Decisions in Economics and Finance
https://doi.org/10.1007/s10203-024-00437-1

Ellsberg 1961: text, context, influence

Ivan Moscati[1,2,3]

Received: 12 September 2023 / Accepted: 13 February 2024

Abstract
In 1961 Daniel Ellsberg published an article titled "Risk, Ambiguity, and the Savage Axioms" in the *Quarterly Journal of Economics*, which became a seminal contribution to the theory of decision-making under uncertainty. This paper analyzes Ellsberg's 1961 classic, situates it within the context of decision-making theory in the 1950s and early 1960s and within the development of Ellsberg's ideas, and provides an overview of the experimental and theoretical literature to which it gave rise.

Keywords Ellsberg · Decision theory · Uncertainty · Ambiguity · Ellsberg paradox

JEL Classification B21 · B31 · D81 · C90

1 Introduction

Daniel Ellsberg (1931–2023) was an American decision theorist, Marine, military analyst, and, after disclosing the Pentagon Papers in 1971, a political activist. His research in decision theory focused on the theory of individual decision-making under uncertainty, and in 1961 he published an article titled "Risk, Ambiguity, and the Savage Axioms" in the *Quarterly Journal of Economics* (*QJE*), which became a seminal contribution to the field.

In that article, Ellsberg envisaged a choice situation, later referred to as the Ellsberg paradox, in which a decision maker has to express her preferences between gambles with uncertain outcomes. The gambles yielded either \$0 or \$100 depending on the color of a ball drawn from an urn containing balls of different colors. Ellsberg observed that several reasonable decision makers expressed deliberate preferences that violated the axioms of the then-dominant theory of decision-making under uncertainty, the

Ivan Moscati
ivan.moscati@uninsubria.it

1 Department of Economics, University of Insubria, Varese, Italy

2 Baffi Carefin Center, Bocconi University, Milan, Italy

3 Centre for Philosophy of Natural and Social Science, LSE, London, UK

Published online: 20 May 2024

version of expected utility theory advanced by Leonard J. Savage ([1954] 1972). For Ellsberg, these preferences show that a class of situations of uncertainty, which he called situations of "ambiguity," exists wherein Savage's theory loses its descriptive and normative validity. In the final part of his article, Ellsberg advanced an alternative model capable of accounting for choices in conditions of ambiguity.

This paper analyzes Ellsberg's *QJE* classic, situates it within the context of the theory of decision-making of the 1950s and early 1960s and within the development of Ellsberg's ideas, and provides a bird's-eye view of the experimental and theoretical literature that stemmed from Ellsberg's work.

The paper caters to both readers who are unfamiliar with Ellsberg's classic and those already acquainted with it. The former group will likely gain an appreciation of the sophisticated nature of Ellsberg's arguments and may be encouraged to read the original article. For those already familiar, this paper unveils aspects of Ellsberg's contribution that become evident only when the origin, context, and evolution of his ideas are elaborated.

The remainder of this article is structured as follows. Section 2 outlines Ellsberg's multifaceted biography. Section 3 reviews the dominant theories of decision-making under risk and uncertainty between 1945 and 1960, and so provides the backdrop for Ellsberg's research. Section 4 traces the development of Ellsberg's ideas, from the two articles derived from his Harvard BA thesis to the initial version of the *QJE* article. Sections 5 and 6 delve deeply into the *QJE* article. Section 7 examines Ellsberg's PhD dissertation, which largely drew on the *QJE* article but diverged from it in some important respects. Section 8 considers Savage's stance on Ellsberg's work. Section 9 overviews the experimental and theoretical literature originated by Ellsberg's article. Section 10 concludes.

2 The many Ellsbergs

Daniel Ellsberg was born in 1931 in Chicago.[1] He attended schools in Bloomfield Hills, Michigan, and began his studies at Harvard in 1948. He graduated in 1952 with a BA degree in economics, earning summa cum laude honors. His thesis, titled "Theories of Decision-making Under Uncertainty: The Contributions of von Neumann and Morgenstern," was supervised by John Chipman, then an assistant professor at Harvard. In the academic year 1952–1953, Ellsberg studied for one year at King's College, Cambridge University on a Woodrow Wilson Fellowship. During his stay at Cambridge, he developed his BA thesis into papers later published in the *Economic Journal* (Ellsberg 1954) and the *American Economic Review* (Ellsberg 1956). Even these early articles display some of the elements that characterize Ellsberg's attitude toward decision theory and that are also present in his *QJE* article: a critical stance on existing theories for decision-making under risk and uncertainty and a focus on the normative dimension of decision analysis.

[1] More details about Ellsberg's biography in Ellsberg (2002, 2006, 2011). Among historians of economics, Carlo Zappia has extensively written on Ellsberg; see in particular Zappia (2016, 2018, 2021a, b).

Upon returning from Cambridge, what at that point seemed to be an academic career well under its way was interrupted for the first time: in 1953, Ellsberg volunteered for the US Marine Corps. From 1954 to 1957, he served as operation officer and company commander, and during the Suez Crisis of 1956, he was deployed to the Mediterranean with the Sixth Fleet.

In 1957, Ellsberg resumed his studies at Harvard on a prestigious three-year Junior Fellowship from the Society of Fellows. During this time, he continued to reflect on choice situations for which the existing theories of decision-making appeared problematic. He presented a paper on these topics at a meeting of the Econometric Society in September 1957 (Ellsberg 1958). Besides his academic activities, in March 1959 Ellsberg delivered a series of public lectures at the Lowell Institute in Boston (Ellsberg 1959). In these lectures, he addressed Cold War themes related to the threat of a nuclear war between the US and the Soviet Union by using game and decision theory.

In summer 1959, Ellsberg joined the RAND Corporation in Santa Monica, California, as a strategic analyst. RAND is a think tank established at the end of World War II by the US Army Air Force with the goal of bringing together civil scientists to work on research projects with possible military applications. While at RAND, Ellsberg pursued his research on decision-making under uncertainty and worked on his Ph.D. dissertation for Harvard. Concurrently, he became a consultant for the US Defense Department, the State Department, and the White House, advising on nuclear deterrence and crisis decision-making.

