

Daniel Ellsberg on the Ellsberg Paradox

CARLO ZAPPIA

Department of Economics and Statistics
University of Siena

zappia@unisi.it

Abstract

Even though Daniel Ellsberg's 1961 article "Risk, ambiguity and the Savage axioms" is well-known and increasingly quoted in current decision theory, introducing the counterexample to Bayesian decision-making that got the normative value of Savage's theory into trouble, its philosophical background remains totally unknown. This paper examines Ellsberg's motivations in presenting his critique first to his fellow decision theorists at Harvard and RAND in the late 1950s and it goes into his reasons for giving a philosophical justification and defence of the paradox in his doctoral thesis of 1962. By concentrating mainly on Ellsberg's all-encompassing analysis of decision-making in his thesis, the paper shows that a number of relevant issues connected to the paradox can be thrown light on. These range from its historical background to the way to test the normative value of decision theory through experiments, and a taxonomy of decision rules based on alternative probabilistic set-ups. Crucially, the paper argues that Ellsberg subscribed to a generalised version of the Bayesian approach, one that informs the developments of the multiple prior approach in current decision theory, but finds its origins in Keynes's *Treatise on Probability*.

September 2015

“So, what’s the paradox?” Convincing them [the laymen] that they do not always act as if they assigned precise numerical probabilities to uncertain events needs no demonstration for nearly anyone (other than ordained Bayesian statisticians).

Ellsberg (2011, p. 221)

1. Introduction

Daniel Ellsberg is a truly unconventional figure. In 2011, at the age of 80, he was interviewed in major American television networks and newspapers as a supporter of Julian Assange’s release of classified US documents through the website Wikileaks. Ellsberg’s involvement was motivated by the fact that he was the man who forty years earlier had released to the press several classified documents – came to be known as the “Pentagon Papers” – that showed that the Johnson Administration had systematically lied to the US Congress about a subject of “transcendent national interest and significance” such as the actual achievements on the field of the Vietnam War (Apple 1996).¹

The 1971 episode explains why the career of the former US Marine and military analyst and consultant at the RAND Corporation and the US Defence Department, who also served two years in Vietnam during the war, was put to an end. But this episode may also explain why Ellsberg did not pursue an academic career. A Harvard economics graduate, Ellsberg published two papers on utility and game theory in the mid-1950s. After serving in the Army for three years, he resumed graduate studies at Harvard in 1957, and he published a few more articles in economic journals, including the famous 1961 *Quarterly Journal of Economics* essay on ambiguity in decision-making. He completed his Ph.D. at Harvard in 1962, while already involved with RAND. But his decision to serve as military consultant made him unable to keep contributing to economics. As a result, the paradox of choice now known as the Ellsberg Paradox – a result that threw decision theory in deep disarray, and yet remained unaddressed for long, “simply because researchers at the time were helpless to address [it]” (Machina 2001, p. xxxix) – was presented in the early 1960s without the economic theorist behind it becoming involved in its defence and development. Even today, with hundreds of articles in major economic journals referring to the Ellsberg Paradox each year, and leading decision theorists still in the process of working out the nature and implications of ambiguity aversion, the overwhelming majority of economists think of Ellsberg only as a paradox and not as a fellow-economist.²

If we do look at Ellsberg as an economist, we discover another important unorthodox aspect. Even though his 1961 article “Risk, ambiguity and the Savage axioms” is well-known and

¹ Even very recently his support for Edward Snowden and his role in sparking the debate that forced US Congress to change surveillance law has been given attention in the international press (MacAskill et al. 2015).

² Almost all current studies on decision-making under uncertainty originate as attempts to account for Ellsberg’s ambiguity, particularly after the experimental evidence confirming his results became overwhelming (Camerer and Weber 1992). Even the recent interest in multiple-prior models and non-additive probability measures was initially brought about by the observation that decision-makers usually exhibit an aversion towards the ambiguous situations pointed out by Ellsberg (Wakker 2008; Gilboa 2009).

increasingly quoted, its philosophical background remains totally unknown.³ In addition, very few attempts to devise a decision rule that can solve the paradox refer to Ellsberg's own ideas about a possible solution. This is not very surprising: Ellsberg's 1961 *QJE* article is rather sketchy on both aspects. In his Ph.D. thesis, however, he dealt with both aspects in detail. Up until the recent printing of the thesis (Ellsberg 2001), conventional wisdom, even among historians of recent economics, seems to have been that the article was an excerpt of the thesis, while in fact the thesis was finished after the article had been published.⁴

The main aim of this paper is to put the Ellsberg Paradox in the right context. This paper examines Ellsberg's motivations in presenting it first to his fellow decision theorists at Harvard and RAND in the late 1950s and then to the economics profession in 1961. It also goes into his reasons for giving a methodological justification and defence in his doctoral thesis of 1962. By concentrating mainly on Ellsberg's deep and all-encompassing analysis of decision-making in his thesis, the paper shows that a number of relevant issues connected to the paradox can be thrown light on. More in particular, the paper argues that Ellsberg subscribed to a generalised version of the Bayesian approach, one that informs the developments of the multiple prior approach in current decision theory. The paper aims to show that even more than 50 years after the publication of the Ellsberg Paradox, an assessment of Ellsberg as decision theorist is necessary and important.

2. Ellsberg's early years

Ellsberg was born in Chicago in 1931 and attended Harvard University on a scholarship. He graduated in economics under the supervision of John Chipman in 1952, with a thesis titled "Theories of Rational Choice under Uncertainty: The Contribution of von Neumann and Morgenstern." He then spent a year as a visiting scholar at the University of Cambridge, UK, where he worked on a revision of part of his thesis, later published as Ellsberg (1954 and 1956). On his return to the US he entered the Marine Corps Basic School and served as a platoon leader, leaving the Corps in 1957 as a First Lieutenant.

Ellsberg started graduate school at Harvard University in 1957. As a Junior Fellow of the Society of Fellows at Harvard he started working on his critique of the Savage axioms. His graduate studies were interrupted when he joined the RAND Corporation as a strategic analyst in 1959. He also served as a consultant to the Defence Department, where he specialized in problems of the command and control of nuclear weapons, nuclear war plans, and crisis decision-making. He completed his Ph.D. in Economics at Harvard in 1962, with a thesis entitled *Risk, Ambiguity and Decision*, which partly drew on the 1961 essay "Risk, ambiguity and the Savage axioms." In 1964 Ellsberg joined the Defence Department as Special Assistant to Assistant Secretary of Defence (for International Security Affairs) John McNaughton, where he worked on the escalation of the war in Vietnam. He transferred to the State Department in 1965 and served for two years at the US Embassy in Saigon, where he was involved in the evaluation of pacification initiatives in the field. On his return to the RAND Corporation in 1967, Ellsberg worked on the top secret McNamara study about US decision-making in Vietnam (1945-68). In 1969, convinced that the McNamara study revealed that the US Administration had expanded their military operations without approval from Congress and had misled the public about the real intentions they pursued during the war,

³ A search for articles citing Ellsberg's 1961 paper in the single year 2014 using Google Scholar returns 515 quotations (search made on Sept 1, 2015).

