphasize the importance of Knightian uncertainty for decision-making, Ellsberg proposed a counterexample to Savage, while Fellner suggested an alternative representation of individual degrees

²⁸ Allais’s (1979 [1953], pp. 79-80) first detailed report of his objections to Savage at the 1952 Paris Conference hints at a similar methodological viewpoint when he discusses an “experimental definition of rationality” as implying that “rationality can be defined by having regard to the behaviour of persons who are commonly consider as rational”: when one does not find “desirable to use an abstract definition of rationality, the only option is to rely on experience, *and to observe the behaviour of men one has reason in other respects to believe act rationally.*” See also Allais 1953, p. 504.

of beliefs. Howard Raiffa (1961) contributed an immediate reply, arguing in favor of Savage and his subjective approach. Notwithstanding a brief round of debate and a first experimental paper confirming the significance of Ellsberg's counterexample (Becker and Brownson 1964), the issue was not addressed for many years.²⁹ Raiffa's position was almost unanimously endorsed at the time: Raiffa (1961, p. 693) argued that he ran experiments in class replicating Ellsberg's examples, but these simply showed that "there is a need to teach people how to cope with uncertainty in a purposive and reflective manner."³⁰

It should be noted that the Allais-Savage debate was not referred to by any of those participating to this new round of debate: Allais's criticism focused on utility axioms in an unambiguous context with objective probabilities, while the concern now was the limitations of probability axioms in contexts in which only subjective probabilities were available.³¹ But Raiffa implicitly assumed Savage's solution to Allais as the conventional way to deal with hypothetical examples: referring to experimental sessions he conducted using adaptations of Ellsberg's example, Raiffa (1961, pp. 692-693) reported that experimental subjects who were initially inclined to violate the axioms behaved consistently with the axioms when problems were presented to them in a different form. As for Savage, he never examined Ellsberg's problems in his writings, probably because he was consumed by the debate in what he perceived as the much more relevant, and hostile, statistical arena (Wallis 1981).³²

Of the two critical papers written in 1961, Ellsberg's is the more celebrated one since it introduces new counterexamples to the mainstream view. Although presented as thought-experiments, Ellsberg's counterexamples received overwhelming support in the laboratory, as testified in a first survey by Camerer and Weber (1992), and have remained central for contemporary experimental studies (Halevy 2007). Moreover, Ellsberg introduced the notion of "ambiguity," to be interpreted as an additional dimension of the problem of choice beyond utility and probability. He argued that the

²⁹ Papers by Brewer (1963) commenting on Fellner and Roberts (1963) on Ellsberg, with short replies (Fellner 1963, Ellsberg 1963), did not ignite a substantive theoretical debate. Apart from some notable exceptions, which we shall see later, even experimental analysis languished until the 1980s.

³⁰ Ellsberg's reply to Raiffa constituted part of his doctoral thesis, but, as we shall see in section 7, the thesis was not published at the time, and did not have an impact on following developments.

³¹ Traditionally, it has been assumed that Allais discussed risk, while Ellsberg uncertainty (Wakker 2008).

³² The only words dedicated to Ellsberg's and Fellner's criticism in Savage's printed works appear in Savage's (1970, p. 25) annotated bibliography prepared for the 2nd edition of the *Foundations* in which the *QJE* Symposium is presented as "an account of an important line of dissent from the theory of personal probability and utility."

violations of rationality he observed in casual testing suggested aversion to ambiguity, since they were related to “situations where available information is scanty or obviously unreliable or highly conflicting; or where expressed expectations of different individuals differ widely; or where expressed confidence in estimates tends to be low” (Ellsberg 1961, p. 660-661). The idea of ambiguity caught such a number of relevant phenomena that, eventually, it became as relevant as uncertainty in recent advances in decision theory (Machina and Siniscalchi 2014).

It is apparent from the beginning of his paper that Ellsberg’s main aim was to contrast the normative value of Savage’s theory: he introduced his analysis by claiming that he did not consider Savage’s approach to be a normative guide in decision contexts perceived as “ambiguous” by decision-makers. While reviewing the mainstream view he remarked that “a number of sets of constraints on choice-behavior under uncertainty have now been proposed ... having the implication that—for a ‘rational’ man—all uncertainties can be reduced to risks,” and that “the propounders of these axioms tend to be hopeful that the rules will be commonly satisfied, at least roughly and most of the time, because they regard these postulates as normative maxims.” Crucially, while referring to Savage’s attitude towards possible violations of his axioms, Ellsberg (1961, p. 646) also noted: “these axioms are believed to predict certain choices that people will make when they take plenty of time to reflect over their decision, in the light of the postulates.”

Ellsberg remarked that a novel characteristic of Savage’s approach was that, unlike the traditional Knightian approach, it offers an operation meaning to the proposition that decision-makers act as though they assign probabilities to uncertain events. But this makes it also possible to identify instances in which individuals do not intend to act as the theory prescribes. Since Savage’s methodology requires the elicitation of subjective probabilities from choices concerning events, an example showing that there is no way to infer definite probabilities would make theories representing uncertainty in terms of probabilities “quite inapplicable” (Ellsberg 1961, p. 646). In order to prove his argument, then, Ellsberg suggested examining a very simple set of choices concerning balls to be drawn from urns, and eliciting probabilities from bets on these urns.

