

Carlo Zappia

## Daniel Ellsberg and the Validation of Normative Propositions

---

### Avertissement

Le contenu de ce site relève de la législation française sur la propriété intellectuelle et est la propriété exclusive de l'éditeur.

Les œuvres figurant sur ce site peuvent être consultées et reproduites sur un support papier ou numérique sous réserve qu'elles soient strictement réservées à un usage soit personnel, soit scientifique ou pédagogique excluant toute exploitation commerciale. La reproduction devra obligatoirement mentionner l'éditeur, le nom de la revue, l'auteur et la référence du document.

Toute autre reproduction est interdite sauf accord préalable de l'éditeur, en dehors des cas prévus par la législation en vigueur en France.

**revues.org**

Revues.org est un portail de revues en sciences humaines et sociales développé par le Cléo, Centre pour l'édition électronique ouverte (CNRS, EHESS, UP, UAPV).

---

### Référence électronique

Carlo Zappia, « Daniel Ellsberg and the Validation of Normative Propositions », *Æconomia* [En ligne], 6-1 | 2016, mis en ligne le 31 mars 2016, consulté le 29 avril 2016. URL : <http://oeconomia.revues.org/2276> ; DOI : 10.4000/oeconomia.2276

Éditeur : Association Æconomia

<http://oeconomia.revues.org>

<http://www.revues.org>

Document accessible en ligne sur : <http://oeconomia.revues.org/2276>

Ce document est le fac-similé de l'édition papier.

Les contenus d'Æconomia sont mis à disposition selon les termes de la Licence Creative Commons Attribution - Pas d'Utilisation Commerciale - Pas de Modification 4.0 International.

# Daniel Ellsberg and the Validation of Normative Propositions

Carlo Zappia\*

In the history of decision theory Daniel Ellsberg is known because his seminal paper “Risk, ambiguity and the Savage axioms” presented the counterexample to Bayesian decision-making that got the normative value of the theory into trouble. An assessment of Ellsberg’s contribution in the more encompassing framework of his long unnoticed 1962 doctoral thesis suggests that, although he did not take part to the ensuing debate, Ellsberg provided a thorough philosophical and methodological background for his critique of Savage’s axiomatization of decision theory. By concentrating mainly on Ellsberg’s analysis of decision-making in his thesis, this paper examines the way he conceived of the possibility to test normative propositions and tries to identify the kind of normative value he attributed to his own suggested solution for the Ellsberg Paradox.

Keywords: Ellsberg (Daniel), decision theory, ambiguity, Ellsberg Paradox

## Daniel Ellsberg et la validation des propositions normatives

Au sein de l’histoire de la théorie de la décision, Daniel Ellsberg est connu pour son article fondateur « Risk, ambiguity and the Savage axioms », dans lequel il élabore le contre exemple à la théorie bayésienne de la décision qui a ébranlé la valeur normative de celle-ci. Une étude de la contribution de Ellsberg qui prend en compte le contexte plus large de sa thèse de doctorat de 1962, longtemps passée inaperçue, suggère que, s’il n’a pas pris part aux débats subséquents, Ellsberg appuie sa critique de l’axiomatisation de Savage de la théorie de la décision sur un arrière plan conceptuel, philosophique et méthodologique, rigoureux. En se concentrant sur la théorie de la décision présente dans sa thèse, l’article examine la manière dont Ellsberg a conceptualisé la possibilité de tester les propositions normatives et tente d’identifier le type de valeurs normatives qu’il attribue à sa propre solution au paradoxe d’Ellsberg.

Mots-clés : Ellsberg (Daniel), théorie de la décision, ambiguïté, paradoxe d’Ellsberg

JEL: B21, B41, D81

---

\*University of Siena, zappia@unisi.it

A previous version of this paper was presented at the Conference “Economics and Psychology in Historical Perspective”, Maison des Sciences Economiques, Paris, December 2014, and at the Annual Conference of the Italian Association for the History of Political Economy (STOREP), Campus Luigi Einaudi, Torino, June 2015. Comments by Richard Arena, Robin Hogarth, Ivan Moscati, Enrico Petracca and two anonymous referees are gratefully acknowledged.

Together with the Allais Paradox, the Ellsberg Paradox is the best-known example of a violation of the prescriptions of the theory of choice in economics literature. But even though Daniel Ellsberg's 1961 article "Risk, Ambiguity and the Savage Axioms" is textbook reference in chapters on decision under risk and uncertainty (Mas-Colell *et al.*, 1995) and extensively quoted—with hundreds of articles in economic journals referring to the Ellsberg Paradox each year<sup>1</sup>—the philosophical background underlying it remains totally unknown. For instance, very few attempts to devise a decision rule that can solve the paradox refer to Ellsberg's own view about a possible solution, despite the fact that most of later studies on decision-making under uncertainty originate as attempts to account for his notion of ambiguity (Camerer and Weber, 1992). This is not very surprising though, since Ellsberg's 1961 article is rather sketchy about motivations and possible developments. Moreover, he never became an academic. An analyst at the RAND Corporation since 1959, after completing his Ph.D. in Economics at Harvard in 1962, Ellsberg joined the US Defence Department to work on the strategic aspects of the escalation of the Vietnam War, and did not keep contributing to economics. His decision in the late 1960s to reveal to the press several classified documents about US Government's misconduct during the war—came to be known as the 'Pentagon Papers' after *The New York Time* started publishing them in June 1971—put an end to his career. Ellsberg then became a political activist and writer, and he still is as of today.<sup>2</sup>

In his long unpublished Ph.D. thesis, though, Ellsberg ([1962] 2001) dealt with a number of issues related to his paradoxical example and this makes it possible to provide a thorough assessment of his contribution to the theory of decision-making. This paper examines Ellsberg's own motivations in presenting the paradox in his doctoral thesis by concentrating mainly on the methodological justifications of Ellsberg's critique of the Bayesian mainstream encapsulated in Savage's *Foundations of Statistics* (Savage, 1954). The paper presents Ellsberg's rationale for the validation of normative propositions and discusses how Ellsberg motivated his subscription to a normative rejection of the strict application of the Bayesian approach. It is shown that Ellsberg accepted a procedure of normative falsification he attributed to Savage himself, arguably emerging from Savage's rationale for his rejection of the Allais Paradox. In doing so Ellsberg offered a paradigmatic example of what has been termed a "quasi-empirical" test of rationality (Guala, 2000).

The paper aims to show that, although it is rarely quoted in the literature on rationality and normative falsification—mainly devoted to

<sup>1</sup> A search for articles citing Ellsberg's 1961 paper in 2015 using Google Scholar returns 488 quotations (search made on February 1, 2016).

<sup>2</sup> See Ellsberg (2003; 2006). On the controversial aspects of Ellsberg's life see Samuels (2002).

the Allais Paradox and the ensuing debate (Starmer, 2005; Mongin, 2009)—Ellsberg’s analysis appears even more significant than Allais’s. Crucially, Ellsberg was aware that, in order to avoid that his argument were classified among instances of irrational behaviour, he could not confined his investigation to the presentation of a counter-example, and that an alternative normative theory to the consolidated mainstream was needed. An investigation of how Ellsberg motivated and defended his paradox in the thesis evidences that he accomplished this aim through an analysis of both the probabilistic foundations of his proposal and its characteristics as a normative guide for decision-making.<sup>3</sup>

The paper is structured as follows. The next section illustrates the way Ellsberg, initially endorsing the rationality axioms of von Neumann and Morgenstern’s theory, started considering uncertainty issues in decision-making, and it presents his famous paradox. Section 2 illustrates the normative turn in decision theory delineated by Savage’s (1954) formulation of expected utility. Section 3 introduces the issue of counter-examples for normative falsification and analyzes the significance Ellsberg attributed to his paradoxical results from a methodological viewpoint. Section 4 presents the way in which Ellsberg interpreted the procedure of falsification he was proposing. Section 5 offers a few concluding remarks.