⁴ As recalled by Ellsberg himself (2011, p. 225), "most of those who have become aware that the thesis existed, either before or after it was published, have had the mistaken belief that it was written before the 1961 *QJE* article, which, they suppose, represents a concise summary of the thesis."

Ellsberg made photocopies of the 7,000 page study and gave it to the Senate Foreign Relations Committee. After he had been unable to persuade the Senators he had contacted to make it public, he sent it to the *New York Times*, the *Washington Post* and 17 other newspapers. This was in 1971. After the publication of these “Pentagon Papers,” Ellsberg publicly turned himself in to the US Attorney’s Office for the District of Massachusetts in Boston, well aware that he was facing charges of conspiracy under the Espionage Act of 1917, which carried a total maximum sentence of 115 years. Luckily, during the hearings of the trial Ellsberg’s defence was able to prove gross governmental misconduct and illegal evidence gathering, and the judge officiating the case dismissed all charges in 1973. Ellsberg then became a political activist and writer.⁵

Ellsberg started publishing in major economic journals soon after his graduation from Harvard in 1952. These publications are extensively quoted in a variety of different fields. In his first article, an elaboration of the bachelor’s thesis, Ellsberg (1954) compares von Neumann and Morgenstern’s notion of utility to the classical notion of a cardinal utility function and clarifies the nature of the cardinality used in expected utility theory. The article, written while he still supported a mainstream view of decision-making, is praised in the literature for its clarity in explaining the different derivation and significance of the two notions of cardinal utility (Fishburn 1989; Giocoli 2003). In his second paper, which was also taken from his thesis, Ellsberg (1956) is critical of the notion of minimax strategy equilibrium in two-person zero-sum games. The motivations justifying the choice of a “reluctant duellist” not to follow the prescription of the theory, as she doubts her adversary really to play a maximin strategy, are part of “a small literature (beginning with Ellsberg 1956) and a much larger oral tradition which argues that Nash behaviour is neither a necessary consequence of rationality nor a reasonable empirical proposition” (Bernheim 1984, p. 1007). It is worth noting that Ellsberg presents his argument emphasising that these “particular uncertainties – as to the other players’ beliefs about oneself – are almost universal” (Ellsberg 1956, p. 43).

Ellsberg’s openness to game theoretical insights, combined with his refusal to apply equilibrium notions in an uncritical manner, is confirmed by his third 1950s essay, a RAND research paper on the theory and practice of blackmail (Ellsberg 1959) that is quoted extensively in political science literature (Zagare 2004). This is the most remarkable of a series of studies on conflict and bargaining, including an analysis of the “deterrence policy” a superpower can enforce avoiding to make “attack” a dominant strategy for the opponent, in the same vein of Schelling (1960) (Ellsberg 1961b). The motivation for distancing his analysis from a more conventional game-theoretical one appears to be the attention Ellsberg pays to the informational conundrum that decision-makers are plunged into when acting in a real strategic context – a state of affairs he had experienced personally as a consultant to the US government.⁶ As a matter of fact, the issue of how to represent the state of information of a decision-maker acting under uncertainty already became the major focus of his analysis in the late 1950s. In his office at RAND, by means of “endless trial and error with paper and pen,” Ellsberg (2011, p. 223) recalls, he was

“searching for choices between gambles – actions with uncertain outcomes – that would give operational meaning behaviorally 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.’”⁷

⁵ See Ellsberg (2003; 2006). On the controversial aspects of Ellsberg’s life see Samuels (2002).

⁶ The quality of his involvement in actual decision-making is exemplified, for instance, by taking part in two of the three working groups reporting to the Executive Committee of the National Security Council during the Cuban Missile Crisis in 1962 (Ellsberg 2006).

⁷ Ellsberg’s (1958) review of Davidson *et al* (1957) is the first publication in which the paradox is alluded to. Ellsberg (1958, pp. 1010-1011) prizes the methodology presented in the reviewed volume, one of the

3. Introducing the Ellsberg Paradox

In order to highlight the significance of Ellsberg's analysis of decision-making put forward in his 1962 doctoral thesis, this section provides a brief summary of how the paradox of choice was introduced in the original *QJE* article. Ellsberg's 1961 essay, first presented at the December Meeting of the Econometric Society in St. Louis in 1960, was published as part of a symposium on "Decision under Uncertainty" that also included a paper by William Fellner on subjective probability distortions and a critical comment by Howard Raiffa. Two further comments on the articles of the symposium were published in 1963, the first by Brewer (1963), who addressed mostly Fellner's paper, and the second by Roberts (1963), who reacted to Ellsberg's. Both Ellsberg (1963) and Fellner (1963) briefly replied to their critics.⁸ Brewer and Fellner (1965) summarised the debate. This also led to the first paper that set forth an experiment demonstrating the relevance of the violations Ellsberg had described (Becker and Brownson 1964). The outcome of this brief round of discussion, it should be reminded, was in line with Raiffa's main point: violations simply showed that people needed to be thought about the theory. As a matter of fact, abandoning 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 mainly faded away at least until Einhorn and Hogarth (1986) proposed a descriptive model of judgement under uncertainty accounting for ambiguity in the same vein of Kahneman and Tversky (1979), and the available experimental evidence became overwhelming (Camerer and Weber 1992).

Ellsberg opens his 1961 article with a reference to Knight's distinction between measurable and unmeasurable uncertainty. He observes that contemporary developments of subjective expected utility theory motivated a growing scepticism about the behavioural significance of this distinction.⁹ Ellsberg acknowledges that Leonard Savage's *Foundations of Statistics* provides the axiomatic structure to substantiate the point that, in the subjectivist interpretation of probability theory, risk and uncertainty become indistinguishable since actual decisions are interpreted as revealing degrees of belief, even when no statistical frequencies of events relevant to the decision are available. Indeed, Savage (1954) showed that one can deduce a unique subjective probability distribution over events with unknown probabilities from a decision-maker's choice behaviour, provided the behaviour satisfies certain axioms. But Ellsberg argues nonetheless that one cannot deny that there are probability relations about which decision-makers feel less "sure" as compared to others. Moreover, this holds true not only in complex situations like investment decisions or war gaming, but even when examining how decision-makers act in simple urn problems.

In a first example with two urns, Ellsberg argues that many decision-makers refrain from

first works intended to test experimentally the significance of the Bayesian approach under uncertainty, but he objects to the nature of the axioms the authors have chosen to test. The authors are criticized because "they tend to suggest that their basic hypothesis is uniquely plausible, covering reasonable behaviour on all occasions of uncertainty", while "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 behavior."

⁸ Ellsberg's reply is his single published comment on the debate opened by the 1961 article.