A main example concerned two urns, and examined decision-makers who faced the choice of betting on the drawing of a red or black ball from either a first urn containing 100 red and black balls in an unknown proportion, or a second urn containing 50 red and 50 black balls. Ellsberg defined the first urn to be ambiguous with respect to the second, risky one. When confronting each urn separately, individuals tended to be indifferent between betting on red or black.³³ But when asked to choose the

³³ This is to be expected for the second urn, because their probabilities are objective in a frequentist sense. But is supposed to be true as well for the first one considering the conventional epistemic

urn from which they would prefer to bet that a red (or a black) ball would be drawn, most respondents revealed a preference for the second one instead of the first. Apparently, to confront a risky, but unambiguous urn is preferred than to confront an ambiguous one. But, Ellsberg observes, this is contradictory.³⁴ Ellsberg's intuition was that individuals pay serious attention to considerations such as vagueness and scantiness of information and that, if not uncertainty proper, that would imply considering ambiguity.³⁵

Remarkably, while presenting his examples, Ellsberg followed Savage's suggested methodology. The decision-maker is informed that her degrees of belief do not correspond to a usual probability prior, and then is asked to pause to reconsider her replies. If she repents her violations, she should decide that her choices "implying conflicts with the axioms were 'mistakes'", and that her "'real' preferences, upon reflection, involve no such inconsistencies:" the decision-maker would confirm that the Savage postulates are, if not descriptive rules, her "normative criteria in these situations" (Ellsberg 1961, p. 655). But, Ellsberg argues, "this is by no means a universal reaction; on the contrary, it would be exceptional." In a suggestive passage, Ellsberg's (1961, p. 656) reports the "important finding" that "after rethinking all their 'offending' decisions in the light of the axioms, a number of people who are not only sophisticated but reasonable decide that they wish to persist in their choices." His experience, after testing people at departmental meetings at Harvard and at the RAND Corporation, was that this group of "deliberate violators" included "people who previously felt a 'first-order commitment' to the axioms, many of them surprised and some dismayed to find that

interpretation of probabilities, that is, for reasons either of symmetry or insufficient information (Hacking 1975).

³⁴ This kind of choice is incompatible with the assignment of a single additive probability distribution: while the indifference of decision-makers between betting on the red (r) or blue (b) drawn from urn I or II examined separately means that their subjective probabilities p are such that $p(r_I)=p(b_I)=1/2$ and $p(r_{II})=p(b_{II})=1/2$, when they choose to bet on red (or black) from the second urn they reveal $p(r_{II})>p(r_I)$ (or $p(b_{II})>p(b_I)$). It is following on this remark, that the non-additive probability approach of Schmeidler (1989) proposes an axiomatic system for subjective probabilities such that $p(r_I) + p(b_I) < p(r_I \cup b_I)$. In contemporary decision theory the individual who prefers to bet on urn II is said to show ambiguity aversion (Eichberger and Kelsey, 2009; Machina and Siniscalchi 2014).

³⁵ Ellsberg proposed also a second example, with a single urn whose composition is partly known and partly unknown, and found a similar attitude among people tested, with many decision-makers who refrained from betting on the drawing of balls belonging to the unknown part of the urn. Technically this second example is even more significant, as it offers a direct test of Savage's sure-thing principle.

they wished, in these situations, to violate the Sure-thing Principle.” Ellsberg’s report is that this group included Leonard Savage himself.³⁶

It is worth noting that Ellsberg did not intend to leave his counterexamples in a vacuum. For individuals who act in conflict with Savage’s axioms “deliberately, without apology,” he proposes to apply a decision rule alternative to the maximization of EU. The quest for a new normative rule is apparent in his suggestion to incorporate a parameter for the confidence in the subjective probability assessment in a new decision rule. And despite being unable to provide an axiomatic foundation for his suggestion, his methodological point was clearly stated when arguing that “once certain patterns of ‘violating’ behaviour [are] distinguished and described in terms of a specified decision rule” (Ellsberg 1961, p. 669), like the one he suggests, “the question of the optimality of this behavior would gain more interest.”³⁷

Thus, all the elements of an articulated critique to Savage’s Bayesian view were already alluded to in the original 1961 article. This is also apparent in the first experimental papers testing Ellsberg’s claim. Indeed, experimental evidence was limited in the early years, and a clear interest in the topic became manifest only in the 1980s, after Kahneman and Tversky (1979) re-opened the debate on violations and particularly with attempts to extend prospect theory to uncertainty (Einhorn and Hogarth 1986) and axiomatic developments allowing for Ellsberg-type behaviour (Schmeidler 1989). But the few investigations in the laboratory of the 1960s and 1970s discussing violations of the axioms of rationality typically consisted of two distinct steps: first, experimental subjects were asked to make their choices, then, to value their own answers in view of the theory. MacCrimmon (1968) and Slovic and Tversky (1974) exposed experimental subjects to written arguments for and against Ellsberg and found that ambiguity aversion may persist even after exposure to theoretical arguments.