## 1. The Origins of the Ellsberg Paradox

Ellsberg graduated in economics at Harvard under the supervision of John Chipman in 1952, with a thesis on the contribution of von Neumann and Morgenstern to the theory of rational choice under risk. An outgrowth of his thesis, Ellsberg (1954) compared von Neumann and Morgenstern’s notion of expected utility (EU) to the classical notion of a cardinal utility function and clarified their different derivation and significance in a way still praised in the economics literature (Fishburn, 1989; Giocoli, 2003). The paper testifies of Ellsberg’s initial adherence to a mainstream view of decision-making. He considered the axiom of “Strong Independence” introduced by Samuelson (1952) as “indubitably the most plausible” of the set of axioms necessary to derive an EU function, since “even people who did not follow it in practice would probably admit, on reflection, that they should” (Ellsberg, 1954, 544).

After serving as a US Marine, Ellsberg started graduate studies at Harvard University in 1957 where he began working on decision-

---

<sup>3</sup> Throughout the paper the use of the term normative does not refer to any ethical motivation behind theorizing, but simply to the action-guiding or prescriptive motivations of analysis, that is, to the logic of decision-making, as it became usual among decision theorists since the early 1950s (on this distinction see Hands, 2012).

making under uncertainty. During his graduate studies he joined the RAND Corporation as a strategic analyst. While at RAND the issue of how to represent the state of information of a decision-maker acting under uncertainty became the major focus of his investigation. His recollection is that, in his office in the late 1950s, by means of “endless trial and error with paper and pen”, he was

searching for choices between gambles—actions with uncertain outcomes—that would give operational meaning behaviourally for the first time to Frank Knight’s distinction between “risk”—roulette-like gambles with “known”, precise probabilities—and “uncertainty”, when no such probabilities were “known” (Ellsberg, 2011, 223).

This approach would mean criticizing the work of Leonard Savage who, following on the probabilistic approach of Bruno de Finetti ([1937] 1964), had shown how von Neumann and Morgenstern’s EU theory could be extended from objective to subjective probabilities, and then made appropriate to account for decisions conditioned on uncertain events as much as on risky ones. Savage provided the axiomatic structure to be satisfied for representing the beliefs of individual agents by means of a (subjective) additive probability distribution. On the grounds of his representation theorem, even when objective probabilities of events conditioning the payoffs of individual agents are not available, one can elicit from choices the probabilities to be used for the maximization of (subjective) EU. The distinction between risk and uncertainty then lost relevance among decision theorists and subjective EU theory came to be characterized as the new “Bayesian” mainstream (Raiffa and Schlaifer, 1961).<sup>4</sup>

It is conceivable that Ellsberg’s (1958) review of Davidson *et al.* (1957)—one of the first works intended to test experimentally the significance of the Bayesian approach under uncertainty—is the first publication in which his discomfort about the negation of the distinction between risk and uncertainty implicit in Savage’s axioms is alluded to. Ellsberg (1958, 1010) praised the reviewed volume for its application of the experimental approach, but saw a major limitation in “the nature of the hypotheses the authors have chosen to test”. Ellsberg maintained that

whereas they [the authors] tend to suggest that their basic hypothesis is uniquely plausible, covering reasonable behaviour on all occasions of uncertainty, I believe that there are important classes of uncertain situations in which normal people will systematically *violate* these axioms, and in which other hypotheses will better describe their behaviour. (Ellsberg, 1958, 1010-1011)

---

<sup>4</sup> On the adoption of the term Bayesian in the theory of decision-making and statistical analysis in the 1950s see Fienberg (2006).

Ellsberg's conclusion, that "human behaviour under uncertainty is undoubtedly more various than has been imagined in that discussion", is suggestive of his forthcoming investigation.

Ellsberg's famous counter-example to Bayesian decision theory was presented at the December Meeting of the *Econometric Society* in St. Louis in 1960 and published as part of a symposium on "Decision under Uncertainty" in the November 1961 issue of the *Quarterly Journal of Economics*, which also included a paper by William Fellner (1961) on the "distortion" towards non-additivity of subjective probability distributions. A comment by Howard Raiffa (1961) offered arguments to the contrary from the mainstream viewpoint. After a brief round of discussion (Brewer, 1963; Roberts, 1963; Ellsberg, 1963; Fellner, 1963) and an experimental paper demonstrating the relevance of the violations Ellsberg had described (Becker and Brownson, 1964) the issue remained unaddressed for long. It was acknowledged that there certainly were situations in which it was troublesome to elicit standard probabilities from choices (Brewer and Fellner, 1965), but Raiffa's main point—that violations simply show that people need to be taught about the theory, and that when taught they would act accordingly—was accepted as the most significant outcome of the debate. As a matter of fact, abandoning the Savage's axioms would mean depriving decision theory of its foundation elements, and no alternative could be envisaged yet. As a result, interest in the topic faded away "simply because researchers at the time were helpless to address [it]" (Machina, 2001, xxxix).

Ellsberg completed his Ph.D. in Economics in 1962. His thesis, entitled *Risk, Ambiguity and Decision*, drew on the 1961 article. In fact, most of the thesis was written after the article had been published. Ellsberg referred to the thesis in his 1963 reply to Roberts, hinting at the fact that it presented a more elaborate justification of his critique of Savage on many issues. But the thesis was not commented on in decision theory debate, and remained completely ignored until its publication in 2001.<sup>5</sup>

Before turning to the methodological issues Ellsberg saw in the acceptance of the Bayesian viewpoint, a cursory presentation of the Ellsberg Paradox is in order. In his 1961 essay Ellsberg presented two examples claiming that the choice behaviour of most decision-makers considering these examples, tested in non-experimental conditions, failed to conform to the prescriptions of the subjective EU approach. In a first example with two urns, Ellsberg found that many decision-makers refrain from betting on the drawing of balls from an "ambiguous" urn when a "risky" one is available. He concentrated on the choice of betting on the drawn of a red (or black) ball from either an

---

<sup>5</sup> Ellsberg's doctoral thesis is now reproduced in its original, unedited version in Ellsberg ([1962] 2001), including an introduction by Isaac Levi and a bibliographical note edited by Mark Machina.

urn containing 100 red and black balls in an unknown proportion, or an urn known to contain 50 red and 50 black balls. When confronting each urn separately, individuals tested by Ellsberg tended to be indifferent between betting on red or black. But when asked to choose the urn from which they would prefer to bet that a red (or a black) ball would be drawn, most respondents revealed to Ellsberg a preference for the second, risky but unambiguous urn, with respect to the first, ambiguous one. Since this kind of choice is incompatible with the assignment of a sharp additive probability distribution,<sup>6</sup> Ellsberg's intuition was that agents paid serious attention to second-order considerations, such as vagueness and scantiness of information. He concluded that the degree of confidence in a probability assessment cannot be excluded from analysis.