At the Econometric Society meeting held in St. Louis in December 1960, Ellsberg presented his paper on “Risk, Ambiguity, and the Savage Axioms.” The final version of the paper was published the following year in the November issue of the *Quarterly Journal of Economics* (*QJE*) (Ellsberg 1961a).

Between 1961 and early 1962, while increasingly engaged in consultancy for the US Defense Department, Ellsberg managed to complete his Ph.D. thesis on “Risk, Ambiguity and Decision.” He submitted it to the Harvard Economics Department in April 1962 and defended it in May (Ellsberg 2001). Thomas Schelling, a Harvard professor and a colleague of Ellsberg’s at RAND, served as the advisor on the thesis, though his role was mostly a formal one. The thesis largely drew on the 1961 article but diverged from it in some important respects.

The completion and defense of the PhD thesis marked the end of Ellsberg’s active engagement with decision theory. In 1962, he did not publish further articles, also because he was increasingly involved in consultancy to the State Department and the White House, especially during the Cuban missile crisis of October 1962. In 1963, he briefly replied to a critical comment on his *QJE* article (Ellsberg 1963). This is Ellsberg’s last academic paper, before he fully immersed himself as a military analysist for the Defense and State Departments.

Starting from 1952, when he began his academic research journey by completing his B.A. thesis, and accounting for his three years in the Marines, Ellsberg’s involvement in theoretical research within decision analysis spanned less than a decade.

Beginning in 1964, at RAND Ellsberg worked principally on issues related to the Vietnam war. In mid-1965, he volunteered to assess pacification programs at the US embassy in Saigon. During his time in Vietnam, he also accompanied combat units in ground operations. After returning to RAND in June 1967, Ellsberg participated in the task force established by Secretary of Defense Robert McNamara to compose a comprehensive history of the US involvement in the Indochina and Vietnam wars between 1945 and 1967. Completed in January 1969, the study was classified as top-secret. Unintentionally, it revealed how successive US administrations had misled Congress and the public regarding the management and prospects of the Vietnam war.

Two copies of this study, known as the Pentagon Papers, were sent to RAND. Ellsberg, who at that point had developed strong antiwar convictions, in late 1969 secretly photocopied the study. In 1971, he leaked it to *The New York Times* and *The Washington Post*. Although he faced charges of espionage and conspiracy, all charges were dismissed in 1973. From then until his passing in June 2023, Ellsberg continued his political activism, opposing US military involvements around the world and advocating for press freedom, while supporting other whistleblowers such as Julian Assange.

3 Decision theories between 1945 and 1960: a review

Ellsberg's (1961a) article can be situated within the burgeoning of theories for decision-making under risk and uncertainty prompted by the publication of *Theory of Games and Economic Behavior* by John von Neumann and Oskar Morgenstern ([1944] 1953). According to a classification introduced by Frank Knight (1921), in situations of risk the decision maker knows the probabilities of the payoff relevant events, whereas in situations of uncertainty she does not. This section reviews the theories of decision-making under risk and uncertainty that dominated the field in the period 1945–1960 and constitute the background of Ellsberg's research.

3.1 Risk: von Neumann and Morgenstern's EU

In *Theory of Games*, von Neumann and Morgenstern advanced a novel version of a theory for decision-making under risk that had been originally put forward by Daniel Bernoulli in the eighteenth century: expected utility theory (EU). They showed that if (and only if) a decision maker's preferences between risky prospects satisfy certain specific axioms, she will prefer the prospect with the highest expected utility. The expected utility of a prospect, expressed by the formula $\sum u(x_i)p(x_i)$, is the average of the utility values $u(x_i)$ of the potential outcomes of the prospect weighted by their respective probabilities $p(x_i)$.

From the late 1940s until the early 1950s, an intense debate developed about the exact assumptions underlying von Neumann and Morgenstern's EU, the normative plausibility of these assumptions and henceforth of EU, and the descriptive power of the theory. All the major economic and decision theorists of the time became involved. The outcomes of the discussion can be distilled as follows.[2] Jacob Marschak, Paul Samuelson, Savage, and other leading scholars accepted EU as a normative theory of rational behavior under risk, though they remained skeptical about its descriptive power. Other scholars, including Milton Friedman and Armen Alchian, accepted EU

[2] For more on the debates about EU between 1945 and 1955, see Moscati (2016, 2018, chapters 10–13).

because they considered it a simple and descriptively valid theory that was also supported by experimental studies, such as one by Frederick Mosteller and Philip Nogee (1951). In any case, by the early 1950s, the majority of economists and decision theorists had come to accept EU as a theory of individual decision-making that is normatively and/or descriptively sound.

During this period, the main challenger of EU was the French economist Maurice Allais, who rejected EU on both normative and descriptive grounds (Allais 1953). Part of Allais's argument hinged on a choice situation later known as the Allais paradox. As discussed in Sect. 7.2, Ellsberg saw the possible connections between his paradox and Allais's only after completing the *QJE* article.

3.2 Complete ignorance: from minimax to Hurwicz's criteria

With respect to decision-making under uncertainty, in the early 1950s decision theorists focused on situations of "complete ignorance," in which the decision maker possesses no information or belief about the likelihood of different events. For such situations, three main decision criteria were proposed: the minimax, the minimax regret, and some version of Hurwicz's criterion.[3]

Von Neumann and Morgenstern ([1944] 1953) adopted the minimax criterion for analyzing zero-sum two-person games. Abraham Wald (1950) extended it to situations of nonstrategic uncertainty. This criterion states that, in conditions of complete ignorance, a decision maker should choose the action minimizing the maximum potential loss.