³⁶ Ellsberg’s claim is that Savage failed his test in February 1958. The two remained in contact later and Savage was sent a draft of Ellsberg’s dissertation in Spring 1962. There is no evidence, either in published works or in the private exchanges archived at the Leonard Jimmie Savage Papers at Yale University, that Savage ever denied Ellsberg’s claim, not even in his exchange with de Finetti (see Zappia 2018).

³⁷ Ellsberg suggests that a decision-maker should choose the action x associated with the highest value of the index $\rho E(x) + (1-\rho)\min(x)$, where $E(x)$, the expected utility of action x with respect to the “best guess” probability, coinciding with Savage’s prior probability. $E(x)$ is weighted with respect to the degree of confidence, ρ , encapsulating the influence of the ambiguity surrounding the decision process. If no ambiguity is perceived the rule coincides with Savage’s. In case ambiguity is perceived (i.e., $0 < \rho < 1$) the best guess distribution is weighted with Wald’s maxmin, assumed as the conventional decision rule for probabilistic ignorance. It must be noted that Eichberger and Kelsey’s (1999) axiomatization of a decision rule for ambiguity builds exactly on Ellsberg’s intuition.

And although the idea to check how much experimental subjects found axioms reasonable was considered controversial and later abandoned, these studies suggested a clear interest in both the empirical verification of “paradoxical” choices and their normative relevance. A main aim of these studies was to analyze whether individuals abandon principles or stick to them when these principles conflict with specific choices they find appealing (Camerer and Weber 1992, pp. 334-337).³⁸

It is then difficult to argue, as Gilboa substantially does, that Ellsberg was simply proposing a counterexample aiming to dispute the descriptive value of Savage. Instead, Ellsberg insisted that his examples could not be classified as mere descriptive anomalies of an otherwise acceptable normative theory. This is why he concentrated on “deliberate decisions,” acknowledging acceptance of the rationale of Savage’s reaction to Allais’s (1953) example and trying to provide a violation able to survive Savage’s test.³⁹ What is more, Ellsberg’s proposal for an alternative decision rule hints at a normative, rather than a descriptive, issue.

7. Daniel Ellsberg’s methodological critique of Savage

As noted, one cannot read Ellsberg’s famous paper simply as a presentation of a counter-example intended to highlight a descriptive violation of Savage’s theory, since many relevant methodological issues for normative falsification were pointed out from the outset. But Ellsberg’s attitude is even more evident in his doctoral thesis, submitted to the Harvard Economics Department in May 1962, and referred to as “a much fuller discussion of the nature, validation and functions of normative theories of choice” in reply to Roberts’s (1963) criticism (Ellsberg 1963, p. 337). The draft of the thesis circulated among major contributors to the debate—including de Finetti, Marschak and

³⁸ MacCrimmon and Larson (1979, pp. 382-384) summarized this first round of experimental studies—comparing individuals’ actual choices with their level of agreement with decision rules—with the claim that in “real situations of uncertainty,” such as stock market investments, deliberate Ellsberg-type violations occurred at such a high rate to demand further theoretical study on the non-additivity of beliefs. As noted by Mongin (2014, pp. 766-768), before Kahneman and Tversky (1979) empirical studies were used to test reflective reasoning also in Allais, confirming that the issue of deliberate choices was intended as crucial under both risk and uncertainty.

³⁹ As mentioned, Ellsberg did not refer to Allais in his paper and the discussion with Raiffa and Roberts contains no comparison between the paradoxes, testifying of the limited interest in Allais’s argument in the early 1960s. However, we shall see that the connection became apparent to Ellsberg, and worth dealing with, in his thesis.

particularly Savage, to whom it is dedicated—but due to Ellsberg’s decision not to enter academics⁴⁰ it remained unpublished until 2001, having basically no impact on the following developments.⁴¹

In his doctoral thesis, Ellsberg (2001 [1962]), tried to show that his ideas apply well beyond the paradox, from both a methodological and theoretical viewpoint. Moreover, while working on the thesis, Ellsberg realized that his attempt to criticize Savage was in agreement with a marginal, but well-articulated viewpoint already present in probability literature, that included scholars such as John Maynard Keynes (1973 [1921]), Bernard Koopman (1940), Irving Good (1952) and Cedric Smith (1961). He found it significant that these authors had provided a “theoretical approach that admit vagueness as an explicit factor without apology and provides a formal vocabulary to deal with it” (Ellsberg 2001 [1962], p. 10) and tried to show that he was working in their footsteps.