In a second example, with a single urn whose composition is partly known and partly unknown, many decision-makers refrain from betting on the drawing of balls constituting the unknown part of the urn. Ellsberg imagined a single urn containing 90 balls, with 30 red balls and 60 either black or yellow balls. Ellsberg reported that when asked to bet on the drawing either of a red ball or of a black ball individuals usually prefer to bet on the red ball. But when asked to bet on the drawing either of a red *and* yellow or of a black *and* yellow they prefer to bet on black and yellow. He showed that this entailed that individuals were inclined to violate Savage's axioms. Indeed, Ellsberg's second example yields a direct test of Savage's "sure-thing principle", the axiom Savage intended as a substitute for von Neumann and Morgenstern's independence principle outside risky environments such as lotteries.<sup>7</sup>

Ellsberg concluded that the choices he had observed were contradictory with the sharp probability assignment assumed by Savage on

---

<sup>6</sup> The indifference of decision-makers between betting on the red or blue drawn from urn *I* or *II* examined separately means that their subjective probabilities  $p$  are such that  $p(\text{red}_I)=p(\text{black}_I)=1/2$  and  $p(\text{red}_{II})=p(\text{black}_{II})=1/2$ . However when they choose to bet on red (or black) from the second urn they reveal  $p(\text{red}_{II})>p(\text{red}_I)$  (or  $p(\text{black}_{II})>p(\text{black}_I)$ ). Following on Ellsberg, the non-additive probability approach of Schmeidler (1989) proposes an axiomatic system for subjective probabilities such that  $p(\text{red}_I)+p(\text{black}_I)<p(\text{red}_I \cup \text{black}_I)$ . In current decision theory the individual who prefers to bet on urn *II* is said to show ambiguity aversion (Eichberger and Kelsey, 2009).

<sup>7</sup> A suitable representation of the example shows that the choice of whether to bet on the drawing of red or the drawing of black (choice *I* or *II*) and on red and yellow or black and yellow (choice *III* or *IV*) should—for a decision-maker obeying to Savage's sure thing principle—be independent of the event that a yellow ball is drawn. This is reported by Ellsberg not to be the normal choice. Ellsberg's (1961, 655) used a column representation that is analogous to the one introduced by Savage (1954, 103) to rationalize his change of attitude towards Allais's example. While Savage's used the column representation to show his confidence in the sure thing principle, Ellsberg's contention was that, in his one-urn example, it shows its apparent violations by the subjects he tested.

the basis of his system of axioms. Crucially, Ellsberg also reported that in both examples a considerable number of subjects did not change their choices even after thorough reconsideration of the significance of the theory, thus showing to be “unrepentant violators” of the Savage Axioms. His examples, he argued, could not be simply suggesting descriptive violations. As we shall see, his idea was that this gave stronger significance to his counter-examples than to previous ones like Allais’. Moreover, since the group of people he tested included a number of fellow decision theorists, their reluctance to repent violations showed that the normative content of the theory was not as compelling as generally agreed.<sup>8</sup> In other words, urn examples are offered in order to show that the choices of “a number of people who are not only sophisticated but reasonable” do not conform to a sharply defined numerical probability, even upon reflection. The attack on the consolidating mainstream represented by subjective EU theory is explicitly put as follows (Ellsberg, 1961, 646):

I propose to indicate a class of choice-situations in which many otherwise reasonable people neither wish nor tend to conform to the Savage postulates, nor to the other axiom sets that have been devised.

## 2. Savage and the Normative Turn in Decision Theory

Even though he presented his remarks on the reasons why violations are so relevant as tentative, Ellsberg (1961) forcefully argued that he did not consider Savage’s approach to be a normative guide in ambiguous contexts, and suggested that this holds true for all deliberate violators of the theory. Furthermore, he claimed that “once certain patterns of ‘violating’ behaviour [are] distinguished and described in terms of a specified decision rule” (Ellsberg, 1961, 669), like the one he examined in the second part of his essay, economic theorists have to wonder what kind of normative status can be attributed to this alternative decision rule.

---

<sup>8</sup> Violators included Norman Dalkey, Jacob Marschak and Savage himself, among others (Ellsberg, 1961, 656). Ellsberg illustrated his examples at faculty seminars he gave in the late 1950s, and used the presentations to test people in the audience. The people tested during these seminars, and in private conversations, “under absolutely non-experimental conditions” (Ellsberg, [1962] 2001, 134), were mostly faculty members at the Universities of Chicago, Harvard, Northwestern and Yale, and at the RAND Corporation. Ellsberg’s ([1962] 2001, xlix) claim is that he showed his urn examples to Savage in February 1958, when Savage reacted “encouraging me to believe that the arguments presented ... deserved serious consideration”. Ellsberg did not test Savage later, but remained in contact with him and, in a private letter to Savage of May 1962, he introduced the draft of his thesis sent to him as “a 400-page letter to you, designed to change your mind” (Leonard J. Savage Papers, MS 695, Box 11/260, Manuscript and Archives, Yale University).



Ellsberg's interest in normative issues, as a result, is apparent from his 1961 article. Mainly, he wanted to avoid that his examples were classified as descriptive anomalies of an otherwise acceptable normative theory. Thus, he made clear from the outset that he intended to concentrate on "deliberate decisions", that is, "choices that people make when they take plenty of time to reflect over their decision, in the light of the [Savage] postulates" (Ellsberg, 1961, 646). This is an acknowledgement that, in principle, he accepted Savage's reaction to Maurice Allais' (1953) example: decision theory has a normative status that cannot be criticized on the grounds of descriptive failures *per se*, since a "reasonable" decision-maker who instinctively violates the theory when confronting Allais's example is supposed to reverse her choice after "thorough deliberation" (Savage, 1954, 102). As is well known, when Allais presented his example to Savage at an informal meeting during the 1952 International Colloquium on Risk in Paris, Savage expressed preferences contradicting his axioms. Upon reflection, though, he changed his mind and claimed: "it seems to me that in reversing my preference ... I have corrected an error" (Savage, 1954, 103).<sup>9</sup> It was Savage's position that only "deliberate decisions" are apt to reveal degrees of belief compatible with a probability measure, that is, that must show the kind of consistency the axioms of the theory entail.

Savage's methodological point can be considered the hallmark of the normative turn in the theory of choice that emerged as the outcome of the debate on empirical evidence in the early years of decision theory after von Neumann and Morgenstern (Giocoli, 2013; Hands, 2015). In confronting critics of the empirical significance of EU such as Samuelson (1950) and Baumol (1951), Savage (1951) had endorsed Marschak's (1951) position that the axioms justifying the maximization of EU under risk should be considered like elements of logic and mathematics, and claimed that empirical evidence *per se* was a peripheral subject of investigation. Moreover, in his joint work with Friedman, Savage had formulated a notion of "plausibility" of axioms, unrelated to empirical issues. Friedman and Savage (1952, 463) claimed that the possibility to regard the maximization of EU as a

---

<sup>9</sup> This was the conventional reaction of the "American School" when challenged by the Allais Paradox. As a matter of fact, at least until the consolidation of experimental economics in the 1970s the status attributed to Allais's counter-example was that of an empirical violation similar to a logical mistake. 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. For a historical reconstruction of the Paris Conference, mostly from Allais' viewpoint, see Jallais and Pradier (2005). Heukelom (2014) illustrates how Savage formulated his answer to Allais through an examination of the private correspondence between the two after the conference. Mongin (2014) investigates the normative component of Allais' critique.

“maxim of behaviour” depends “not on its empirical verification for the economic behaviour of men at large, but on its acceptability to individuals who are particularly concerned with such decisions, as a rule guiding ‘wise’ behaviour in the face of uncertainty”. Their fundamental point was that if a “reader” considers the postulates that are sufficient for the derivation of EU in the light of the illustration they provide in the paper “he will concede that the principle is not one he would *deliberately* violate”, and this is why there is “reason for supposing that people do actually tend to avoid *flagrant* violations of the principle” (Friedman and Savage, 1952, 669, emphasis added).<sup>10</sup> To sum up: the claim that something that cannot be empirically proved may still hold true derives from the coherence of the EU hypothesis “with the body of economic theory and, more important, from the plausibility of postulates ... rather than from repeated success in prediction” (Friedman and Savage, 1952, 474). Moreover, emphasis is placed on deliberate action, rather than on action *per se*, and deliberate violations are excluded on the basis of a logical argument.