Savage (1951) introduced a variant named minimax regret. Regret is defined as the difference between the highest utility the decision maker could have obtained by choosing an alternative action and the utility obtained from the chosen action. Minimax regret prescribes the action minimizing the maximum possible regret.

Leonid Hurwicz, then a researcher at the Cowles Commission for Research in Economics, proposed a weighted combination of the best and worst outcomes. Let max_a denote the maximum or best possible outcome associated with action a and min_a its minimum or worst possible outcome. In one of his papers, Hurwicz (1951a) suggested that the decision maker should select the action with the highest value of the index $\alpha max_a + (1 - \alpha) min_a$. Here, α is a parameter reflecting the decision maker's level of optimism or pessimism. This model became widely known as the Hurwicz pessimism–optimism criterion. In particular, if $\alpha = 0$, the decision maker is maximally pessimistic, and the criterion coincides with minimax.

In a separate paper, Hurwicz (1951b) advanced a probabilistic variation of the pessimism-optimism model, which he termed the "generalized Bayes–minimax criterion."[4] Within this model too, the decision maker should choose the action with the

[3] A fourth criterion is based on the "principle of insufficient reason" or "principle of indifference." According to this criterion, when a decision maker lacks information about the relative likelihood of different events, she should consider all events as equally likely and make decisions accordingly. In the 1950s, this criterion faced numerous criticisms and garnered limited support in the field (see Luce and Raiffa 1957, 284–286).

[4] One may argue about whether Hurwicz's model is genuinely Bayesian but such a debate and, more generally, a discussion about the true meaning of Bayesianism, falls beyond the scope of the present article.

highest value of the index $\alpha max_a + (1 - \alpha) min_a$. However, the meaning of max_a and min_a is different. Hurwicz now assumed that the decision maker is not totally ignorant but knows that the true probability distribution over the uncertain events belongs to a certain set Y of probability distributions. With this information, the decision maker can calculate the expected payoff of any given course of action a with respect to each probability distribution in the set Y. Within this model, max_a and min_a signify the maximum and minimum expected payoffs, respectively, of course of action a, whereby these values are calculated over the set of probability distributions in Y.

As discussed in Sects. 6 and 7, the models proposed by Ellsberg to account for decision-making in conditions of ambiguity draw significantly on Hurwicz's generalized Bayes–minimax criterion.

3.3 The direct and indirect approaches to subjective probability

In several uncertain situations, the decision maker is not in a state of complete ignorance but possesses some information about the distinct events and holds beliefs about their relative likelihood. In the literature of the 1950s, this circumstance was called a state of "partial ignorance."

Under certain assumptions, the decision maker's beliefs about the relative likelihood of events can be quantified, and the number attached to each belief can be interpreted as the "subjective" or "personal" probability ascribed by the decision maker to that event. When subjective probabilities can be derived, the problem of decision-making under partial ignorance reduces to the problem of decision-making under risk, and EU can be used again. Historically, two distinct approaches have been advanced to develop a theory of subjective probability: the intuitive-probability approach and the preference-based approach.

The intuitive-probability approach, also termed the direct approach to subjective probability, was advocated by probability theorists such as Bernard O. Koopman (1940) and Irving J. Good (1952), among others. Within this approach, the individual's beliefs are the primitive element of the analysis: it is assumed that the individual holds beliefs about the relative likelihood of events prior to and independently of her preferences between alternative courses of action and that he can access these beliefs through introspection. This approach, which Ellsberg eventually came to favor (see Sect. 7.3), has always remained a minority position within the economic theory of decision-making.

Since at least the 1950s, the dominant approach in economics has been the preference-based approach, also referred to as the indirect approach to subjective probability. This approach was separately introduced by Bruno de Finetti ([1931] 1993) and Frank Ramsey (1931) and further advanced by Savage in *The Foundations of Statistics* ([1954] 1972). Within this line of research, the decision maker's preferences

Footnote 4 continued
When Hurwicz, Ellsberg, and other scholars under consideration refer to Bayes or Bayesian theory, I adhere to their terminology without questioning its appropriateness. For a comprehensive discussion of the diverse and changing meanings of the term "Bayesian" in statistics and decision theory see Fienberg (2006) and Marinacci (2015).

between alternative courses of actions stand as the primitive element of the analysis, and subjective probabilities are inferred only indirectly from those preferences. Generally, the alternative courses of actions are alternative monetary bets.[5]

3.4 Savage's EU and the sure-thing principle

Savage demonstrated that if the decision maker's preferences between courses of actions with uncertain outcomes satisfy seven postulates, designated as P1–P7, it becomes possible to (1) identify a unique probability measure defined over the set of uncertain events and interpret it as expressing the subjective probabilities that the decision maker attaches to the events, and (2) model the decision maker's behavior as though she maximized expected utility, where expected utility is calculated by employing the subjective probabilities identified at the previous step. Because it involves subjective probabilities, Savage's theory is often referred to, also by Ellsberg, as "Bayesian decision theory."[6]

In Savage's axiomatization, postulate P2 plays a central role. It requires that the decision maker's preference between two courses of actions does not change when the payoffs corresponding to events for which both actions yield the same payoff change. For instance, if bets a_1 and a_2 both yield \$100 when a certain event E occurs, and the decision maker prefers a_1 to a_2, she should continue to prefer a_1 even if both bets yield \$0 rather than \$ 100 under event E.

P2 embodies the key part of what Savage termed the Sure-Thing principle. The principle is expressed by P2 combined with postulates P3 and P7. However, in the decision-theoretical literature, it has typically been identified solely with P2. The two counterexamples to Savage's version of EU that Ellsberg presented in the *QJE* article target P2, which Ellsberg also equated with the Sure-Thing principle.