⁹ In his highly influential review article on alternative theories of decision-making in "risk-taking situations," Arrow (1951) had already promoted a new Bayesian consensus on the basis of Savage's early draft of his *Foundations of Statistics*. After assessing Knight's contribution from a subjectivist viewpoint, Arrow (1951, p. 417) commented: "Knight's uncertainties seem to have surprisingly many of the properties of ordinary probabilities, and it is not clear how much is gained by the distinction." See also Marschak (1950).

betting on the drawing of balls from an uncertain urn when a risky one is available.¹⁰ 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.¹¹ In both cases the choice behaviour of most decision-makers fails to conform to the prescriptions of subjective expected utility theory. Crucially, Ellsberg reports that 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. Ellsberg did not run experimental sessions and he presents his urn examples mainly as “hypothetical experiments.” He also illustrated his examples at faculty seminars he gave in the late 1950s: people tested during these seminars were mostly faculty members at the Universities of Chicago, Harvard, Northwestern and Yale, and at the RAND Corporation.

Aiming to avoid for his examples to be classified as a descriptive anomaly of an otherwise acceptable normative theory, Ellsberg (1961a, p. 646) concentrates 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.” This is an acknowledgement that, in principle, he accepts Savage’s reaction to Allais’s (1953) counter-example: decision theory has a normative status that cannot be criticised on the grounds of descriptive failures per se, since a decision-maker who instinctively violates the theory when confronting Allais’s decision set-up is supposed to reverse her choice after “thorough deliberation” (Savage 1954, p. 102). It is only “deliberate decisions” that are apt to reveal degrees of belief compatible with a probability measure, that is, that must show consistency with the axioms of the theory.¹²

Ellsberg concludes that, since the group of many “otherwise reasonable people” he tested who did not conform to the Savage axioms included a number of fellow decision theorists, his examples cannot be simply suggesting descriptive violations; they show that the normative content of the theory is not as compelling as generally agreed. 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

¹⁰ As is well-known, Ellsberg’s first example concerns the choice of betting on the drawn of a red (or black) ball from either an urn containing 100 red and black balls in an unknown proportion, or an urn containing 50 red and 50 black balls. Individuals tested by Ellsberg tended to be indifferent between betting on red or black when confronting each urn separately. But when asked to choose from which urn 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, “unambiguous” urn, with respect to the first, “ambiguous” one. These choices, Ellsberg observed, are contradictory with a definite, sharp probability assignment. Indeed, 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 $p(\text{red}_I)+p(\text{black}_I)<p(\text{red}_I\cup\text{black}_I)$. The individual who prefers to bet on urn *II* is said to show ambiguity aversion in current decision theory.

¹¹ Ellsberg’s second example yields a direct test of Savage’s “sure-thing principle,” Savage’s substitute for von Neumann and Morgenstern’s independence principle outside risky lotteries.

¹² This was the conventional reaction of the “American School” when challenged by the Allais Paradox in the early 1950s (Jallais and Pradier 2005). When he was presented Allais’s example 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, p. 103). See also Morgenstern (1979). It should be noted that at least until MacCrimmon (1968) and Slovic and Tversky (1974) the influence of Allais’s paradox on normative decision theory was almost nihil. The paradox was considered for long time only an isolated example suggesting a descriptive anomaly.

reasonable” do not conform to a numerical probability, even upon reflection. Savage is reported by Ellsberg (1961a, p. 656) to be one of them.¹³ The attack on the consolidating mainstream represented by subjective expected utility theory is explicitly put as follows (Ellsberg 1961a, p. 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”.

These choice-situations, he claims, must involve “uncertainties that are not risks.”

Ellsberg is well aware that the notion of uncertainty hinted at by Knight cannot be restricted to simplified cases like the one he discusses, and that there is no “true” uncertainty in his “unknown” urn. In his brief comments on the significance of his results beyond simple decisional contexts he hints at a “broader field of application” and suggests that, as compared with the effects of “familiar” production decisions or decisions concerning random processes (like coin flipping or urn extractions), “the results of Research and Development, or the performance of a new President, or the tactics of an unfamiliar opponent are all likely to appear ambiguous” (Ellsberg 1961a, p. 661). But in order to avoid being caught in a methodological discussion about the meaning of uncertainty, he claims to concentrate on discussing a “third dimension” of the problem of choice, additional to utility and probability proper, one which concerns the nature of the decision-maker’s information. What is at issue in his decision problem, Ellsberg (1961a, p. 657) claims, is what

“might be called the *ambiguity* of this information, a quality depending on the amount, type, reliability and ‘unanimity’ of information, and giving rise to one’s degree of ‘confidence’ in an estimate of relative likelihoods.”

Ambiguity is a qualitative notion as it may be high (and the confidence in any particular estimate of probabilities low) even where there is ample quantity of information, but “there are questions of reliability and relevance of information, and particularly where there is conflicting opinion and evidence” (Ellsberg 1961a, p. 659).¹⁴ To be sure, Ellsberg concludes, Savage himself alluded to a difficulty in his approach when arguing that “there seems to be some probability relations among which we feel relatively more ‘sure’ as compared to others,” and that “vagueness associated with judgements of the magnitude of personal probability” may constitute a difficulty for the theory of

¹³ Ellsberg’s (2001, p. xlix) recollection is that he showed his examples to Savage in February 1958, when Savage reacted “encouraging me to believe that the arguments presented here deserved serious consideration.” Ellsberg discussed the topic with Savage while writing the thesis, but there is no evidence that the comments concerning a draft of the thesis promised by Savage to Ellsberg in a letter dated February 8, 1962 – in the hope they arrive “soon enough to affect your thesis, insofar you find them pertinent” – have ever been written. The manuscript of Ellsberg’s thesis, sent to Savage a few days before its defence, is announced as “a 400-page letter to you, designed to change your mind” (Letter from D. Ellsberg to L. J. Savage, May 21, 1962, Leonard J. Savage Papers, MS 695, Box 11/260, Manuscript and Archives, Yale University). Ellsberg did not collect personal correspondence he had during his “economics” years (personal communication).

¹⁴ In the economics literature the term ambiguity has been used in different ways, either as second order uncertainty, or “uncertainty about uncertainties” (Einhorn and Hogarth 1986, p. 227), or “uncertainty about probability, created by missing information that is relevant and could be known” (Camerer and Weber 1992, p. 330). The widespread attitude in current decision theory is to use uncertainty as a generic term to describe all states of information about probabilities. While risk is used when the relevant probabilities are known, ambiguity refers to situations where some or all of the relevant information about probabilities is missing. Choices are said to be ambiguous if they are influenced by events whose probabilities are unknown or difficult to determine (Eichberger and Kelsey 2009).

personal probability (Savage 1954, p. 57, 59). But it is only Knight who stated that this is a crucial aspect of decision-making.¹⁵ Clearly, Ellsberg's contention about the degree of confidence in a probability assessment resembles Keynes's approach to uncertainty at least as much as Knight's. But neither the *General Theory* nor the *Treatise on Probability* are mentioned in his paper.¹⁶