On these grounds, in his *Foundations*, Savage presented his derivation of EU theory under uncertainty explicitly as an attempt to extend logical reasoning to situations of uncertainty.<sup>11</sup> His primary aim was not to provide an empirical theory for predicting human behaviour, but to develop logical tools for deciding between alternative courses of action. Decisions made in the face of uncertainty, Savage claimed, require “formal reasoning”. The purpose of his book, then, is “to discuss the implications of reasoning for the making of decisions” (Savage, 1954, 6). Savage used the analogy with logic already referred to by Marschak: since reasoning is usually associated with logic, but the implication of “ordinary logic” are limited in the face of uncertainty,

---

<sup>10</sup> Friedman and Savage’s (1948) early work on the significance of von Neumann and Morgenstern’s (1947) approach—featuring the now classic illustration of a utility function representing different attitudes toward risk—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” (Friedman and Savage, 1948, 282). The 1952 article concentrated instead on the significance of the axioms being proposed, and investigated their potential for normative economics. While the two papers are usually considered as part of a coherent effort at systematization (Jallais and Pradier, 2005; Heukelom, 2014), a fundamental step towards normativism that is put forward in the second paper did not appear in the previous one (Starmer, 2005).

<sup>11</sup> With the notable exception of Giocoli (2013), the fact that Savage’s 1954 *Foundations* were under elaboration at least since 1949, when a preliminary report was presented at a meeting of the Econometric Society in Boulder (Savage, 1950), is seldom noticed in the literature on the Allais-Savage exchange that insists on Savage’s retreat from an empirical interpretation of the theory (Jallais and Pradier, 2005). On the evolution of Savage’s thought towards the *Foundations*, see also Moscati (2016), who clarifies that Savage defended the normative appeal of his sure-thing principle already in his correspondence with Samuelson, in mid-1950. See also Samuelson (1952).

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”.

As just noted, Savage’s viewpoint entails analyzing “deliberate” behaviour only. Despite Allais’ (1953) critique, Savage (1954) argued that theories of rational behaviour have a normative status and should not be modified in the light of evidence showing descriptive violations, which can only be taken as indication of irrational choices. Savage suggested that one could term irrational only decision-makers whose behaviour is inconsistent with their own theory, and that empirical evidence should be examined with respect to decision-makers readiness to revise their choices in the light of their theories of rationality.

Guala (2000) terms the approach codified by Savage a “quasi-empirical test of rationality”. Since empirical evidence is no longer available as a theoretical argument, it remains to explain how a normative theory can be confirmed or refuted. In his joint work with Friedman, and when contrasting the anomalous evidence put forward by Allais in his *Foundations*, Savage provided a critical method for testing normative theories. Savage’s approach, then, can be seen as a solution to the particularly puzzling question of how to distinguish between making a mistake and not obeying subjective EU axioms:

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 his initial impression of the situation or to accept the implications of the theory for it. (Savage, 1954, 102)

Once the falsifying evidence is collected and the plausibility of axioms is therefore questioned, an individual who maintains preferences that are in conflict with a given normative principle “must abandon, or modify, the principle”. Indeed, Savage concluded, “that kind of discrepancy seems intolerable in a normative theory”.

### 3. Ellsberg on the Logic of Normative Falsification

As seen in the previous section, Savage’s approach then amounts to a procedure of falsification of a normative theory. But from a methodological viewpoint an issue of comparable relevance is how to move further. Progress in methodological terms cannot simply rely on falsification. The intuition which suggested the plausibility of the rejected axioms has been put under fire, but what is the intuition underlying the counter-example? Falsifying a normative theory does not necessarily imply questioning the normative adequacy of a theory, since research programmes are typically endowed with a protective belt that is designed to isolate the effect of a counter-example. Following

on the Lakatosian claim that refutation also requires the suggestion of a viable alternative (Lakatos, 1968), it can then be argued that in order to supersede a theory of decision-making one needs to provide an alternative model of rational choice, a “rival theory” able to explain the counter-example and to offer a conjecture about the origins of the violation (i.e. with a rationality principle behind it).<sup>12</sup> This process, which has been termed the “logic of normative falsification and theory-improvement”, entails that the procedure of criticizing a normative theory and the growth of normative knowledge cannot be separated (Guala, 2000, 72-74).<sup>13</sup>

The history of modern decision theory is usually seen as a process of generalization sparked by the acceptance of the Allais Paradox. So called non-EU theories concentrated on the weak link of EU theory, identified with the independence axiom, and were formulated in order to account for its violation. This process was of course influenced by the spreading of the experimental method in economics, which made possible the collection of empirical data to be used not only for testing the EU hypothesis but also to imagine new alternative hypotheses. From a historical viewpoint it must be recalled, though, that Allais’ own conjecture for a rival theory had to wait 25 years before being extensively formulated, and that the alternatives to EU theory have been developed primarily for descriptive rather than normative purposes. The normative status of EU theory remained substantially intact at least until the early 1990s (Fishburn and Wakker, 1995).

As already noted, the dramatic events of Ellsberg’s life made it impossible for him to contribute to the developments related to his paradoxical results. Although for different reasons his case is similar to Allais’: the impact of the Ellsberg Paradox on normative issues was almost irrelevant for long. The first theory of decision under uncertainty explicitly mentioning Ellsberg’s results as a motivating factor is that formulated by Einhorn and Hogarth (1986), who still mainly concentrated on descriptive issues, but only Schmeidler (1989) was able to provide a proper axiomatic generalization of subjective EU accounting for Ellsberg’s results. However, the 1962 thesis addresses a number of relevant issues which have arisen in the following years and are part of the current debate in decision theory.<sup>14</sup> In what fol-

---

<sup>12</sup> In placing emphasis on the significance of Lakatos’s claim, Guala (2000, 71) notes that Morgenstern’s (1979) late reflection on Allais’ paradox pointed at the inability of Allais to formulate a principle of rationality of a broader content than the one he intended to criticize.

<sup>13</sup> Lakatos’ (1970, 94) defence of Popperian methodology against allegations of “naive falsificationism” requires that crucial counter-evidence “can be recognized as such among the scores of anomalies only *with hindsight*, in the light of some superseding theory”. According to this argument, “*refutation without an alternative shows nothing but the poverty of our imagination in providing a rescue hypothesis*”.

<sup>14</sup> See for instance Binmore (2009) and Gilboa (2009).

lows, specific regard is given to the way Ellsberg's thesis deals with the validation of normative propositions.

From the very beginning of his thesis Ellsberg professes his great admiration for Savage's work, one that goes further than a praise for Savage's ability to devise a comprehensive axiomatic structure for decision-making. He values Savage's perspective, first, because in following de Finetti's operational definition of probability Savage paved the way to the settling of the philosophical controversy about the notion of subjective probability. Second, Ellsberg praises the fact that Savage's approach favoured the much-awaited introduction of a collection of techniques for the measurement of definite subjective probabilities and utilities, like the one put forward by Davidson *et al.* (1957). But Ellsberg's respect for Savage is also motivated by methodological reasons.

Ellsberg notes that Savage's perspective is purely normative, but that the test he proposes for confirming the validity of his theory gives it also empirical content. Indeed, Ellsberg argues, Savage's proposal that "when certain maxims are presented for your consideration, you must ask ... how would you react if you noticed yourself violating them" (Savage, 1954, 7, quoted in Ellsberg, [1962] 2001, 24) suggests "definite empirical predictions" as to how violators would react. Savage then subscribed to the idea that it is possible to describe the kind of observed behaviour that is incompatible with the theoretical proposition. Ellsberg finds it "admirable" that "Savage exposes his own theory to empirical 'refutation' by a given reader".<sup>15</sup> To assert that "such introspective exercises are appropriate tests" for normative propositions is "to *define* in operational terms the nature and the goals of a normative logic of choice" (Ellsberg, [1962] 2001, 26).