3.5 Savage's reflective notion of normativity

To fully grasp Ellsberg's (1961a) article, it is crucial to recall that Savage advocated a normative interpretation of postulates P1–P7, considering them as maxims of rational behavior rather than descriptions of actual behavior. According to Savage, the normative, and therefore "rational," nature of the postulates does not derive from any logical or a priori principle but from the circumstance that the decision maker deliberately accepts these postulates as sensible criteria and wants to conform to them. In an often-quoted passage from the *Foundations*, Savage ([1954] 1972, 7) explained:

[5] To understand how the preference-based approach works, consider a situation involving two bets, a_1 and a_2, and two possible events, E_1 and E_2. Bet a_1 yields \$100 if E_1 occurs, and \$0 if E_2 occurs; a_2 is symmetrical to a_1: it yields \$0 if E_1 occurs, and \$100 if E_2 occurs. Under the assumption that the decision maker prefers more money to less money, if she prefers bet a_1, it seems reasonable to infer from her preference that she judges E_1 more likely than E_2. Conversely, if the decision maker prefers a_2, we can infer that she considers E_2 more likely than E_1.

[6] On the diverse meanings of the term "Bayesian" in economics, statistics, and decision theory, see footnote 4.

> I am about to build up a highly idealized theory of the behavior of a "rational" person with respect to decisions. … When certain maxims are presented for your consideration [as "rational"], you must ask yourself whether you try to behave in accordance with them.

As discussed in Sect. 5.1, Ellsberg fully embraced Savage's reflective notion of normativity and attempted to identify situations in which decision makers deliberately refuse to conform to Savage's postulates.

4 Ellsberg's path to the *QJE* article

Since his BA thesis of 1952, Ellsberg's research in decision theory had been closely connected with the research overviewed in the preceding section. Ellsberg received feedback on his papers from influential economists and decision theorists of the period. In addition to John Chipman, his advisor for the Harvard BA thesis, notable figures such as Paul Samuelson (MIT); Oskar Morgenstern and Martin Shubik (Princeton); Frederick Mosteller, Thomas Schelling, and Wassily Leontief (Harvard); Kenneth Arrow (Stanford), Lloyd Shapley (RAND); and Nicholas Kaldor (Cambridge) commented on Ellsberg's works. This section delves into Ellsberg's intellectual trajectory from his BA thesis to his article in the *QJE*.

4.1 Different utility indices

The first article Ellsberg derived from his thesis, titled "Classic and Current Notions of 'Measurable Utility'," contributes to the debate on von Neumann and Morgenstern's EU. In this work, Ellsberg (1954) clarified the conceptual difference between the utility index that early economists such as William Stanley Jevons and Alfred Marshall had employed to analyze choices between riskless options, and the utility index featured in von Neumann and Morgenstern's EU that was used to analyze choices between risky options.

Importantly, even in this first work, Ellsberg expressed reservations about the descriptive and normative validity of EU. At the descriptive level, he noted that alternative approaches to the theory of decisions under risk, such as those based on the mean, variance, or other elements of the distribution of monetary payoffs, "might produce fully as good predictions" as those based on EU (554). At the normative level, he expressed doubts about the consistency requirements imposed on decision makers' behavior by the EU axioms (554–555).

4.2 The reluctant duelist

The other article Ellsberg extracted from his BA thesis is titled "Theory of the Reluctant Duelist" and addresses decision-making in situations involving strategic uncertainty. In this paper, Ellsberg (1956) criticized the minimax criterion advocated by von Neumann and Morgenstern ([1944] 1953) as the only rational norm of behavior in two-person zero-sum games. Ellsberg argued that, in several situations, the minimax strategy is too cautious and defensive. He contended that the minimax criterion is appropriate solely for the special case of the reluctant duelist: an individual "who is forced, reluctantly, to make decisions" and whose sole concern is "to come out with as little loss as possible" (1956, 922).

4.3 Savage arrives in town

During Ellsberg's tenure in the Marines from 1954 to 1957, the main development in decision theory was the publication of Savage's *Foundations* ([1954] 1972). As attested by Duncan Luce and Howard Raiffa's authoritative book on game and decision theory (1957), Savage's approach to subjective probability and his version of EU quickly gained prominence.

Upon returning to Harvard, Ellsberg turned his attention to Savage's approach. While he admired Savage's theoretical construction, he remained skeptical about the possibility of reducing all uncertainties to risks. As Ellsberg (2011, 222–223) later reflected:

> Savage's innovative generalization of the Bernoulli proposition … implied that all uncertainties could be expressed … in the same probabilistic terms as spins of a well-balanced roulette wheel. Compelling as I did find Savage's axioms and logic, I still found this hard to believe.

Ellsberg began looking for uncertain situations of partial ignorance that resisted reduction to risky situations and in which, therefore, one or more of Savage's axioms were violated. Initially, and until the fall of 1957, his suspects did not encompass P2, which he considered the most plausible of Savage's axioms (Ellsberg 2001, 244, fn. 2).

4.4 Russian roulette

Ellsberg's explorations led to a paper titled "Winning at the Russian Roulette," presented at the September 1957 Econometric Society meeting held in Atlantic City (Ellsberg 1958; see also Ellsberg 2001, 267–269). Notably, during the same event, Ellsberg chaired a session on "Subjective Probabilities and Utility Theory", within which Luce (Harvard) and Jacques Drèze (Carnegie Institute of Technology) delivered papers, and Savage served as their discussant (Econometric Society 1958). It is most likely during this session that Ellsberg became acquainted with Savage, who was then affiliated with the University of Chicago.

In the Russian roulette paper, Ellsberg argued that in situations combining risk and uncertainty, the norms for rational behavior under risk as defined by the EU axioms may conflict with the norms for rational behavior under uncertainty as defined by criteria such as minimax, minimax regret, and the Hurwicz criterion. As an illustrative example, Ellsberg considered the case of a fan of Russian roulette who prefers the roulette yielding "Life" with probability 5/6 and "Death" with probability 1/6, to either "Life for sure" or "Death for sure." Ellsberg contended that these preferences violate the EU axioms but can be rationalized by assuming that the Russian roulette fan seeks to minimize maximum regret. However, it is worth noting that the paper never reached publication, likely due to the questionable idea of using a Russian roulette player as a paradigm of rational behavior.