A series of technical aspects that are discussed in the 1961 paper deserve mention as they constitute the subject matter of many later developments in decision theory. First, Ellsberg justifies the search of decision criteria alternative to the maximisation of expected utilities on the grounds that the probabilities that can be elicited from the choices of the decision-makers who prefer to bet on the "known" urn are "inconsistent with the essential properties of probability relationships" (Ellsberg 1961a, p. 651). There is no examination, though, of what the relaxing of some of these properties may imply in terms of an alternative representation. Second, the epistemic state of a decision-maker whose choices cannot be represented by a probability distribution is interpreted as that of a decision-maker who cannot commit herself to a single probability distribution out of a set of possible probability priors. Ellsberg proposes to interpret the beliefs of the violator of the Savage axioms as those of a decision-maker who, out of the set of all possible probability distributions, takes into account a subset of distributions "that still seem 'reasonable,' reflecting judgements that he 'might almost as well' have made, or that his information – perceived as scanty, unreliable, ambiguous – does not permit him confidently to rule out" (Ellsberg 1961a, p. 661). The subjective probability distribution presupposed by a Bayesian theorist is now considered simply a best estimate among the set of "reasonable" probability distributions. Third, a parameter representing the decision-maker's confidence in her best estimate is introduced in order to specify a decision rule alternative to the maximisation of subjective expected utility. Inspired by statistical literature (Hodges and Lehmann 1952), Ellsberg suggests a decision rule that is a mixture of "best guess" and "maximin": the 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 probability estimate, 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 (i.e., $\rho=1$) the uncertain context can be reduced to a risky one, so the rule coincides with Savage's. The case of "complete" ignorance (i.e., $\rho=0$) suggests concentrating on maximising the minimum possible outcome associated with the action as in Wald (1945). In between these two extremes the more relevant case of "considerable" ignorance is contemplated, and the existence of ambiguity (i.e., $0 < \rho < 1$) implies that the best guess distribution is now weighted with the maximin. Ellsberg concludes the presentation of his decision rule with a numerical example that accounts for the

¹⁵ There is no reference to Ellsberg in Savage's printed works apart from Savage's annotated bibliography prepared for the 2nd edition of the *Foundations of Statistics*, in which the papers presented at the 1961 *QJE* Symposium are said to present "an account of an important line of dissent from the theory of personal probability and utility" (Savage 1970, p. 25). Savage (1970, p. 24) also acknowledges that de Finetti and Savage (1962) "discuss a large number of doubts and difficulties about the actual application of personal probability," but while the notion of interval-valued probabilities is examined at length, there is no mention of Ellsberg in that paper. De Finetti's (1967) taxonomy of the varieties of subjective approaches to probability, which draws on de Finetti and Savage (1962), classifies Ellsberg's analysis among subjective "psychological" theories, devoid of normative content. On de Finetti's analysis of interval probabilities see Feduzi *et al.* (2014).

¹⁶ Knight's (1921, p. 233) reference to confidence is peripheral but unequivocal: "The business man himself not merely forms the best estimate he can of the outcome of his actions, but he is likely also to estimate the probability that his estimate is correct. The 'degree' of certainty or of confidence felt in the conclusion after it is reached cannot be ignored, for it is of the greatest practical significance. The action which follows upon an opinion depends as much upon the amount of confidence in that opinion as it does upon the favourableness of the opinion itself."

choices of an hypothetical violator of the Savage axioms.¹⁷ As we shall see in the following section, all three of these technical aspects are the subject matter of a much deeper investigation and development in his doctoral thesis.

Even though he presents the concluding remarks of the 1961 article as tentative, Ellsberg forcefully argues that he does not consider Savage's approach to be a normative guide in ambiguous contexts. He suggests that this must remain true for all deliberate violators of the theory. As a matter of fact, "once certain patterns of 'violating' behaviour [are] distinguished and described in terms of a specified decision rule" (Ellsberg 1961a, p. 669), like the one he suggests, economic theorists have to wonder what kind of normative status can be attributed to this decision rule. Contrary to what economics literature on decision-making inspired by his results has been assuming for over fifty years, Ellsberg was not silent on these issues and his doctoral thesis shows this.

4. Assessing the Ellsberg Paradox

The doctoral thesis submitted to the Economics Department of Harvard provides the philosophical background of Ellsberg's critique of Savage as well as a thorough discussion of the literature on decision-making. As reported in a brief introductory note dated April 1962, the thesis draws heavily upon the 1961 article but "concentrates upon materials and aspects of the analysis considered subsequently to its acceptance for publication" (Ellsberg 2001, p. xl): the overlap between the article and the thesis is said to account for around ten per cent of the new material.¹⁸ This section offers an assessment of a selection of issues discussed in the thesis, with the aim of showing that Ellsberg's ideas apply well beyond the paradox, both from a methodological and a technical viewpoint. More in particular, this section concentrates on the following issues: the probabilistic foundation of his theory, the validation of a normative theory, the notion of consistency in alternative probabilistic approaches, and a taxonomy of decision rules.

Before moving on, though, an important terminological aspect that distinguishes the thesis from the article is worth pointing out. In the thesis, the theorists who follow Ramsey, de Finetti and Savage in their subjective approach to probability, are no longer identified as Bayesians, as it was in the *QJE* article. They are now considered as part of a "neo-Bernoullian" school, namely, scholars who provided the axiomatic structure to justify the Bernoulli Proposition that a rational decision-maker should maximise expected utility in making her decision ("moral expectation," in the English translation of Daniel Bernoulli's (1954) latin term "emolumentum medium"). This is to stress that there are also a number of authors who consider themselves to be Bayesians without subscribing to a sweeping application of Bernoulli's principle of the maximisation of expected utility. The term "neo-Bernoullian" is used, then, to denote the "dominant wing ... of the 'neo-Bayesian' school, which insists upon inferring probability from betting behaviour alone and demands that 'reasonable' behaviour be consistent in every case with definite, uniquely-defined numerical probabilities." This wing must be distinguished from "less exacting 'neo-Bayesians' such as I. J. Good and myself"

¹⁷ Inspired by Ellsberg, Eichberger and Kelsey (1999) provided an axiomatization of this decision rule in the context of non-additive probability theory.

¹⁸ Ellsberg referred to the thesis in his 1963 reply to Roberts, hinting at the fact that it presents a more elaborate justification of his critique of Savage on many issues. But the thesis remained unpublished and was almost completely ignored until its recent publication. The thesis is now reproduced in its original, unedited version, with an introduction by Isaac Levi and a bibliographical note edited by Mark Machina, in Ellsberg (2001). In his review of the recent printing of the thesis, Mukerji (2003, p. F188) observes that "it is simply breathtaking to note how many, many, years ahead of his time this man truly was."

(Ellsberg 2001, p. 165 fn.).¹⁹ Moreover, in view of its rejection of the assumption that, in decision under uncertainty, there always exists a complete ordering of probability, Keynes's epistemic approach to probability is put at the beginning of this tradition of thought.