Ellsberg's argument is that this is a crucial step forward from the conventional exposition of von Neumann and Morgenstern's EU that takes for granted that the axioms defining rational choice cannot be questioned on empirical grounds. Marschak in particular is criticized for providing the methodological justification of this point in his claim that "to discuss a set of norms of reasonable behaviour ... is a problem in logic, not in psychology. It is a normative, not a descriptive problem" (Marschak, 1951, 493, quoted in Ellsberg, [1962] 2001, 27). Ellsberg objects 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 cer-

---

<sup>15</sup> As a matter of fact, the maximization of subjective EU axiomatized by Savage has empirical content in so far as it rules out certain probability profiles. The operational procedure of eliciting probabilities from choices makes possible to rule out those probability profiles not conforming to the theory, and then to conclude that a standard probability function cannot explain the choices of an un-repentant violator of Savage sure-thing principle (see Ellsberg, [1962] 2001, 40-41). See also footnote 6 above.

tain maxims of behaviour must have a descriptive component nonetheless. The extension of “*familiar rules of logic and arithmetic [to decision under uncertainty] is not a normative but a descriptive problem*” (Ellsberg, [1962] 2001, 27). A “system of logic” may satisfy “internal consistency” or certain “aesthetic considerations of form, style, semantic usage”. But if it does not apply to the deliberate decisions of a decision-maker, it may induce the decision-maker to find “more rational to satisfy his preferences and let the axioms satisfy themselves”, to use a sentence that Samuelson (1950) made famous when still an opponent of EU theory (Ellsberg, [1962] 2001, 29). Ellsberg forcefully argues that basic assumptions concerning logically consistent behaviour like the Savage axioms can be considered “eminently reasonable”, but not “*uniquely* reasonable”, as for instance Raiffa and Schlaifer (1961) suggested in their book systematizing the Bayesian viewpoint.

Savage, then, must be praised because, after endorsing Marschak’s point on the logic of decision theory, he moved further to suggest how to test a normative theory like his own.<sup>16</sup> But while Allais’ counter-example for risky situations has already been rejected as a normative falsifier by Savage and his followers on these grounds (Luce and Raiffa, 1957, 32), Ellsberg thinks his own, designed to investigate the domain of uncertainty, can pass the test. So the identification of a clear-cut counter-example is instrumental to a discussion of the normative value of the theory. Ellsberg does not use the word falsifiability, but his main criterion for evaluating a normative theory is to ask whether its prescriptions can be carried out. The methodological argument he proposes follows a typically (even though crude) falsificationist approach, with the aim to reject Savage’s theory, and to provide heuristic advice about how to develop a new normative theory, arguably “less likely to be confronted with deliberate and persistent violation” when tested (Ellsberg, [1962] 2001, 25).

The fact that this is his aim, can help explain why Ellsberg concentrates on the new notion of ambiguity instead of on uncertainty proper. He claims that he is interested in situations in which prior probabilities are difficult to elicit, since related to instances like investment decisions or war-gaming, and he is well aware that the notion of uncertainty hinted at by Frank Knight (1921)—one in which probabilities are “unknown”—cannot be restricted to simplified cases like the one he discusses, for there is no “true” uncertainty in his “ambiguous” urn.<sup>17</sup> But in order to avoid being caught in a philosophical dis-

---

<sup>16</sup> Although concerned with the Allais Paradox, Mongin (2009, 328) makes the same point when arguing that “a thoughtful adherence to the paradoxical choice under the same informational circumstances” implies that “Savage’s alleged condition for rationality may well work like a boomerang”.

<sup>17</sup> In his definition of the “spectrum of uncertainties” a theory of decision must encompass, he lists a range of events that starts with the effects of “familiar pro-

cussion about the meaning of “true” uncertainty, he restricts attention to how decision-makers act in urn problems, where the issue is related to the vagueness of the choice context more than to uncertainty proper. In the ambiguous urn the proportion of balls of different colours is not properly unknown, but known within limits. As a result, Duhemian problems of indeterminateness—concerning whether the main hypothesis under investigation has been falsified or the falsification concerns the auxiliary hypotheses used in the construction of the counter-example—are limited as much as possible by confining the decision context to simple urn problems.

Ellsberg’s argument also suggests that he accepts a distinction between falsification and rejection. A counter-example may provide a necessary condition for the falsification of certain normative propositions, but this is not enough to reject a theory. In order to reject a theory, a conjecture about the origins of the anomaly, and a sketch of a new theory capable to give account of the anomaly, is needed as well (Guala, 2000). It is apparent that Ellsberg’s aim is to provide heuristic advice about how developing a new normative theory and that his investigation does not stop after finding psychological elements justifying the violations of the current theory. The following section of this paper concentrates on his suggestions.

#### 4. The Methodological Foundations of the Ellsberg Paradox

As already noted, the doctoral thesis submitted to Harvard Economics Department in April 1962 provides the philosophical background of Ellsberg’s critique of Savage, concentrating on the significance of the process of normative testing. This section tries to show that the thesis presents also a thorough discussion of the contemporary literature on decision-making, including which theoretical developments Ellsberg envisages as viable alternative to subjective EU theory, that is, able to provide normative rules for deliberate violators. Crucially, Ellsberg offers a cogent presentation of the intuition behind his paradoxical result and elaborates his proposal as a way forward.

As in the 1961 article, Ellsberg opens his doctoral investigation with a reference to Knight’s (1921) distinction between measurable and unmeasurable uncertainty. He observes that while contemporary developments of subjective EU theory motivated a growing scepticism about the behavioural significance of Knightian uncertainty, he

---

duction decisions or well-known random processes (coin flipping or roulette, certain chemical/physical phenomena)”, but ends with “the results of long-range Research and Development, the performance of a new President, the consequences of major departmental reorganization, the tactics of an unfamiliar opponent, or the secrets concealed by an efficient security system” (Ellsberg, [1962] 2001, 5). See also footnote 18.

finds counter-intuitive the denial that there are probability relations about which decision-makers feel less “sure” as compared to others: in uncertain contexts, the degree of confidence in a probability assessment must have a role. However, in the thesis the emphasis is more on John Maynard Keynes than on Knight, and on the authors who followed Keynes in stressing vagueness and imprecision of subjective probability assessments—none of whom were quoted in the article. Ellsberg ([1962] 2001, 9-11) attributes to Keynes’s ([1921] 1973) *Treatise on Probability* the merit of having introduced “formally the notion of *non-comparability* of beliefs”, and recognizes that the notion of weight of argument “seems closely related to our notion of ‘ambiguity’”. He also recalls that Keynes placed emphasis on the idea that, in situations where information is perceived to be vague, the traditional approach to probability is inadequate; and that Keynes’s discussion of non-numerical probabilities inspired Bernard O. Koopman (1940) and Irving J. Good (1950) in their investigations of an axiomatic structure compatible with intervals of probability priors. He praises the fact that the theories of intuitive probability professed by Keynes and his followers constitute “a theoretical approach that admits vagueness as an explicit factor without apology and provides a formal vocabulary for discussing it”.