4.5 From Chipman's boxes to Ellsberg's urns

A pivotal source of inspiration for identifying uncertain situations that cannot be reduced to risky situations came from an experimental study conducted by Ellsberg's advisor for the Harvard BA thesis, namely Chipman. In July 1957, Chipman, who had moved from Harvard to the University of Minnesota in 1955, undertook a study (later published as Chipman 1960) that tested a probabilistic version of the transitivity axiom for choices.[7] Chipman devised boxes containing kitchen matches broken in two parts, heads and stems. One set of boxes, called "known boxes," contained heads and stems in known proportions: 60–40, 50–50, and 40–60. Another set of boxes, called "unknown boxes," contained heads and stems in unknown proportions; however, participants in the experiment could extract a sample of 10 matches from each box and ascertain the proportion of heads and stems in the sample. Subsequently, the participants were presented with a wager: if a head was drawn, they would earn 25 cents. They were then asked to choose between drawing from a "known" box or an "unknown" box.

In a specific session, participants were required to choose between a known box containing 50 heads and 50 stems, and an unknown box from which a sample yielding 5 heads and 5 stems had been drawn (Chipman 1960, 88). As Ellsberg (1961a, 651, fn. 9) acknowledged, the setup of Chipman's experiment is "almost identical" to the setup of his first experiment with two urns (see Sect. 5.2).

It remains unclear when Ellsberg became acquainted with Chipman's study. Chipman presented the work at the Philadelphia meeting of the Econometric Society held in December 1957 (Econometric Society 1958). Therefore, it seems plausible that Ellsberg gained access to a copy of Chipman's paper by the end of 1957 or the beginning of 1958. Regardless, we do know that by February 1958 Ellsberg had devised his own choice situations using urns, because in that month he presented them to Savage (Ellsberg (2001, xlix).

4.6 Toward St. Louis

As mentioned in Sect. 2, in March 1959 Ellsberg delivered a series of lectures at the Lowell Institute in Boston. One lecture provided the basis for a paper titled "The Crude Analysis of Strategy Choices," which Ellsberg presented at the December 1960 meeting of the American Economic Association held in St. Louis (Ellsberg 1961b). In this short essay, Ellsberg proposed to use the tools of game theory to illuminate the decision problem faced by the US and the Soviet Union about launching or not launching a nuclear attack on the opponent.

[7] The probabilistic version of the transitivity axiom states that if the decision maker chooses option x over option y with a frequency greater than 0.5, and she chooses y over z with a frequency greater than 0.5, then she should choose x over z with a frequency greater than 0.5.

Finally, we arrive at Ellsberg's paper titled "Risk, Ambiguity, and the Savage Axioms." Between 1959 and 1960, he presented this paper in seminars held at Harvard, Chicago, Yale, and Northwestern Universities, the Interdisciplinary Colloquium on Mathematics in the Behavioral Sciences at UCLA, and RAND (Ellsberg 2001, 134, fn. 2). On 28 December 1960, he presented the paper at the Econometric Society meeting, also held in St. Louis in connection with the American Economic Association meeting, in a session named "Economic Theory and Method" in which participated, among others, Tjalling Koopmans (Yale), Lionel McKenzie (Rochester), and Hurwicz (who had moved from the Cowles Commission to the University of Minnesota) (Econometric Society 1961; Ellsberg 2001, 179, fn. 1).

A revised version of the essay featuring an alternative model for decision-making under ambiguity was published in August 1961 as RAND document P-2173. The final version of the paper, nearly identical to the RAND version, appeared in the November 1961 issue of the *QJE*, in a symposium on "Decisions under Uncertainty." This symposium also included an article by William Fellner (Yale) and a commentary on Ellsberg's and Fellner's papers by Raiffa (Harvard).

5 The *QJE* article, part I: Ellsberg's urns

5.1 Normative falsification and hypothetical experiments

Ellsberg embraced Savage's normative interpretation of postulates P1–P7 and Savage's reflective notion of normativity (see Sect. 3.5). In the *QJE* article, Ellsberg aims to identify a class of choice situations in which a significant proportion of reasonable decision makers, upon reflection and deliberation, refuse to behave in accordance with Savage's axioms: "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" (1961a, 646). This type of goal has been termed a goal of normative falsification (see Guala 2000; Zappia 2016). For Ellsberg, however, Savage's postulates retain their normative validity in choice situations beyond the specific class he envisaged.

Ellsberg noted that actual choice experiments such as those conducted at the time by Mosteller and Nogee (1951), Davidson et al. (1957), and Chipman (1960), are not suited to normative falsification because they also record careless, instinctive, or even random choices. For Ellsberg, the appropriate method for normative falsification involves using "purely hypothetical experiments" (1961a, 651, fn. 9). In these, selected individuals who appear otherwise reasonable are presented with hypothetical choice situations and are interrogated about which options they prefer. The individuals are given ample room to ponder their answers and evaluate the consequences of their stated preferences and are allowed to correct the preferences if they wish to do so.

5.2 Experiment #1: Urn I or Urn II?

Ellsberg considered two hypothetical choice situations, both featuring urns containing balls of different colors. In the first situation, the less renowned of the two, the decision maker faces two urns containing red and black balls and labelled as Urn I and Urn II. The decision maker is told that Urn I contains 100 red and black balls in an unknown ratio, whereas Urn II contains exactly 50 red and 50 black balls. As already noted, this scenario resembles that of Chipman's experiment, with Ellsberg's Urn II corresponding to Chipman's known match box and Urn I corresponding to Chipman's unknown match box.

In Ellsberg's experiment, the decision maker is asked to express her preferences between different pairs of bets. In the first pair, bet a_1 is "draw from Urn I, and receive \$100 if the drawn ball is red, and \$0 if it is black," and bet a_2 is "draw from Urn II, and receive \$100 if the drawn ball is red, and \$0 if it is black." In the second pair, bet a_3 is "draw from Urn I, and receive \$100 if the drawn ball is black, and \$0 if it is red," and a_4 is "draw from Urn II, and receive \$100 if the drawn ball is black, and \$0 if it is red."