A close scrutiny of Ellsberg’s thesis cannot be provided here (but see Levi 2001 and Zappia 2016). However, the thesis elaborates further on the issue of empirical relevance of a normative theory. Ellsberg notes that while Savage’s perspective is purely normative, the test for confirming the validity of his theory he suggests may provide hints as to its empirical content. Indeed, when Savage asserted that introspective exercises such as asking how a decision-maker value an axiom are appropriate tests for normative propositions, he helped define “in operational terms the nature and the goals of a normative logic of choice” (Ellsberg 2001 [1962], p. 26).

A normative theory has empirical content, then, in so far as it makes it possible to describe the kind of observed behaviour that is incompatible with the theoretical proposition that individuals maximize SEU—that is, to identify probability profiles not conforming to the theory, for instance non-additive probability profiles, as in his urn examples. Here Ellsberg points out that this is a crucial step forward from the conventional exposition of von Neumann and Morgenstern’s EU, one that takes for granted that the axioms defining rational choice cannot be questioned on empirical grounds. As already noted, this was Marschak’s position, and indeed Ellsberg now criticizes Marschak for providing the methodological justification of this point when he claimed that “to discuss a set of norms of reasonable behavior ... is a problem in logic, not in psychology. It is a normative, not a

⁴⁰ Ellsberg, already at the RAND Corporation while working on the thesis, moved to the US Secretary of Defence in 1964, and never came back to academics. In the political arena he is better known for his releasing to the press of the Pentagon Papers in 1971. From then on he has been, and still is, an independent writer and political activist.

⁴¹ The correspondence included in the Savage Papers is illustrative of Savage’s sympathy for Ellsberg’s investigation. But Ellsberg’s inability to provide axiomatic justification for his result brought as a consequence Savage’s unwillingness to endorse a less firm subjectivist perspective while still fighting for the acceptance of the core of his viewpoint among statisticians (Zappia 2018).

descriptive problem” (Marschak 1951, p. 493, quoted in Ellsberg 2001 [1962], p. 27). To this position, Ellsberg notes that, while it is true that the consistency of a certain set of axioms is a logical, normative problem, and not a psychological, descriptive one, the degree of acceptance and authority of certain maxims of behaviour must have a descriptive component nonetheless. For Ellsberg (2001 [1962], p. 27), the extension of “*familiar rules of logic and arithmetic [to decision under uncertainty] is not a normative but a descriptive problem.*”

As seen in the previous sections, the current criticism of the Bayesian viewpoint concentrates on the extent to which a normative theory can dispense with empirical claims and still be authoritative: “Internal coherence of beliefs is important, but so is external coherence: having coherent beliefs that have nothing to do with evidence and data cannot be considered rational” (Gilboa et al. 2009, p. 292). It is striking to see its similarities with Ellsberg’s classic way of arguing against the Bayesian viewpoint. Ellsberg’s (2001 [1962], p. 27) position is clear cut: “internal consistency is not an adequate test.” A “system of logic” that satisfies “internal consistency” or certain “aesthetic considerations of form style, semantic usage,” but does not apply to the deliberate decisions of a decision-maker, may induce the decision-maker to prefer satisfying her preferences and comfortably let the axioms unapplied (Ellsberg 2001 [1962], p. 29).⁴² As a result, Ellsberg concludes, Savage’s basic assumptions concerning logical consistent behaviour can surely be considered “eminently reasonable,” but not “uniquely reasonable,” as Raiffa (1961) had suggested in his comments.

Ellsberg’s newly provided idea in the 1962 version of his argument is that Savage’s is a strict version of Bayesianism, unable to deal with ambiguity, and that a generalized Bayesian approach should be pursued. In the probabilistic framework he puts forward there is, on the one hand, the strictly-Bayesian decision-maker, who may keep in mind a whole set of “reasonably acceptable” probability distributions before acting, but who eventually settles upon a single distribution (or acts as if she did). The generalized-Bayesian decision-maker, on the other hand, retains all those probability distributions that do not definitely contradict her “vague” opinion when relevant information is perceived as ambiguous.⁴³

This constructive effort, only sketched out in 1961, resolves now into a taxonomy of different

⁴² Compare with Gilboa *et al.*’s (2009, p. 292) claim: “Internal coherence of beliefs is important, but so is external coherence. Having coherent beliefs that have nothing to do with evidence and data cannot be considered rational.”

⁴³ Ellsberg’s now much more elaborate positive analysis hinges on the formal structure guaranteed by Koopman’s and Good’s research on partially ordered probabilistic beliefs and constitutes the link with authors such as Levi (1980) and Gärdenfors and Sahlin (1982), who have worked on the philosophical background of representations of degrees of belief through a set of probability distributions.

decision rules, trying to find a role not only for “confidence” in probabilities, but also for “optimism/pessimism” about possible outcomes when probabilities are unavailable or deemed to be unreliable. The proposed classification tries to mix Keynes’s (1973 [1921]) considerations on confidence in a probability assessment with Hurwicz’s (1951) proposal to mitigate Wald’s insistence on worst outcomes, with attention also to best outcomes.⁴⁴ The search for alternative decision rules with normative content advocated by Ellsberg, then, substantially coincides with Gilboa’s quest for axiomatically based alternatives to the Bayesian approach.