4.1 Probabilistic foundations

As in the 1961 article, Ellsberg opens his doctoral investigation with a reference to Knight. However, in the doctoral thesis the emphasis is more on Keynes and the authors who followed him in stressing vagueness and imprecision of subjective assessments, none of whom was quoted in the article. Ellsberg (2001, pp. 9, 11) attributes to Keynes's *Treatise on Probability* the merit of having introduced "formally the notion of *non-comparability* of beliefs," and recognises that the notion of weight of argument "seems closely related to our notion of 'ambiguity'." He also recalls that Keynes emphasised 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 (1952) in their investigations of an axiomatic structure compatible with an interval of probability priors. In particular, Ellsberg distinguishes Koopman's and Good's theories of "consistent judgments" from Ramsey's, de Finetti's and Savage's theories of "consistent choices," praising the fact that theories of intuitive probabilities constitute "a theoretical approach that admits vagueness as an explicit factor without apology and provides a formal vocabulary for discussing it" (p. 10). Another crucial reference is to Cedric A. Smith's (1961) paper on consistency in statistical inferences, on which more later.

Furthermore, Ellsberg's discussion of ambiguity as a consequence of "highly conflicting" information starts with Keynes. Ellsberg (2001, p. 11) notes that in order to exemplify the possibility of partially ordered probabilities, Keynes referred to conflicting evidence. In a vivid attempt to exemplify the effects of conflicting evidence Keynes argued that, before going out for a walk, it is an "arbitrary matter to decide for or against the umbrella ... if the barometer is high, but the clouds are black" (Keynes 1921, p. 32). Ellsberg associates this example with his own intuition about the situation of an "intelligence panel of experts advising the President on Soviet military inventory, each member giving a precise, definite probability distribution, each distribution different," a situation that he himself, as a military analyst, experienced.

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 decision-making from a probabilistic viewpoint. It also serves to demonstrate that this was a variegated lot, including economists like Keynes, Knight, and Shackle, mathematicians like Koopman and Wald, and statisticians like Good, Hodges, Lehmann and Smith. Ellsberg (2001, p. 12) argues that these authors share the same distrust about the significance of assuming that subjective beliefs can be always represented by a single and fully reliable additive probability function.²⁰

¹⁹ Irving J. Good, was a British mathematical statistician known for his attempt to reject precise probabilities in an otherwise typically Bayesian framework (Good 1952 and 1962). The term "neo-Bernoullian" was possibly borrowed by Good himself (Ellsberg 2011, p. 221), but the insistence on the Bernoulli Proposition as the pillar of mainstream decision theory is due to Allais (1953).

²⁰ George Shackle's (1949) non-probabilistic theory of decision, cited by Arrow (1951) as the sole formalised alternative to probabilistic decision-making, is also a major reference in the thesis. Ellsberg (1961a, p. 643) had criticized Shackle's "extreme form of the Knightian position" for his rejection of a representation of uncertainty via numerical probabilities in any situation of crucial importance to the decision-maker, including those clearly conditioned on an aleatory mechanism. But in the thesis he now

The strength of the claim the doctoral thesis puts forward – namely, that a strict Bayesian approach such as the one endorsed by neo-Bernoullians is unable to distinguish between ambiguous and unambiguous decision contexts – suggests that Ellsberg’s aim was not limited to the testing of Savage’s theory. In his investigation Ellsberg shares with a number of scholars the cognitive unease of the theorist who accepts a theoretical approach in principle, but finds it restrictive in its application, and is inclined to question its normative value as a result.²¹

4.2 The testing of a normative theory

From the very beginning of his thesis Ellsberg professes his great admiration for Savage’s work. His praise does not hinge solely on Savage’s ability to devise a comprehensive axiomatic structure for a theory of decision-making. He values Savage’s perspective 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. Savage’s contribution made it possible to abandon psychological introspection in decision-making under uncertainty, in analogy with Samuelson’s revealed-preference approach in the theory of choice. Moreover, 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).

Therefore, Ellsberg’s respect for Savage is also motivated by methodological reasons. Ellsberg stresses that, in Savage’s perspective normative decision theory is not devoid of empirical content since it suggests the kind of observed behaviour that is compatible with the theoretical proposition. Specifically, the procedure of eliciting probabilities from choices makes it possible to rule out certain patterns of choice as not conforming to the theory. The test of the acceptability of a normative theory lies simply in telling an individual making decisions: “when certain maxims are presented for your consideration, you must ask ... how would you react if you noticed yourself violating them” (Savage 1954, p. 7, quoted in Ellsberg 2001, p. 24). As in the 1961 article, then, the test does not concern “unreflective behaviour,” but only situations in which the decision-maker is allowed to make, or change, her decision after thorough consideration. The empirical investigation that Savage’s approach makes possible, creates the possibility of questioning a theory when the decision-maker wishes to persist in her violation of the “maxims” of behaviour.²² So the identification of a clear-cut counterexample is instrumental to a discussion of the normative value of the theory.

admits that “when ambiguity is extreme, by any of his indices: relevant information sparse, or obviously unreliable and contradictory; wide differences in the expressed expectations of different individuals; low confidence in available estimates,” Shackle’s “somber reflections” seem “too ominously relevant to the very circumstances upon which this study focuses to be dismissed” (Ellsberg 2001, pp. 16-17). On the links between Ellsberg’s ambiguity and Shackle’s uncertainty see Basili and Zappia (2010) and Zappia (2014).

²¹ Regrettably, Ellsberg does not acknowledge that Keynes’s rejection of the criterion of maximising “moral expectations” in the *Treatise on Probability* precludes to an insightful analysis of the “making of practical decisions” (Basili and Zappia 2009). Also, apparently Ellsberg was not aware that Keynes had made a consistent use of the two-urn example to justify the influence of confidence in decision-making, on which see Feduzi (2010).

²² Ellsberg notes that Savage suggested to him in conversation that the procedure he was following corresponded to that pursued by Bernoulli (1954) when suggesting his maxim of behaviour. When Bernoulli questioned the use of expected value in decision theory he made it through an empirical observation, known as the St. Petersburg Paradox. The rejection of the expected value was made on the grounds that no gambler would regret not to bet a high amount of money in St. Petersburg game, even when shown that the expected value of the game has an infinitely high value.

Ellsberg's point (2001, p. 27) is 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 – as first argued by Marschak (1951) in his presentation of von Neumann and Morgenstern's utility axioms – the degree of acceptance and authority of certain maxims of behaviour must have a descriptive component nonetheless. A "system of logic" which satisfies logical consistency but does not apply to the deliberate decisions of a decision-maker may induce the decision-maker to consider, "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 expected utility theory (Ellsberg 2001, p. 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. The extension of "*familiar rules of logic and arithmetic* [to decision under uncertainty] *is not a normative but a descriptive problem.*"

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 falsificationist approach. But his aim is not confined to reject Savage's theory: he intends 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 2001, p. 25). It is apparent that Ellsberg's investigation does not stop after finding psychological element justifying the violations of the current theory.²³

It is worth noting that in his thesis Ellsberg also insists on true empirical work. Since the theories he wants to criticize show the possibility of measuring subjective variables, important factors that justify the choices of deliberate violators that they do not account for should also be measurable. In order to achieve this, empirical research in the laboratory "should strongly complement and interact with the continuing analysis of normative principle" (Ellsberg 2001, p. 34). Ellsberg approvingly quotes Becker and Brownson's preliminary experimental study (published later as Becker and Brownson 1964), that was based on his 1961 article and aimed at the measurement in monetary units of the desire by experimental subjects to avoid drawing from ambiguous urns. Ellsberg's (2001, p. 35) brief, but unequivocal assessment of empirical analysis concludes by foreboding the importance that laboratory experiments would assume in the following years: "empirical findings, as, hopefully, they begin to appear in more abundance, will undoubtedly change the views of every party to the current discussion in important ways now unforeseeable."