Ellsberg observed that most subjects preferred a_2 to a_1 and a_4 to a_3.That is, most subjects tended to favor drawing from Urn II in both instances. However, a minority of subjects displayed the opposite preference pattern, a_1 preferred to a_2 and a_3 preferred to a_4, favoring drawing from Urn I. Both preference patterns challenge the preference-based approach to subjective probability. Here, we discuss the majority pattern, as similar arguments apply to the minority pattern.

Under the plausible assumption that the decision maker prefers more money to less, her preference for a_2 over a_1 indicates that she considers the event "drawing a red ball from Urn II" (call this event Red_{II}) more likely than the event "drawing a red ball from Urn I" (call it Red_I). If this likelihood judgement could be represented by probability numbers p, the decision maker's first preference implies therefore that for her $p(Red_{II}) > p(Red_I)$. However, the second preference ranking (a_4 preferred to a_3) indicates exactly the opposite: that for the decision maker $p(Red_{II}) < p(Red_I)$. In fact, if she prefers a_4 to a_3, this means that she judges the event "drawing a back ball from Urn II" ($Black_{II}$) more probable than the event "drawing a black ball from Urn I" ($Black_I$): $p(Black_{II}) > p(Black_I)$. As both urns contain only red and black balls, by the principle of additivity of probability theory, $p(Red_I)+p(Black_I)=1$ and $p(Red_{II}) + p(Black_{II}) = 1$. But if $p(Black_{II})>p(Black_I)$ then, by the additivity principle, $1-p(Red_{II})>1-p(Red_I)$ or, what was to be shown, $p(Red_{II}) < p(Red_I)$.

According to Ellsberg, this finding shows that in some situations the de Finetti–Ramsey–Savage approach to subjective probability fails, that is, in some situations, inferring subjective probabilities from the decision maker's preferences is impossible. Imagining a direct conversation with the decision maker, referred to as "you," Ellsberg (651) stated:

> An observer, applying the basic rule of the Ramsey–Savage approach, … must conclude that your choices are not revealing judgments of 'probability' at all. So far as these events are concerned, it is *impossible* to infer probabilities from your choices; you must inevitably be violating some of the Savage axioms.

In particular, Ellsberg argued that the decision maker's preference for a_2 over a_1 and a_4 over a_3 violated either Savage's postulate P1, the assumption that her preferences between courses of action are complete and transitive, or P2, or both.

The inability to infer subjective probabilities from the decision maker's preferences holds significant consequences for the applicability of EU. In fact, if subjective probabilities cannot be identified, calculating the expected utility of each course of action becomes impossible, and therefore modeling the decision maker's behavior as though she were maximizing expected utility is infeasible. Ellsberg (655) emphasizes that:

> it is impossible, on the basis of such choices, to infer … probabilities for the events in question. Moreover … it is impossible to find probability numbers in terms of which these choices could be described – even roughly or approximately – as maximizing the mathematical expectation of utility.

5.3 Experiment #2 (Ellsberg's paradox): red, black, or yellow?

Ellsberg's second hypothetical experiment is the most renowned of the two, as it directly tests a single postulate of Savage's theory, P2 or the Sure-Thing Principle. When people refer to "the Ellsberg Paradox," they typically mean this specific choice situation.

In this scenario, the decision maker is presented with a single urn containing 90 balls of three colors: red, black, and yellow. Specifically, the urn contains 30 red balls and 60 balls that are either black or yellow, with the proportions of the latter two being unknown. The decision maker is once again asked to express her preferences between different pairs of bets whose outcomes depend on the colors of balls drawn from the urn. In the first pair, bet a_1 yields \$100 if the ball drawn is red and \$0 if it is black or yellow; bet a_2 yields \$100 if the ball drawn is black and \$0 if it is red or yellow. Bets a_1, a_2 and their payoffs are displayed in decision matrix 1:

Decision matrix 1

	Red	*Black*	*Yellow*
a_1	\$100	\$0	\$0
a_2	\$0	\$100	\$0

In the second pair of bets, bet a_3 yields \$100 if the ball drawn is red or yellow and \$0 if it is black; bet a_4 yields \$100 if the ball drawn is black or yellow and \$0 if it is red. Decision matrix 2 represents bets a_3 and a_4:

Decision matrix 2

	Red	*Black*	*Yellow*
a_3	\$100	\$0	\$100
a_4	\$0	\$100	\$100

Ellsberg observed that the most common response pattern was a_1 preferred to a_2, and a_4 preferred to a_3. However, the opposite pattern, a_2 preferred to a_1, and a_3 preferred to a_4, was also observed at times. The crux of the matter lies in the fact that both patterns violate Savage's axioms, and more precisely the Sure-Thing principle. As discussed in Sect. 3.4, this principle stipulates that, if a_1 is preferred to a_2, then a_3 should be preferred to a_4 and vice versa, and this is because a_1 and a_3 on one hand, and a_2 and a_4 on the other hand, differ only in their third-column payoffs, which is the same for each pair of bets (either \$0 or \$100).

As in the first experiment, this violation makes it impossible to infer probabilities from the decision maker's preferences. Her preferring a_1 to a_2 indicates that she considers the event "drawing a red ball" more probable than the event "drawing a black ball". However, her preference for a_4 over a_3 indicates just the opposite, namely that she considers the event "drawing a black ball" more probable than the event "drawing a red ball".[8] As in the first choice situation, if subjective probabilities cannot be inferred from the decision maker's preferences, her choice behavior cannot be modeled as though she were maximizing expected utility by using EU.