Before coming to the conclusions, a last remark on Allais is worth making. We have seen that Ellsberg (1961) did not refer to Allais’s challenge of Savage, as he was assuming that the issue of ambiguity was related to the difficulty to elicit definite probabilities, a problem Allais’s examples do not present. Savage’s axioms were acceptable to Ellsberg as normative criteria in any situation with objective probabilities, like Allais’s. However, in a detailed section of the thesis, Ellsberg acknowledges that the pattern of violations Allais expects from his examples is “quite similar” to his own. Ellsberg comments on Savage’s solution to simplify Allais’s problem into a convenient representation before answering it, and examines whether such a “process of abstraction” suppress significant information in the original data. He notes that the idea on which Allais constructed his example was to contrast a prospect of a risky reward with an offer of a sure gain, and quotes from Allais (1979 [1953]) when he signalled “the considerable psychological importance that the advantage of certainty, as such, can have” (Ellsberg 2001 [1962], p. 257). Ellsberg argues that by turning away from the problem as presented to him by Allais, Savage is suppressing from his consideration the very difference that Allais considers essential: a “taste for certainty” that is crucial for Allais, is removed in Savage’s (1954, p. 102-103) table representation.⁴⁵ On careful consideration of Allais’s hypothetical examples, Ellsberg concludes, there is a similarity he did not notice in 1961 but he now finds of “considerable heuristic value:” it appears to him that “*Certainty is to Risk, as Risk*

⁴⁴ Ellsberg’s suggestion for a comprehensive decision rule is intended to accommodate a wide variety of rational behaviour attributable to individuals who deliberately violate Savage’s axioms. Using the same notation of footnote 30, Ellsberg suggests now that a decision-maker should choose the action x associated with the highest value of the index $\rho E(x) + (1-\rho)[\alpha \max(x) + (1-\alpha) \min(x)]$, where the best guess distribution is weighted with a mixture of max and min of outcomes with respect to optimism α (ranging from 0 to 1), as in Hurwicz (1951). It is worth noting that even this rule has been axiomatized in contemporary literature, as α -Maximization of EU (Ghirardato *et al.* 2004).

⁴⁵ Ellsberg (2001 [1962], p. 229) uses a similar argument against Raiffa, arguing that the way he offered the opportunity to choose mixed strategies in his replications of the urn examples banished ambiguity, in fact producing “crucially different” experiments from the ones he had suggested in his 1961 article.

is to Ambiguity” and a preference for certainty when risk is available can be equated to a preference for risk when ambiguity is to confronted (Ellsberg 2001 [1962], p. 263). In both cases, Ellsberg concludes, the change maybe so significant to certain decision-makers to induce them not to rely on Savage’s axioms, even upon reflection. The need for a new normative standard is evidenced once more.

8. Concluding remarks

This paper has tried to provide a comparison of critical views on the still dominant Bayesian approach to decision-making in economics. A brief review of the current criticism of Itzhak Gilboa and associates has shown that that a number of relevant issues for today’s theories of decision-making under uncertainty were already examined by Daniel Ellsberg in his classic criticism of Leonard Savage’s axiomatization.

I have argued, first, that the definition of rationality suggested by Gilboa in a recent paper (and in a series of previous works with co-authors, characterizing the core of today’s theory of decision under uncertainty) was already introduced by Savage in his *Foundations of Statistics*. Indeed, the application of this notion of rationality guided Savage in his examination of Allais, and may have guided him in his evaluation of Ellsberg, though there is no evidence that he rejected Ellsberg’s argument. Moreover, early experimental analysis testing the validity of axioms concentrated on deliberate choices.

Second, I have noted that Gilboa’s idea that Bayesianism as a normative theory of decision should be put to a test concerning decision-makers’ behaviours, and cannot be simply interpreted as the extension of formal logic to the uncertainty domain, was already advanced by Ellsberg in his early 1960s critique of Marschak and Savage. While it is generally agreed that Ellsberg provided the fundamental counter-example to Bayesian decision theory, Gilboa follows in a tradition that has misrepresented Ellsberg as interested only in showing descriptive violations of the theory, whereas instead he was arguing in favour of its abandonment as a normative theory of behaviour and of a redesign of rational choice theory away from a strict form of Bayesianism.

Most recent advances in decision theory under uncertainty are the outcome of more than fifty years of attempts to make sense of the Ellsberg Paradox, with “leading researchers ... still in the process of working out the nature and implications of what has come to be known as ‘ambiguity aversion’” (Machina *et al.* 2011). This has been a long theoretical process to which Ellsberg did not contribute further. But from a methodological viewpoint the core elements of today’s developments were already

put forward by Ellsberg himself. And even the historical background was evident to Ellsberg, who noted that the epistemic approach to probability put forward by Keynes in his *Treatise on Probability*—and related attempts to view probabilities as a partial order when analysing uncertainty—offered a compelling interpretational framework for an *ante litteram* criticism of Bayesianism.