4.3 Consistency in decision-making

Ellsberg recalls that de Finetti (1937) solved the problem of how to obtain 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. De Finetti's contribution was to identify a system of betting prices apt to elicit "definite probabilities" from choices, and then to show that the requirement that these probabilities are "coherent" implies that they satisfy the axioms of probability.²⁴ Although he objects to the fact that this can be done with respect to "*every*

²³ Methodologically, there is a distinction to make between falsification and rejection. A counterexample may provide a necessary condition for the falsification of a normative theory. But this is not sufficient 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 (Lakatos 1970, Guala 2000). Ellsberg's thesis aims to provide heuristic advice about how to develop a new normative theory.

²⁴ In his *Foundations* Savage (1954, pp. 27-29) was not as explicit as de Finetti on the analogy with betting

conceivable event,” Ellsberg (2001, p. 55) claims that, in principle, he accepts de Finetti’s approach and the Bayesian methodology of eliciting probabilities from choices. Ellsberg’s acceptance, though, is coupled with a rejection of de Finetti’s “condition of coherence.”

Ellsberg argues that there is no rationale for the assumption that the decision-maker is always willing to take each side of a bet. For purposes of measurement, de Finetti regarded probabilities as “prices” and assumed that the highest price an individual is ready to pay for betting in favour of an event coincides with the lowest she is ready to accept for taking the other side of the bet. No inequality between upper and lower betting prices is contemplated by de Finetti. But “the explicit assumption ruling this inequality out as a possibility for reasonable behaviour,” Ellsberg (2001, p. 68) contends, “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.”

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), an approach that distinguishes the possibility to elicit “definite” probabilities from the fact that these probabilities must show the kind of consistency de Finetti assumed. Ellsberg notices 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. Indeed, consistency does not apply only to a set of “precise, definite beliefs,” since “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 2001, p. 88). Arguments based on de Finetti’s notion of consistency “contain, implicitly, the special, additional, assumption” that the decision-maker who is willing to pay a price p to bet on an event is also willing to pay $q = 1-p$ to bet on its complement. But Smith demonstrated that if the requirement of consistency of the betting prices is to avoid that a book is made against the decision-maker when $q \neq 1-p$, this implies simply that $p + q \leq 1$. In order for a book to be avoided, $p+q = 1$ is only a special case. Ellsberg (2001, p. 85) approvingly quotes Cox and Good, who, in their discussion on Smith’s paper, remark that the “reluctance to gamble” that is allowed by Smith when he identifies a non-betting interval “seems intuitively to correspond to recognizable behaviour-patterns associated with circumstances in which opinions are exceptionally indefinite.”²⁵

To sum up. An argument based on coherence has great strength since introduces an “objective” criterion of “reasonable” behaviour, by excluding that it is possible “to make a fool” of the decision-maker. But, Ellsberg (2001, p. 88) concludes, 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.”

behaviour, but later he endorsed this procedure for explanatory purposes (Savage 1962).

²⁵ It is worth noting that 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, and object to it only on the grounds that in order to provide a “zone of indecision” Smith’s theory defines an interval of probabilities with sharp edges that are at least as questionable as the sharpness of the point probabilities it criticises. A similar point is made by Savage (1962) and quoted by Ellsberg (2001, p. 73). Ellsberg objects to this view on the grounds that his examples imply a “zone of indefiniteness” that is distinguishable from any “zone of indecision,” since he does not consider that vagueness as necessarily implying indecision.

4.4 Decision rules for a generalized Bayesian approach

As we have seen, Ellsberg's aim in the thesis is to provide an alternative theory for Savage's "Bernoullian," strictly-Bayesian approach. His alternative theory is thought to be compatible with a generalised Bayesian viewpoint and endowed with a clear normative content. The 1961 paper is suggestive of such a theory when rejects the maximization of expected utility in ambiguous contexts, and introduces a criterion for choice depending on the degree of confidence in the probability assessment. In the thesis, Ellsberg tries to accomplish the more relevant task of justifying the proposed decision rule on the basis of a theoretical approach that is as general as to conceive an analysis of situation in which the decision-maker "does not know" the probability distribution of the pay-off relevant events, without conflating uncertainty into risk.

This aim is accomplished in two steps. First, Ellsberg provides a Bayesian justification for the choices of the unrepentant violator of Savage postulates to be included in the rational realm. Second, he derives a decision rule. In the probabilistic framework he wants to introduce, we find, on the one hand, the Bernoullian, strictly-Bayesian decision-maker, who may keep in mind a whole set of "reasonably acceptable" probability distributions before acting, but 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, especially when relevant information is perceived as scanty, unreliable, contradictory, ambiguous. As seen above, Ellsberg claims that the formal structure that guarantees the consistency of such probabilistic beliefs is provided by Koopman, Good and Smith.²⁶

The second step is a thorough analysis of what decision rule best suits this generalised probabilistic viewpoint. As mentioned earlier, Ellsberg concluded his 1961 article with an analysis of how to act "reasonably" when more than one personal probability distribution is acceptable and proposed the use of an index including the degree of confidence in a probability assessment. The 1962 thesis features a comparison between alternative criteria for explaining the observed departures from Savage's axioms, including decision rules not discussed in the article. Ellsberg observes that a number of alternatives to Savage's maximisation of subjective expected utility have been devised to deal with complete ignorance and that an extensive examination of the axiomatic structure underlying these decision rules is available (Milnor 1954, Chernoff 1954, Luce and Raiffa 1957). Among other decision rules, Ellsberg concentrate on what he calls Hurwicz α -maximin and proposes to follow Hurwicz's suggestion (1951b) to apply it to a context of partial ignorance.

Hurwicz's decision rule – referred to as Hurwicz α -criterion by Luce and Raiffa (1957) and better known as the Arrow-Hurwicz criterion after Arrow and Hurwicz (1972) – considers the behaviour associated with Wald's 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 α – 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. In his attempt to find a decision rule that can account for the behaviour of unrepentant violators of Savage's axioms, Ellsberg (2001, pp. 182-183) examines Hurwicz α -maximin in detail, and discusses which values of α make the behaviour of a decision-maker compatible with his paradox. So the idea of a set of distributions over the states of the world is maintained, but now it is the Hurwicz criterion that is applied to the restricted set of plausible distributions.

²⁶ This is in the main the road taken by authors unpersuaded by the strict Bayesian perspective codified by Savage. Among these see Levi (1974), Gardenförs and Sahlin (1982), Gilboa and Schmeidler (1989) and Walley (1991). See also Levi (2001).