Ostensibly, Ellsberg’s assessment of previous treatments of decision-making under uncertainty is not motivated by the necessity to provide an introductory overview of the state of the art. Rather, it is meant to show that, before the new subjectivist mainstream consolidated, there had been a lively tradition of thought discussing vagueness in decision-making from a probabilistic viewpoint. Ellsberg places emphasis on the fact that a number of scholars shared the cognitive unease of the theorist who accepts a theoretical approach in principle, but finds it restrictive in its application, and therefore is inclined to question its normative value and to provide heuristic advice about how to develop a new theory.<sup>18</sup> As recalled by Levi (2001) in his introduction to the publication of Ellsberg’s doctoral thesis, this is overall the road taken by all authors endorsing a Bayesian set-up, but not persuaded by the strict Bayesian perspective codified by Savage’s axioms, and eager to provide a generalized Bayesian perspective.<sup>19</sup>

---

<sup>18</sup> George Shackle’s (1949) non-probabilistic theory of decision, cited by Arrow (1951) as the sole formalized alternative to probabilistic decision-making, but dismissed for its lack of axiomatic bases, is also a major reference in the thesis. Ellsberg (1961, 643) had criticized Shackle’s outright rejection of numerical probabilities. But in the thesis he now admits that “when ambiguity is extreme” Shackle’s “somber reflections” seem “too ominously relevant to the very circumstances upon which this study focuses to be dismissed” (Ellsberg, [1962] 2001, 16-17). On the links between Ellsberg’s ambiguity and Shackle’s uncertainty see Basili and Zappia (2010) and Zappia (2014).

<sup>19</sup> Among these authors, see Levi (1974), Gardenförs and Sahlin (1982) and Walley (1991).



That Ellsberg's aim was much more ambitious than simply presenting a counter-example, and that he wanted to offer the backbone of a rival theory, is confirmed by the way in which the thesis defends an interval-valued probability approach. In particular two aspects are worth noting here: the first, concerning the kind of consistency in decision-making his viewpoint can entail, and the second, the kind of decision rule devised to replace the maximization of EU. As for the first aspect, Ellsberg reminds that de Finetti ([1937] 1964) had solved the problem of getting numerical probabilities that represent degrees of belief in a certain event through an analysis of the conditions under which a decision-maker would be prepared to bet on the event, and that although he did not use this approach in the *Foundations*, Savage (1962) started using it in his later presentations of the subjective approach. De Finetti's identification of a system of betting prices apt to elicit "definite probabilities" from choices included the requirement that these prices should be "coherent", in order to avoid a Dutch Book made against the decision-maker. Ellsberg ([1962] 2001, 55) accepts de Finetti's approach and the Bayesian methodology of eliciting probabilities from choices, but rejects the generality of de Finetti's "condition of coherence". Ellsberg argues that while it can be admitted that a decision-maker is in principle prepared to bet on any event, there is no rationale in the assumption that the decision-maker is always willing to take each side of the bet conditioned on that event.

For purposes of measurement, de Finetti regarded probabilities as prices, and assumed that the highest price an individual is willing to pay for betting in favour of an event coincides with the lowest price she is willing to accept for taking the other side of the bet. While no inequality between upper and lower betting prices is contemplated by de Finetti, Ellsberg ([1962] 2001, 68) contends,

the explicit assumption ruling this inequality out as a possibility for reasonable behaviour seems fairly gratuitous; at least, I cannot see that it has any direct intuitive appeal although it would follow as a theorem from the Savage postulates.<sup>20</sup>

Ellsberg is endorsing here the approach of deriving lower and upper betting probabilities presented by Cedric Smith at the Annual Meeting of the Royal Statistical Society in October 1960 (Smith, 1961), approvingly quoting Good, who, in his discussion on Smith's paper, remarked that the "reluctance to gamble" that is allowed by Smith

---

<sup>20</sup> Ellsberg's contention is that probabilities that can be derived from the choices of a deliberate violator of Savage axioms in his urn example are "indefinite" between limits, but cannot be termed "irrational" by means of the consistency argument, as this does not apply only to a set of "precise, definite beliefs." Indeed, "beliefs that must be treated as 'indefinite' within limits can still be precise enough to determine decisions in betting, and susceptible of quantitative expression in terms of inequalities" without the decision-maker falling prey to a Dutch book (Ellsberg, [1962] 2001, 88).

when he identifies a non-betting interval “seems intuitively to correspond to recognizable behaviour-patterns associated with circumstances in which opinions are exceptionally indefinite.”<sup>21</sup> Ellsberg ([1962] 2001, 88) concludes, then, that it is possible

to announce a set of betting quotations that is comfortably ‘coherent’, precluding any danger of falling prey to a Dutch book, yet which is not consistent with the existence of a definite probability measure over all events.

The second qualifying aspect of Ellsberg’s positive contribution in the thesis is the suggestion for an alternative decision rule. Moving forward from the proposals already examined in the 1961, in his 1962 thesis Ellsberg tries to accomplish the more relevant task of justifying the proposed decision rule on the basis of an approach that allows for intuitively more realistic statements, such as that the decision-maker does not know the probability distribution of the pay-off relevant events. This aim is pursued in two steps. The first step is to specify the representation of a probabilistic framework in which the decision-maker is allowed to keep in mind a whole set of “reasonably acceptable” probability distributions before acting. A Bayesian decision-maker accepting indeterminateness acts as if she retains all those probability distributions that do not definitely contradict her “vague” opinion, especially when relevant information is perceived as unreliable and contradictory. Once this is accepted, second-order considerations related to the degree of confidence in a probability assessment or with an optimistic/pessimistic attitude in choosing action can be represented through a (convex) set of probabilities.

Ellsberg observes that a number of alternatives to Savage’s maximisation of subjective EU have been devised to deal with complete ignorance and that an extensive examination of the axiomatic structure underlying these decision rules is available (Chernoff, 1954; Luce and Raiffa, 1957). His second step, then, is to concentrate on Hurwicz’s  $\alpha$ -maximin and to follow Hurwicz’s (1951b) suggestion, applying it to a context of partial ignorance.<sup>22</sup> Ellsberg examines which values of  $\alpha$  make the behaviour of a decision-maker compatible with the paradoxical choices of decision-makers confronting his examples. Thus, he starts from the idea that a set of distributions over the states

---

<sup>21</sup> It is worth noting that even de Finetti and Savage were very much impressed by Smith’s argument. In de Finetti and Savage (1962)—a joint paper written in Italian, never translated into English—they discussed at length Smith’s perspective, somewhat approvingly. On this point see Feduzi *et al.* (2014).

<sup>22</sup> Hurwicz’s (1951a) decision rule—better known as the Arrow-Hurwicz criterion after Arrow and Hurwicz (1972)—considers the behaviour associated with Wald’s (1945) maximin as overly pessimistic and introduces a parameter representing the degree of optimism of a decision-maker who is completely ignorant of the relevant probabilities. If  $\alpha$ , a coefficient ranging from 0 to 1, represents optimism the Hurwicz’s index is:  $\alpha \max(x) + (1-\alpha) \min(x)$ . Both the worst and the best possible outcomes associated with action  $x$  are considered.

of the world define partial ignorance and applies the Hurwicz Criterion to the restricted set of plausible distributions.<sup>23</sup> It is also worth noting that Ellsberg suggests that an axiomatic rationale can be found for this procedure. Indeed, he ([1962] 2001, 184) points out that Hurwicz (1951a) and Arrow (1953) had indicated a number of axioms on the basis of which “any criterion satisfying a certain subset of these requirements takes into account *only the minimum and maximum utility associated with each act*”. So, the use of the best/worst pairs is justified by Ellsberg because of its normative appropriateness. This accords well with the emphasis the thesis lays on the search for a new normative rationale underlying decision-making under uncertainty.

## 5. Concluding Remarks

In a recent paper on the way mainstream decision theory understands rationality when dealing with uncertainty, Gilboa *et al.* (2008, 286-287) argue:

The Bayesian approach could .... be viewed as an elegant but imperfect method for representation of uncertainty, one among many to be used depending on the application. Indeed, this is the way that it is viewed by many in diverse fields such as statistics, philosophy and computer science. However, within economic theory the Bayesian approach is the sole claimant to the throne of rationality.