REFERENCES

- Allais, M., 1953. Le comportement de l'homme rationnel devant le risque: Critique des postulats et axiomes de l'école américaine. *Econometrica*, 21, 503–46.
- Allais, M., 1979 [1953]. Fondements d'une théorie positive des choix comportant un risque et critique des postulats et axiomes de l'école américaine, *Économétrie*, pp. 257–331. Translated into English as: The foundations of a positive theory of choice involving risk and a criticism of the postulate and axioms of the American school. In: Allais and Hagen 1979, 27–145.
- Allais, M., and Hagen, O., eds., 1979. *Expected Utility Hypotheses and the Allais Paradox*. Dordrecht: Reidel Publishing Company.
- Al-Najjar N. I., and Weinstein, J., 2009. The ambiguity aversion literature: A critical assessment. *Economics and Philosophy*, 25, 249–284.
- Baumol, W. J., 1951. The Neumann-Morgenstern utility index: An ordinalist view. *Journal of Political Economy*, 59, 61–66.
- Becker, S. W., and Brownson, F. O., 1964. What price ambiguity? Or the role of ambiguity in decision making. *Journal of Political Economy*, 72, 63–73.
- Binmore, K., 2009. *Rational Decisions*. Princeton: Princeton University Press.
- Blume, L. E., and Easley, D., 2008. Rationality. In: L. E. Blume, and S. N. Durlauf, eds. *The New Palgrave: A Dictionary of Economics*, 2nd Edition. London: Palgrave Macmillan, 884–893.
- Brewer, K. R. W., 1963. Decision under uncertainty: comment. *Quarterly Journal of Economics*, 77, 159–61.
- Camerer, C., and Weber, M., 1992. Recent developments in modelling preferences: uncertainty and ambiguity. *Journal of Risk and Uncertainty*, 5, 325–70.
- de Finetti, B., 1951. Recent suggestions for the reconciliation of theories of probability. In: J. Neyman, ed. *Proceedings of the Second Berkeley Symposium on Mathematical Statistics and Probability*. Berkeley: University of California Press, 217–25

- de Finetti, B., 1964 [1937]. La prévision: Ses lois logiques, ses sources subjectives. *Annales De L'institut Henri Poincaré*, 7, 1-68. Translated into English as: Foresight: its logical laws, its subjective sources. In: H. E. Jr. Kyburg, and H. E. Smokler, eds. *Studies in Subjective Probability*. New York: Wiley, 93–158.
- Eichberger, J., and Kelsey, D., 1999. E-capacities and the Ellsberg Paradox. *Theory and Decision*, 46, 107–38.
- Eichberger, J., and Kelsey, D., 2009. Ambiguity. In: P. Anand, P. Pattanaik, and C. Puppe, eds. *The Oxford Handbook of Rational and Social Choice*. Oxford: Oxford University Press, 113–39.
- Einhorn, H. J. and Hogarth, R., 1986. Decision making under ambiguity. *Journal of Business*, 59, 225–50.
- Ellsberg, D., 1961. Risk, ambiguity, and the Savage axioms. *Quarterly Journal of Economics*, 75, 643–69.
- Ellsberg, D., 1963. Risk, ambiguity, and the Savage axioms: reply. *Quarterly Journal of Economics*, 77, 336–42.
- Ellsberg, D., 2001 [1962]. *Risk, Ambiguity and Decision*. New York: Routledge.
- Fellner, W., 1961. Distortion of subjective probabilities as a reaction to uncertainty. *Quarterly Journal of Economics*, 75, 670–89.
- Fellner, W., 1963. Slanted subjective probabilities and randomization: A reply to Howard Raiffa and K. R. W. Brewer. *Quarterly Journal of Economics*, 77, 676–90.
- Fishburn, P. C., and Wakker, P., 1995. The invention of the independence condition for preferences, *Management Science*, 41, 1130–44.
- Friedman, M., and Savage, L. J., 1948. The utility analysis of choices involving risk, *Journal of Political Economy*, 56, 279–304.
- Friedman, M., and Savage, L. J., 1952. The Expected-Utility Hypothesis and the measurability of utility. *Journal of Political Economy*, 60, 463–474.
- Gärdenfors, P., and Sahlin, N.-E., 1982. Unreliable probabilities, risk taking and decision making. *Synthese*, 53, 361–86.
- Ghirardato, P., Maccheroni, F., and Marinacci, M., 2004. Differentiating ambiguity and ambiguity attitude. *Journal of Economic Theory*, 118, 133–73.
- Gilboa, I., 2009. *Theory of Decision Under Uncertainty*. Cambridge: Cambridge University Press.
- Gilboa, I., 2015. Rationality and the Bayesian Paradigm, *Journal of Economic Methodology*, 22, 312–334.