Ellsberg (656) examined the possibility that the majority behavior in both choice situations could be rationalized using minimax, minimax regret, or Hurwicz's (1951a) pessimism–optimism criterion. However, he swiftly dismissed this possibility: because each pair of bets within the choice situations had the same minimum (\$0) and maximum (\$100) values, the three criteria imply that decision makers should be indifferent between the various bets, which was not the case.

5.4 Deliberate violators

The group of "otherwise reasonable people" (1961a, 646) to whom Ellsberg presented his two hypothetical choice situations consisted of the participants to the seminars where he presented the paper between 1959 and 1960 (see Sect. 4.6) and included some of the major theorists of the period: Paul Samuelson, Jacob Marschak (UCLA), Howard Raiffa, Gerard Debreu (Yale), Robert Schlaifer (Harvard), Norman Dalkey (RAND), and Savage himself, who was tested in February 1958 (Ellsberg 1961a, 654, fn. 4, 655–656; Ellsberg 2001, xlix). As discussed in Sect. 5.1, Ellsberg was not seeking decision makers whose preferences merely violated Savage's axioms; rather, he sought what he later termed "deliberate violators" (Ellsberg 1963, 337). These are decision makers who, after careful consideration, after having well understood the meaning of Savage's axioms and the fact that their preferences violated them, consciously decide to uphold their preferences.

Whereas some individuals tested, including Samuelson, Debreu, and Schleifer, did not violate the axioms, others such as Raiffa tended to violate them but felt "guilty about it" and went back "into further analysis" (656). Yet others violated the axioms and

[8] In more detail, the decision maker's preference for a_1 over a_2 indicates that for her $p(Red) > p(Black)$. However, her preference for a_4 over a_3 indicates that she deems $p(Black\ or\ Yellow) > p(Red\ or\ Yellow)$. Given that the urns only contain red, black, and yellow balls, the additivity principle of probability requires that $p(Black\ or\ Yellow) + p(Red) = 1$ and $p(Red\ or\ Yellow) + p(Black) = 1$. Hence, $p(Black\ or\ Yellow) > p(Red\ or\ Yellow)$ implies that $1 - p(Red) > 1 - p(Black)$ or, in other words, $p(Red) < p(Black)$.

maintained their choices. The group of deliberate violators included, besides Ellsberg himself, Marschak, Dalkey, and even Savage:

> Some violate the axioms cheerfully, even with gusto (J. Marschak, N. Dalkey); others sadly but persistently, having looked into their hearts, found conflicts with the axioms and decided … to satisfy their preferences and let the axioms satisfy themselves. … Since this group included L.J. Savage, when last tested by me …, it seems to deserve respectful consideration. (656)

6 The *QJE* article, part II: modelling ambiguity

In the final section of his paper, Ellsberg aimed to clarify the distinguishing features of the two choice situations he envisaged, and he advanced a decision model that could account for the preferences stated by a majority of decision makers in these situations.

6.1 Ambiguity

For Ellsberg (657–658), the two choice situations involving urns are characterized by a form of uncertainty that lies between the partial ignorance that can be reduced to risk by using Savage's theory and the complete ignorance for which the minimax criterion or its variations were conceived. In the urn examples, decision makers are not dealing with risk, because it is not feasible to attach probabilities to certain events, neither directly through introspection, nor indirectly by inferring them from stated preferences. Nonetheless, these decision makers are not in the realm of complete ignorance either. In the first experiment, they know that Urn II contains exactly 50 red and 50 black balls, and that Urn I contains only red and black balls. In the second experiment, the subjects know that 30 over 90 of the balls in the urn are red, and that the rest are either black or yellow. The problem lies in the fact that some of the information they have about the urn composition is vague or scant, leading to low confidence in making decisions based on it.

Borrowing the term from psychological literature (Frenkel-Brunswik 1949; Hamilton 1957), Ellsberg termed this kind of situations "ambiguous." For him, ambiguity has to do with situations in which.

> so many of the probability judgments an individual can bring to bear upon a particular problem are either "vague" or "unsure" that his confidence in a particular assignment of probabilities, as opposed to some other of a set of "reasonable" distributions, is very low. (660)

Ellsberg claimed that, in ambiguous situations, consistent and deliberate choice behavior that violates the Savage axioms "may commonly occur" (660).

6.2 Ambiguity, not ambiguity aversion

Within the present literature in decision theory, ambiguity is typically coupled with, if not equated to, "ambiguity aversion," that is, with a preference for non-ambiguous options over ambiguous ones. Yet, Ellsberg was concerned with the overall effects of ambiguity on decision-making, rather than solely focusing on ambiguity aversion. In particular, he considered the minority of subjects who deliberately opted for the more ambiguous alternatives in the two-urn or three-color-urn cases no less reasonable than those who preferred the less ambiguous options.

In a brief comment written on the occasion of the fiftieth anniversary of the publication of the *QJE* article, Ellsberg (2011, 225) called attention to this aspect of his research:

> I have never personally regarded the phenomenon I was investigating as "ambiguity aversion" … In the *QJE* article … I repeatedly mentioned that some subjects deliberately and consistently chose the more ambiguous alternative …. I regarded these choices as no less reasonable … than the opposite behavior.

6.3 Two models for decision-making under ambiguity

In the version of the paper that he presented at the St. Louis meeting in December 1960, Ellsberg advanced a model to account for decision-making in ambiguous situations that is largely analogous to Hurwicz's (1951b) generalized Bayes–minimax criterion (see Sect. 3.2). Hurwicz, who attended Ellsberg's presentation, noted the similarity between their models (Ellsberg 2001, 179, fn. 1).

After the St. Louis meeting, Ellsberg modified his model for decision-making under ambiguity, and in the *QJE* article, he presented a different model that drew on the decision criterion proposed by statisticians Joseph Hodges Jr. and Erich Lehmann (1952). The *QJE* model has the following components:

(1) As in Hurwicz's generalized Bayes–minimax criterion, the decision maker is not totally uninformed but knows that the true probability distribution over the uncertain events falls within a specific set Y of probability distributions. For instance, in the three-color urn experiment, Y includes all probability distributions of the form $\left(\frac{1}{3}, \lambda, \frac{2}{3} - \lambda\right)$, where $\frac{1}{3}$ is the probability of drawing a red ball, λ is the probability of a black ball, $\frac{2}{3} - \lambda$ is the probability of a yellow ball, and λ ranges between 0 and $\frac{2}{3}$.