It is worth noting that Ellsberg (2001, p. 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.*”²⁷ 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 rational underlying decision-making under uncertainty.

Ellsberg (2001, p. 195) concludes his investigation with a taxonomy of criteria in terms of confidence and optimism. A synthesis between the 1961 and the 1962 suggestions is also offered as a proposal for a new normative decision rule. A criterion combining Bayesian and minimax principles, that can generate all the choice patterns observed in urn experiments, called “restricted Bayes-Hurwicz criterion,” is proposed as eminently general. In order to account for both confidence and optimism, the normative proposal is to prescribe that the decision-maker chose the action x that maximises the index $\rho E(x) + (1-\rho)[\alpha \max(x) + (1-\alpha) \min(x)]$. But the new normative theory, does not suggest how to fix the value of ρ and α in given circumstances for a given decision-maker, in much the same way the neo-Bernoullians would not dictate the shape of the utility function.

5. Concluding remarks

In a recent paper on the way mainstream decision theory understands rationality when dealing with uncertainty, Gilboa *et al.* (2008, pp. 286-87) 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 counterexample to Bayesian decision-making that got the normative value of the theory into trouble. But he could not contribute to the foundations of a new normative theory. An assessment of his contribution in the more encompassing framework of his doctoral thesis suggests that, following in a tradition of thinkers that were critical of the maximisation of Bernoulli’s “moral expectation” and supported a generalised version of the Bayesian approach, Ellsberg provided a thorough philosophical, methodological and theoretical background for his critique.

This paper has simply tried to argue that, contrary to the conventional wisdom emerging from the current literature in decision-making, Ellsberg is not simply a paradox. Had Daniel Ellsberg’s general analysis been available when economists discussed probability and decision theory, it might well have renewed interest in theories of decision under uncertainty much earlier than it has been the case.

²⁷ Ellsberg (2001, p. 206) explains his interest in Hurwicz’s proposal also by arguing that “a normative criterion acceptable for my own decision-making [in highly uncertain contexts] would certainly take account of favourable possibilities in addition to ‘worst’ possibilities.” It is inevitable to note that certain recent developments in decision theory elaborate on the intuition behind the Hurwicz criterion exactly in the direction suggested by Ellsberg. See the class of so-called α MEU criteria, where the optimism parameter is introduced in the multiple priors context (Girardato *et al.* 2004, Chateauneuf *et al.* 2007).

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.
- Apple, R. W. Jr. (1996), "25 years later: Lessons from the Pentagon Papers," *New York Times*, June 23rd.
- Arrow, K. J. (1951), "Alternative approaches to the theory of choice in risk-taking situations," *Econometrica*, 19: 404-37.
- Arrow, K. J. (1953), "Hurwicz's optimality criterion for decision making under ignorance," *Technical Report 6*, Stanford University.
- Arrow, K. and Hurwicz, L. (1972), "Decision making under ignorance," in C. F. Carter and J. L. Ford (eds.), *Uncertainty and Expectations in Economics. Essays in Honour of G. L. S. Shackle*. Oxford: Basil Blackwell.
- Basili, M. and Zappia, C. (2009), "Keynes's 'non-numerical' probabilities and non-additive measures," *Journal of Economic Psychology*, 30: 419-30.
- Basili, M. and Zappia, C. (2010), "Ambiguity and uncertainty in Ellsberg and Shackle," *Cambridge Journal of Economics*, 34: 449-74.
- Becker, S. and Brownson, F. (1964), "What price ambiguity? Or the role of ambiguity in decision making," *Journal of Political Economy*, 72: 63-73.
- Bernheim, B. D. (1984), "Rationalizable strategic behaviour," *Econometrica*, 52: 1007-1028.
- Bernoulli, D. (1954, 1738), "Exposition of a new theory on the measurement of risk," *Econometrica*, 22: 23-36.
- Binmore, K. (2009), *Rational Decisions*. Princeton: Princeton University Press.
- Brewer, K. (1963), "Decision under uncertainty: comment," *Quarterly Journal of Economics*, 77: 159-61.
- Brewer, K. and Fellner, W. (1965), "The slanting of subjective probabilities: agreements on some essentials," *Quarterly Journal of Economics*, 79: 657-63.
- Camerer, C. and Weber, M. (1992), "Recent developments in modelling preferences: uncertainty and ambiguity," *Journal of Risk and Uncertainty*, 5: 325-70.
- Chateauneuf, A., Eichberger, J. and Grant, S. (2007), "Choice under Uncertainty with the Best and Worst in Mind: Neo-additive Capacities," *Journal of Economic Theory*, 137: 538-67.
- Chernoff, H. (1954), "Rational selection of decision functions," *Econometrica*, 22: 422-43.
- Davidson, D., Suppes, P. and Siegel, S. (1957), *Decision Making*. Stanford: Stanford University Press.
- de Finetti, B. (1937), "Foresight: its logical laws, its subjective sources," in H. E. Kyburg and H. E. Smokler (eds.), *Studies in Subjective Probability*. New York: Wiley, 1964.
- de Finetti, B. (1967), "Probability: Interpretations," in *International Encyclopedia of Social Sciences*. New York, Macmillan.
- de Finetti, B. and Savage, L. J. (1962), "Sul modo di scegliere le probabilità iniziali," *Biblioteca del Metron, Serie C*, 1: 81-147.
- Eichberger, J. and Kelsey, D. (1999), "E-capacities and the Ellsberg Paradox," *Theory and Decision*, 46:107-138.
- 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.
- Einhorn, H. J. and Hogarth R. M. (1986), "Decision making under ambiguity," *Journal of Business*, 59: 225-50.
- Ellsberg, D. (1954), "Classic and current notions of 'measurable utility'," *Economic Journal*, 64: 528-56.
- Ellsberg, D. (1956), "Theory of the reluctant duelist," *American Economic Review*, 46: 909-23.
- Ellsberg, D. (1958), "Review of Decision Making by D. Davidson, P. Suppes and S. Siegel"