- Gilboa, I., and Marinacci, M., 2016. Ambiguity and the Bayesian Paradigm. In: H. Arló-Costa, V. F. Hendricks, and J. van Benthem, eds. *Readings in Formal Epistemology*, vol. I. New York: Springer, 385–439.
- Gilboa, I., Postlewaite, A. W., and Schmeidler, D., 2008. Probability and uncertainty in economic modelling. *Journal of Economic Perspectives*, 22, 173–88.
- Gilboa, I., Postlewaite, A. W., and Schmeidler, D., 2009. Is it always rational to satisfy Savage's axioms? *Economics and Philosophy*, 25, 285–96.
- Gilboa, I., Postlewaite, A. W., and Schmeidler, D., 2012. Rationality of belief or: why Savage's axioms are neither necessary nor sufficient for rationality. *Synthese*, 187, 11–31.
- Gilboa, I., and Schmeidler, D., 1989. Maxmin expected utility with non-unique prior. *Journal of Mathematical Economics*, 18, 141–53.
- Gilboa, I., and Schmeidler, D., 2001. *A Theory of Case Based Decision Theory*. Cambridge: Cambridge University Press.
- Gillies, D., (2000). *Philosophical Theories of Probabilities*. London: Routledge.
- Giocoli, N., 2003. *Modeling Rational Agents: From Interwar Economics To Early Modern Game Theory*. Cheltenham: Edward Elgar Publishing.
- Giocoli, N., 2013. From Wald to Savage: homo economicus becomes a Bayesian statistician. *Journal of the History of the Behavioral Sciences*, 49, 63–95.
- Good, I. J., 1952. Rational decisions. *Journal of the Royal Statistical Society, Series B*, 14, 107–114.
- Guala, F., 2000. The logic of normative falsification: rationality and experiments in decision theory. *Journal of Economic Methodology*, 7, 59–93.
- Hacking, I., 1975. *The Emergence of Probability*. Cambridge: Cambridge University Press.
- Halevy, Y., 2007. Ellsberg revisited: An experimental study. *Econometrica*, 75, 503–536.
- Hammond, P. J., and Zank, H., 2014. Rationality and dynamic consistency under risk and uncertainty. In: M. Machina, and W. K. Viscusi, eds. *The Handbook of the Economics of Risk and Uncertainty*, Vol. I. Amsterdam: Elsevier, 41–98.
- Hands, D. W., 2015. Normative rational choice theory: past, present and future. Mimeo. University of Puget Sound. Available at SSRN: <http://ssrn.com/abstract=1738671>
- Hansen, L. P., 2014. Nobel Lecture: Uncertainty outside and inside economic models. *Journal of Political Economy*, 122, 945–87.
- Herfeld, D., 2015. Rational choice as a toolbox for the economist: an interview with Itzhak Gilboa. *Erasmus Journal for Philosophy and Economics*, 7, 116–141.
- Herfeld, D., 2018. From theories of human behavior to rules of rational choice: tracing a normative turn at the Cowles Commission, 1943–54. *History of Political Economy*, 2018, 50, 1-48.

- Heukelom, F., 2014. *Behavioural Economics: A History*. Cambridge: Cambridge University Press.
- Hirshleifer, J., and Riley, J. C., 1992. *The Analytics of Uncertainty and Information*. Cambridge: Cambridge University Press.
- Hurwicz, L., 1951. Some specification problems and application to econometric models (abstract). *Econometrica*, 19, 343–44.
- Jallais, S., and Pradier, P.-C., 2005. Changes in expected utility theory and the Allais experiment. In: P. Fontaine, and R. Leonard, eds. *The Experiment in the History of Economics*. London: Routledge, 21–42.
- Jeffreys, H. 1948. *Theory of Probability*. Oxford: Clarendon Press.
- Kahneman, D., and Thaler, R. H., 2006. Utility maximization and experienced utility. *Journal of Economic Perspectives*, 20, 221–34.
- Kahnemann, D., and Tversky, A., 1979. Prospect theory: An analysis of decision under risk. *Econometrica*, 47, 263–91.
- Kelsey, D., and Quiggin, J., 1992. Theories of Choice under Ignorance and Uncertainty, *Journal of Economic Surveys*, 6, 133–153.
- Keynes, J. M., 1973 [1921]. *A Treatise on Probability. The Collected Writings of John Maynard Keynes, Vol. VII*, edited by D.E. Moggridge. London: Macmillan.
- Klibanoff, P., Marinacci, M. and Mukerji, S., 2005. A smooth model of decision making under ambiguity. *Econometrica*, 73, 1849–92.
- Koopman, B. O., 1940. The axioms and algebra of intuitive probability. *Annals of Mathematics*, 41, 269–92.
- Levi, I., 1980. *The Enterprise of Knowledge*. Cambridge, MA: MIT Press.
- Levi, I., 2001. Introduction. In: Ellsberg 2001, ix–xxxvii.
- Luce, D., and Raiffa, H., 1957. *Games and Decisions*. New York: Wiley.
- MacCrimmon, K., 1968. Descriptive and normative implications of the decision-theory postulates. In K. Borch, K., and J. Mossin, eds. *Risk and Uncertainty: Proceedings of a Conference Held by the International Economic Association*. London: Macmillan and Co, 3–23.
- MacCrimmon, K., Larsson, S., 1979. Utility theory: Axioms versus ‘paradoxes’. In: Allais and Hagen 1979, 333–409.
- Machina, M., 1987. Choice under uncertainty: Problems solved and unsolved. *Journal of Economic Perspectives*, 1, 121–54.
- Machina, M., Ritzberger, K. and Yannelis, N. 2011. Introduction to the symposium issue. *Economic Theory*, 48, 219–221.