Daniel Ellsberg offered the counter-example to Bayesian decision-making that got the normative value of the theory into trouble. But his 1961 paper was never interpreted as an attempt to contribute to the foundations of a new normative theory. Ellsberg’s involvement with the military, first, and his brave attempt to put an end to the Vietnam War with the release of the Pentagon Papers, later, made impossible for him to devote time to the defence of his viewpoint on methodological grounds.

An assessment of Ellsberg’s contribution in the more encompassing framework of his long unnoticed 1962 doctoral thesis adds to our understanding of the role the Ellsberg Paradox may have had in the contest of a theoretical approach to decision-making. This paper has argued that—following in a tradition of thinkers that were critical of the maximisation of Bernoulli’s ([1738] 1954) “moral expectations” even before von Neumann and Morgenster’s EU theory—Ellsberg provided a thorough methodological and theoretical background for his critique. Ellsberg insisted that the phenomena he concentrated on could not be explained through the postulation of additional psychological effects, rather they require normative ideas allowing for irre-

---

<sup>23</sup> It should be noted that certain recent developments in decision theory elaborate on the intuition behind the Hurwicz criterion exactly in the direction suggested by Ellsberg (Ghirardato *et al.*, 2004; Eichberger and Kelsey, 2009).

ducible second order characteristics like ambiguity, something admitted in contemporary decision theory only recently.

Ellsberg did not refer to falsifiability, but the methodological argument he proposed followed a falsificationist approach, with the aim to reject Savage's theory and to provide heuristic advice about how to develop a new normative theory, arguably less likely to be confronted with deliberate violation when tested. The paper has tried to show that Ellsberg's application of normative falsification was not naïve, showing awareness that a counter-example provides a necessary condition for the falsification of a normative theory, but it is not a sufficient condition to reject it. As a matter of fact, Ellsberg's conjecture about the origins of the anomaly is intended to offer a sketch of a new normative theory capable to account for what in his view no longer appears as an anomaly.

## References

- Allais, Maurice. 1953. Le comportement de l'homme rationnel devant le risque: critique des postulats et axiomes de l'école américaine. *Econometrica*, 21(4): 503-546.
- Allais, Maurice and Ole Hagen (eds). 1979. *Expected Utility Hypotheses and the Allais Paradox*. Dordrecht: Reidel Publishing Company.
- Arrow, Kenneth J. 1951. Alternative Approaches to the Theory of Choice in Risk-Taking Situations. *Econometrica*, 19(4): 404-437.
- Arrow, Kenneth J. 1953. Hurwicz's Optimality Criterion for Decision Making under ignorance, *Technical Report 6*, Stanford University.
- Arrow, Kenneth J. and Leonid Hurwicz. 1972. Decision Making Under Ignorance. In Charles F. Carter and Jim L. Ford (eds), *Uncertainty and Expectations in Economics. Essays in Honour of G. L. S. Shackle*. Oxford: Basil Blackwell, 1-11.
- Basili, Marcello and Carlo Zappia. 2010. Ambiguity and Uncertainty in Ellsberg and Shackle. *Cambridge Journal of Economics*, 34(3): 449-474.
- Baumol, William J. 1951. The Neumann-Morgenstern Utility Index—An Ordinalist View. *Journal of Political Economy*, 59(1): 61-66.
- Becker, Selwyn W. and Fred O. Brownson. 1964. What price ambiguity? Or the role of ambiguity in decision making. *Journal of Political Economy*, 72(1): 62-73.
- Bernoulli, Daniel. [1738] 1954. Exposition of a new theory on the measurement of risk. *Econometrica*, 22(1): 23-36.
- Binmore, Ken. 2009. *Rational Decisions*. Princeton: Princeton University Press.
- Brewer, K. R. W. 1963. Decision under uncertainty: comment. *Quarterly Journal of Economics*, 77(1): 159-161.
- Brewer, K. R. W. and William Fellner. 1965. The slanting of subjective probabilities: agreements on some essentials. *Quarterly Journal of Economics*, 79(4): 657-663.

- Camerer, Colin F. and Martin Weber. 1992. Recent developments in modelling preferences: uncertainty and ambiguity. *Journal of Risk and Uncertainty*, 5(4): 325-370.
- Chernoff, Herman. 1954. Rational selection of decision functions. *Econometrica*, 22(4): 422-443.
- Davidson, Donald, Patrick Suppes, and Sidney Siegel. 1957. *Decision Making. An Experimental Approach*. Stanford: Stanford University Press.
- de Finetti, Bruno. [1937] 1964. La prevision: 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 Henry E. Kyburg and Howard E. Smokler (eds), *Studies in Subjective Probability*. New York: Wiley, 97-156.
- de Finetti, Bruno and Leonard J. Savage. 1962. Sul modo di scegliere le probabilità iniziali. *Biblioteca del Metron, Serie C*, 1: 81-147.
- Eichberger, Jurgen and David Kelsey. 2009. Ambiguity. In Paul Anand, Pasanta Pattanaik and Clemens Puppe (eds), *The Handbook of Rational and Social Choice*. Oxford: Oxford University Press, 113-139.
- Einhorn, Hillel J. and Robin M. Hogarth. 1986. Decision making under ambiguity. *Journal of Business*, 59(4): 225-250.
- Ellsberg, Daniel. 1954. Classic and current notions of 'measurable utility'. *Economic Journal*, 64(255): 528-556.
- Ellsberg, Daniel. 1958. Review of *Decision Making. An Experimental Approach* by Davidson, D., Suppes, P. and S. Siegel. *American Economic Review*, 48(5): 1009-1011.
- Ellsberg, Daniel. 1961. Risk, ambiguity, and the Savage axioms. *Quarterly Journal of Economics*, 75(4): 643-669.
- Ellsberg, Daniel. 1963. Risk, ambiguity, and the Savage axioms: reply. *Quarterly Journal of Economics*, 77(2): 336-342.
- Ellsberg, Daniel. [1962] 2001. *Risk, Ambiguity and Decision*. New York: Routledge.
- Ellsberg, Daniel. 2003. *Secrets: A Memoir of Vietnam and the Pentagon Papers*. New York: Penguin Books.
- Ellsberg, Daniel. 2006. *Daniel Ellsberg: Extended Biography* (written for the Right Livelihood Award process), available at [www.ellsberg.net](http://www.ellsberg.net) (consulted on February 1, 2016).
- Ellsberg, Daniel. 2011. Introduction to the symposium issue: notes on the origins of the Ellsberg urns. *Economic Theory*, 48(2): 221-227.
- Fellner, William. 1961. Distortion of subjective probabilities as a reaction to uncertainty. *Quarterly Journal of Economics*, 75(4): 670-689.
- Fellner, William. 1963. Slanted Subjective Probabilities and Randomization: A Reply to Howard Raiffa and K. R. W. Brewer. *Quarterly Journal of Economics*, 77(4): 676-690.
- Feduzi, Alberto, Jochen Runde, and Carlo Zappia. 2014. De Finetti on uncertainty. *Cambridge Journal of Economics*, 38(1): 1-21.
- Fienberg, Stephen E. 2006. When did Bayesian inference become 'Bayesian'? *Bayesian Analysis*, 1(1): 1-40.
- Fishburn, Peter C. 1989. Retrospective on the utility theory of von Neumann and Morgenstern. *Journal of Risk and Uncertainty*, 2(2): 127-158.