- American Economic Review*, 48: 1009-1011.
- Ellsberg, D. (1959), "The theory and practice of blackmail," mimeo, RAND Corporation, Santa Monica.
- Ellsberg, D. (1961a), "Risk, ambiguity, and the Savage axioms," *Quarterly Journal of Economics*, 75: 643-69.
- Ellsberg, D. (1961b), "The crude analysis of strategy choices," *American Economic Review*, 51: 472-78.
- Ellsberg, D. (1963), "Risk, ambiguity, and the Savage axioms: reply," *Quarterly Journal of Economics*, 77: 336-42.
- Ellsberg, D. (2001), *Risk, Ambiguity and Decision*. New York: Routledge.
- Ellsberg, D. (2003), *Secrets: A Memoir of Vietnam and the Pentagon Papers*. Penguin Books: New York.
- Ellsberg, D. (2006), "Daniel Ellsberg: Extended Biography (written for the Right Livelihood Award process)," available at www.ellsberg.net/bio/extended-biography (consulted on Sept. 1, 2015).
- Ellsberg, D. (2011), "Introduction to the symposium issue: notes on the origins of the Ellsberg urns," *Economic Theory*, 48: 221-27.
- Feduzi, A., Runde, J. and Zappia, C. (2014), "De Finetti on uncertainty," *Cambridge Journal of Economics*, 38: 1-21.
- Fellner, W. (1961), "Distortion of subjective probabilities as a reaction to uncertainty," *Quarterly Journal of Economics*, 75: 670-689.
- Fishburn, P. C. (1989), "Retrospective on the utility theory of von Neumann and Morgenstern," *Journal of Risk and Uncertainty*, 2: 127-58.
- 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. and Schmeidler, D. (1989), "Maxmin expected utility with non-unique prior," *Journal of Mathematical Economics*, 18: 141-153.
- Gilboa, I., Postlewaite, A. W. and Schmeidler, D. (2008), "Probability and uncertainty in economic modelling," *Journal of Economic Perspectives*, 22: 173-188.
- Giocoli, N. (2003), *Modeling Rational Agents: From Interwar Economics to Early Modern Game Theory*. Aldershot, Edward Elgar.
- Good, I. J. (1952), "Rational decisions," *Journal of the Royal Statistical Society, Series B*, 14: 107-114.
- Good, I. J. (1962), "Subjective probability as the measure of a nonmeasurable set," in Nagel, E., Suppes, P. and Tarsky, A. (eds.) *Logic, Methodology and Philosophy of Science*. Stanford: Stanford University Press.
- Hands, D. W. (2012), "The positive-normative dichotomy and economics," in Maki, U. (ed.), *Philosophy of Economics*. Amsterdam: Elsevier.
- Hodges, J. L. and Lehmann, E. L. (1952), "The uses of previous experience in reaching statistical decision," *Annals of Mathematical Statistics*, 23: 396-407.
- Hurwicz, L. (1951a), "Optimality criteria for decision making under ignorance," *Cowles Commission Discussion Paper No. 370*, Chicago.
- Hurwicz, L. (1951b), "Some specification problems and application to econometric models" (abstract), *Econometrica*, 19: 343-344.
- Jallais, S. and Pradier, P-C. (2005), "Changes in Expected Utility Theory and the Allais Experiment," in Fontaine, P. and Leonard, R. (eds.), *The Experiment in the History of Economics*. London: Routledge.

- Kahnemann, D. and Tversky, A. (1979), "Prospect theory: An analysis of decision under risk," *Econometrica*, 47: 263-291.
- Keynes, J. M. (1921), *A Treatise on Probability. The Collected Writings of John Maynard Keynes*, Vol. VII. London: Macmillan, 1973.
- Kyburg, H. E. Jr. (1998), "Interval-valued probabilities," in G. de Cooman, P. Walley, and F. G. Cozman (eds.), *Imprecise Probabilities Project*. Available from <http://ippserv.rug.ac.be/>.
- Knight, F. H. (1921), *Risk, Uncertainty and Profit*. Chicago: University of Chicago Press, 1985.
- Koopman, B. O. (1940), "The axioms and algebra of intuitive probability," *Annals of Mathematics*, 41: 269-92.
- Lakatos, I. (1970), "Falsification and the methodology of scientific research programmes," in Lakatos, I. and Musgrave, A. (eds.), *Criticism and the Growth of Knowledge*. Cambridge: Cambridge University Press.
- Levi, I. (1974), "On indeterminate probabilities," *Journal of Philosophy*, 71: 391-418.
- Levi, I. (2001), "Introduction," in Ellsberg 2001.
- Luce, D. and Raiffa, H. (1957), *Games and Decisions*. New York: Wiley.
- MacAskill, E., Roberts D., and Jacobs, B. (2015), "Daniel Ellsberg credits Edward Snowden with catalysing US surveillance reform," *The Guardian*, June 1st.
- Machina, M. (2001), "Bibliographical note," in Ellsberg 2001.
- 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 Neyman, J. (ed.), *Proceedings of the Second Berkeley Symposium on Mathematical Statistics and Probability*. Berkeley: University of California Press.
- Morgenstern, O. (1979), "Some reflections on utility," in Allais, M. and Hagen, O. (eds.), *Expected Utility Hypothesis and the Allais Paradox*. Dordrecht: Reidel.
- Mukerji, S. (2003), "Review of Risk, Ambiguity and Decision by D. Ellsberg," *Economic Journal*, 113: 187-188.
- 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. (1931), "Truth and probability," in *Foundations: Essays in Philosophy, Logic, Mathematics and Economics*, edited by D. H. Mellor. London: Routledge & Kegan Paul, 1978.
- Roberts, H. (1963), "Risk, ambiguity, and the Savage axioms: a comment," *Quarterly Journal of Economics*, 77: 327-36.
- Samuels, W. J. (2002), "Daniel Ellsberg: analysis and redemption of an economist," *History of Economic Ideas*, 10: 123-54.
- Samuelson, P. A. (1950), "Probability and the attempts to measure utility," *Economic Review*, 1, pp. 167-173. Reprinted with a 1965 postscript as Chapter 12 in J. E. Stiglitz (ed.), *Collected Scientific Papers of Paul A. Samuelson*. Cambridge, Mass: MIT Press, 1966.
- Savage, L. J. (1954), *The Foundation of Statistics*. New York: Wiley, 1972, 2nd ed.
- Savage, L. J. (1962), "Bayesian statistics," in Machol, R. E. and Gray, P. (eds.), *Recent Developments in Decision and Information Processes*. MacMillan: New York.
- Savage, L. J. (1970), "Reading suggestions for the Foundation of Statistics," *The American Statistician*, 24: 23-27.
- Schmeidler, D. (1989), "Subjective probability and expected utility without additivity," *Econometrica*, 57: 571-87.
- Shackle, G. L. S. (1949), *Expectation in Economics*. Cambridge: Cambridge University Press. 1952, 2nd ed.

- Shelling, T. (1960), *The Strategy of Conflict*. Cambridge, MA: Harvard University Press.
- Smith, C. A. B. (1961), "Consistency in statistical inference," *Journal of the Royal Statistical Society*, 23: 1-25.
- von Neumann, J. and Morgenstern O. (1944), *Theory of Games and Economic Behaviour*. Princeton: Princeton University Press.
- Wakker, P. (2008), "Uncertainty," in Blume, L. and Durlauf, S. N. (eds.), *The New Palgrave: A Dictionary of Economics*, 2nd. Edition. London: Palgrave Macmillan.
- Wald, A. (1945), "Statistical decision functions which minimize the maximum risk," *Annals of Mathematics*, 46: 265-280.
- Walley, P. (1991), *Statistical Reasoning with Imprecise Probabilities*. London: Chapman and Hall.
- Zagare, F. C. (2004), "Reconciling Rationality with Deterrence. A Re-Examination of the Logical Foundations of Deterrence Theory," *Journal of Theoretical Politics*, 16: 107-141.
- Zappia, C. (2014), "Non-Bayesian decision theory ahead of its time: the case of G.L.S. Shackle," *Cambridge Journal of Economics*, 38: 1133-54.