- Machina, M., and Siniscalchi, M., 2014. Ambiguity and ambiguity aversion. In: M. Machina, and W. K. Viscusi, eds. *The Handbook of the Economics of Risk and Uncertainty*, Vol. I. Amsterdam: Elsevier, 729–807.
- Marschak, J., 1946. Neumann's and Morgenstern's new approach to static economics. *Journal of Political Economy*, 54, 97–115.
- Marschak, J., 1950. Rational behaviour, uncertain prospects, and measurable utility. *Econometrica*, 18, 111–42.
- Marschak, J., 1951. Why “should” statisticians and businessmen maximize “moral expectation”? In: J. Neyman, ed. *Proceedings of the Second Berkeley Symposium on Mathematical Statistics and Probability*. Berkeley: University of California Press, 493–506.
- Mongin, Philippe. 2014. “Le paradoxe d’Allais: comment lui rendre sa signification perdue.” *Revue économique*, vol. 65 (5): 743–79.
- Moscatti, I., 2016. Retrospectives: How economists came to accept expected utility theory. The case of Samuelson and Savage. *Journal of Economic Perspectives*, 30, 219–236.
- Mukerji, S., 2009. Foundations of ambiguity and economic modelling. *Economics and Philosophy*, 25, 297–302.
- Mukerji, S. and Tallon, J.-M., 2004. An overview of economic applications of David Schmeidler’s models of decision making under uncertainty. In: I. Gilboa, ed. *Uncertainty in economic theory*. London: Routledge, 283–302.
- Raiffa, H., 1961. Risk, ambiguity, and the Savage axioms: a comment. *Quarterly Journal of Economics*, 75, 690–694.
- Raiffa, H., and Schlaifer, R., 1961. *Applied Statistical Decision Theory*. Boston: Harvard University Graduate School of Business Administration.
- Ramsey, F. P., 1964 [1931]. Truth and probability. In: H. E. Jr. Kyburg, and H. E. Smokler, eds. *Studies in Subjective Probability*. New York: Wiley, 61–92.
- Roberts, H. V., 1963. Risk, ambiguity, and the Savage axioms: a comment. *Quarterly Journal of Economics*, 77, 327–36.
- Samuelson P. A., 1938. A note on the pure theory of consumer’s behaviour. *Economica*, 5, 61–71
- Samuelson, P. A., 1950. Probability and the attempts to measure utility. *Economic Review*, 1, 167–73.
- Savage, L. J., 1954. *The Foundations of Statistics*. New York: Wiley.
- Savage, L. J., 1970. Reading suggestions for the Foundations of Statistics. *The American Statistician*, 24, 23–27.
- Schmeidler, D., 1989. Subjective probability and expected utility without additivity. *Econometrica*, 57, 571–87.

- Shafer, G., 1976. *A Mathematical Theory of Evidence*. Princeton: Princeton University Press.
- Slovic, P., and Tversky, A., 1974. Who accepts Savage's axioms? *Behavioural Science*, 19, 368–78.
- Smith, C. A. B., 1961. Consistency in statistical inference. *Journal of the Royal Statistical Society, Series B*, 23, 1–25.
- Sugden, R., 1991. Rational choice: A survey of contributions from economics and philosophy, *Economic Journal*, 101, 751–85.
- Starmer, C., 2005. Normative notions in descriptive dialogues. *Journal of Economic Methodology*, 12, 277–89.
- von Neumann, J., and Morgenstern, O., 1947. *Theory of Games and Economic Behaviour*. 2nd Edition. Princeton: Princeton University Press.
- Wakker, P., 2008. Uncertainty. In: L. E. Blume, and S. N. Durlauf, eds. *The New Palgrave: A Dictionary of Economics*, 2nd Edition. London: Palgrave Macmillan, 428–439.
- Wallis, A. W., 1980. Memorial service tribute. In L. J. Savage, *The Writings of Leonard Jimmie Savage: A Memorial Selection*. Washington, DC: American Statistical Association and Institute of Mathematical Statistics, 11–24.
- Zappia, C. 2016. Daniel Ellsberg and the validation of normative propositions. *Æconomia. History, Methodology, Philosophy*, 6, 57–79.
- Zappia, C. 2018. Inside critics of the Bayesian mainstream in economics and statistics: Letters to Leonard Jimmie Savage. Mimeo, University of Siena.