- Fishburn, Peter C. and Peter Wakker. 1995. The invention of the independence condition for preferences. *Management Science*, 41(7): 1130-1144.
- Friedman, Milton and Leonard J. Savage. 1948. The utility analysis of choices involving risk. *Journal of Political Economy*, 56(4): 279-304.
- Friedman, Milton and Leonard J. Savage. 1952. The Expected-Utility hypothesis and the measurability of utility. *Journal of Political Economy*, 60(6): 463-474.
- Gardenförs, Peter and Nils-Eric Sahlin. 1982. Unreliable probabilities, risk taking and decision making. *Synthese*, 53(3): 361-386.
- Ghirardato, Paolo, Maccheroni, Fabio and Massimo Marinacci. 2004. Differentiating ambiguity and ambiguity attitude. *Journal of Economic Theory*, 118(1): 133-173.
- Gilboa, Itzhak. 2009. *Theory of Decision Under Uncertainty*. Cambridge: Cambridge University Press.
- Gilboa, Itzhak, Andrew W. Postlewaite, and David Schmeidler. 2008. Probability and uncertainty in economic modelling. *Journal of Economic Perspectives*, 22(3): 173-188.
- Giocoli, Nicola. 2003. Modeling Rational Agents: From Interwar Economics to Early Modern Game Theory. Aldershot: Edward Elgar.
- Giocoli, Nicola. 2013. From Wald to Savage: *homo economicus* becomes a Bayesian statistician. *Journal of the History of the Behavioural Sciences*, 49(1): 63-95.
- Good, Irwin J. 1950. *Probability and the Weighing of Evidence*. London: Charles Griffin.
- Guala, Francesco. 2000. The logic of normative falsification: rationality and experiments in decision theory. *Journal of Economic Methodology*, 7(1): 59-93.
- Hands, D. Wade 2012. The positive-normative dichotomy and economics. In Uskali Mäki (ed.), *Philosophy of Economics*. Amsterdam: Elsevier, 219-239.
- Hands, D. Wade 2015. Normative rational choice theory: past, present and future. University of Puget Sound. Available at SSRN: <http://ssrn.com/abstract=1738671>
- Heukelom, Floris. 2014. *Behavioural Economics: A History*. Cambridge: Cambridge University Press.
- Hurwicz, Leonid. 1951a. Optimality criteria for decision making under ignorance. *Cowles Commission Discussion Paper No. 370*, Chicago.
- Hurwicz, Leonid. 1951b. Some specification problems and application to econometric models (abstract). *Econometrica*, 19(3): 343-344.
- Jallais, Sophie and Pierre-Charles Pradier. 2005. Changes in Expected Utility Theory and the Allais Experiment. In Philippe Fontaine and Robert Leonard (eds), *The Experiment in the History of Economics*. London: Routledge, 21-42.
- Keynes, John M. [1921] 1973. A Treatise on Probability. The Collected Writings of John Maynard Keynes, Vol. VII. London: Macmillan.
- Knight, Frank H. 1921. *Risk, Uncertainty and Profit*. Chicago: University of Chicago Press.

- Koopman, Bernard O. 1940. The axioms and algebra of intuitive probability. *Annals of Mathematics*, 41(2): 269-292.
- Lakatos, Imre. 1968. Changes in the problem of inductive logic. In Imre Lakatos (ed.), *The Problem of Inductive Logic*. Amsterdam: North Holland, 315-417.
- Lakatos, Imre. 1970. Falsification and the methodology of scientific research programmes. In Imre Lakatos and Alan Musgrave (eds), *Criticism and the Growth of Knowledge*. Cambridge: Cambridge University Press, 91-195.
- Levi, Isaac. 1974. On indeterminate probabilities. *Journal of Philosophy*, 71(13): 391-418.
- Levi, Isaac. 2001. Introduction. In Daniel Ellsberg, *Risk, Ambiguity and Decision*. New York: Routledge, ix-xxxvii.
- Luce, Duncan and Howard Raiffa. 1957. *Games and Decisions*. New York: Wiley.
- Machina, Mark. 2001. Bibliographical note. In Daniel Ellsberg, *Risk, Ambiguity and Decision*. New York: Routledge, xxxix-xxlviii.
- Marschak, Jacob. 1951. Why 'should' statisticians and businessmen maximize 'moral expectation'? In Jerzy Neyman (ed.), *Proceedings of the Second Berkeley Symposium on Mathematical Statistics and Probability*. Berkeley: University of California Press, 493-506.
- Mas-Colell, Andreu, Michael D. Whinston, and Jerry R. Green. 1995. *Microeconomic theory*. New York: Oxford University Press.
- Mongin, Philippe. 2009. Duhemian themes in expected utility theory. In Anastasios Brenner and Jean Gayon (eds), *French Studies in the Philosophy of Science*. New York: Springer, 303-357.
- Mongin, Philippe. 2014. Le paradoxe d'Allais: comment lui rendre sa signification perdue. Paper presented at the Conference Economics and Psychology in Historical Perspective. *Revue Economique*, 65(5): 743-779.
- Morgenstern, Oskar. 1979. Some reflection on utility. In Maurice Allais and Ole Hagen (eds), *Expected Utility Hypothesis and the Allais Paradox*. Dordrecht: Reidel, 157-184.
- Moscatti, Ivan. 2016. How economists came to accept expected utility theory. The case of Samuelson and Savage. *Journal of Economic Perspectives*, forthcoming.
- Raiffa, Howard. 1961. Risk, ambiguity, and the Savage axioms: a comment. *Quarterly Journal of Economics*, 75(4): 690-694.
- Raiffa, Howard and Robert Schlaifer. 1961. *Applied Statistical Decision Theory*. Boston: Harvard University Graduate School of Business Administration.
- Roberts, Harry V. 1963. Risk, ambiguity, and the Savage axioms: a comment. *Quarterly Journal of Economics*, 77(2): 327-336.
- Samuels, Warren J. 2002. Daniel Ellsberg: analysis and redemption of an economist. *History of Economic Ideas*, 10(2): 123-154.
- Samuelson, Paul A. 1950. Probability and the attempts to measure utility. *Economic Review*, 1(3): 167-173.

- Samuelson, Paul A. 1952. Probability, utility, and the independence axiom. *Econometrica*, 20(4): 670-678.
- Savage, Leonard J. 1950. The role of personal probability in statistics (abstract). *Econometrica*, 18(2): 183-184.
- Savage, Leonard J. 1951. The theory of statistical decision. *Journal of the American Statistical Association*, 46(253): 55-67.
- Savage, Leonard J. 1954. *The Foundation of Statistics*. New York: Wiley.
- Savage, Leonard J. 1962. Bayesian statistics. In Robert F. Machol and Paul Gray (eds), *Recent Developments in Decision and Information Processes*. MacMillan: New York, 161-194.
- Schmeidler, David. 1989. Subjective probability and expected utility without additivity. *Econometrica*, 57(3): 571-587.
- Shackle, George. L. S. 1949. *Expectations in Economics*. Cambridge: Cambridge University Press.
- Slovic, Paul and Amos Tversky. 1974. Who accepts Savage's axioms? *Behavioural Science*, 19(6): 368-378.
- Smith, Cedric. A. B. 1961. Consistency in statistical inference. *Journal of the Royal Statistical Society, Series B*, 23(1): 1-25.
- Starmer, Chris. 2005. Normative notions in descriptive dialogues. *Journal of Economic Methodology*, 12(2): 277-289.
- von Neumann, John and Oskar Morgenstern. 1947. *Theory of Games and Economic Behaviour*, 2<sup>nd</sup> ed. Princeton: Princeton University Press.
- Wald, Abraham. 1945. Statistical decision functions which minimize the maximum risk. *Annals of Mathematics*, 46(2): 265-280.
- Walley, Peter. 1991. *Statistical Reasoning with Imprecise Probabilities*. London: Chapman and Hall.
- Zappia, Carlo. 2014. Non-Bayesian decision theory ahead of its time: the case of G. L. S. Shackle. *Cambridge Journal of Economics*, 38(5): 1133-1154.