(2) For any given course of action a, it is possible to calculate its expected utility with respect to each probability distribution in the set Y. In Ellsberg's experiment, if we assume that $u(\$0) = 0$ and $u(\$100) = 6$, the expected utility of action a_1 ("bet on red") is $6 \times \frac{1}{3} + 0 \times \lambda + 0 \times \left(\frac{2}{3} - \lambda\right) = 2$, and the expected utilities of a_2, a_3, and a_4 are 6λ, $6(1 - \lambda)$, and 4, respectively.

(3) The variable min_a denotes the minimum expected utility of course of action a across the set of probability distributions in Y. For course of action a_1, whose expected utility is 2 for all values of λ, $min_{a_1} = 2$. For a_2, the minimum occurs

when there are no black balls in the urn, i.e., when $\lambda = 0$ and $min_{a_2} = 0$. It is easy to see that $min_{a_3} = 2$ and $min_{a_4} = 4$.

(4) The main difference from Hurwicz's generalized Bayes–minimax criterion is that Ellsberg assumed that, among all probability distributions that the decision maker considers possible, one specific distribution stands out as more plausible than the others. This distribution is referred to as the "estimated distribution" and is labelled as y_0. In the three-color urn experiment, y_0 could be $\left(\frac{1}{3}, \frac{1}{3}, \frac{1}{3}\right)$, for example.

(5) The variable est_a denotes the expected utility of action a calculated by using the estimated distribution y_0. In our example, est_{a_1} is given by $6 \times \frac{1}{3} + 0 \times \frac{1}{3} + 0 \times \frac{1}{3} = 2$; likewise, the estimated expected utilities of the other actions are $est_{a_2} = 2$, $est_{a_3} = 4$, and $est_{a_4} = 4$.

(6) From Hodges and Lehmann (1952), Ellsberg takes the idea of a parameter ρ that denotes the decision maker's degree of confidence in the estimated distribution y_0. Here, ρ ranges between 0 and 1, with increasing values signifying greater confidence. In particular, when $\rho = 0$, the decision maker has no confidence in the estimated distribution, whereas when $\rho = 1$, she has complete confidence in the estimated distribution, which means that there is no ambiguity.

According to Ellsberg (1961a, 664), the simplest model that accounts for decision-making in conditions of ambiguity and also offers normative guidance in such situations posits that the decision maker chooses the course of actions associated with the highest value of the index $\rho \times est_a + (1 - \rho) \times min_a$. Notably, when $\rho = 0$ the model aligns with minimax, whereas for $\rho = 1$ it coincides with EU.

For actions $a_1 - a_4$, Ellsberg's index (referred to as E) takes these values: $E_{a_1} = 2$, $E_{a_2} = 2\rho$, $E_{a_3} = 2\rho + 2$, $E_{a_4} = 4$. He showed that, by assuming that $\rho = \frac{1}{4}$, but in fact for any $\rho < 1$, $E_{a_1} > E_{a_2}$ and $E_{a_4} > E_{a_3}$, which means that a decision maker following Ellsberg's decision rule would prefer a_1 to a_2 and a_4 to a_3, just as the majority of the subjects he tested.

Ellsberg (669) concluded his article by reiterating that judging the behavior of these individuals as irrational would be misguided, and that for them "the Bayesian or Savage approach gives wrong predictions and, by their lights, bad advice."

6.4 Ellsberg vs. Fellner and Raiffa

As mentioned in Sect. 4.6, Ellsberg's article was published as part of a symposium titled "Decisions under Uncertainty," which also included an article by Fellner and a commentary by Raiffa on Ellsberg's and Fellner's articles. Here, we briefly review Fellner's and Raiffa's papers, and Ellsberg's responses to them presented in his PhD dissertation.

Fellner considered choice situations akin to those envisaged by Ellsberg's in his first experiment. In Fellner's experiment, subjects had to choose whether "to make a prize contingent on drawing the color of their choice from a deck whose composition was wholly unknown to them," akin to Ellsberg's Urn I, or "from a deck of guaranteed composition," akin to Ellsberg's Urn II (Fellner 1961, 687). Like Ellsberg, Fellner found that most subjects preferred the guaranteed deck, i.e., the unambiguous option.

However, Fellner's explanation differed from Ellsberg's: Fellner conjectured that, in the presence of uncertainty, subjects employ psychological decision weights that distort genuine subjective probabilities. Fellner termed these decision weights "slanted probabilities."

In his PhD dissertation, Ellsberg (2001, 114–115, 172–176) acknowledged the similarities between his work and Fellner's but criticized the concept of "slanted probabilities," arguing that the phenomena Fellner discussed are better captured by probability intervals (see Sect. 7.3). Moreover, Ellsberg remarked that Fellner's subjects were not encouraged to reach deliberate decisions, so that his experiment was not suitable for normative falsification.

In contrast, Raiffa's paper was indeed about normative falsification or, better, normative validation: he aimed at validating the normative validity of Savage's postulates by showing that, if Ellsberg's choice situations are opportunely reframed, individuals do not breach the postulates and do not want to. In particular, Raiffa (1961) reframed Ellsberg's choice situations by considering probabilistic mixtures of Ellsberg's original bets. For example, Raiffa reframed the choice between bets a_1 and a_2 and the choice between a_3 and a_4 in Ellsberg's three-color experiment as a single choice between option A, in which an unbiased coin is tossed and a_1 is taken if heads appears and a_4 is taken if tails appears, and option B, for which heads leads to a_2 and tails to a_3.

In his PhD thesis, Ellsberg (2001, 108, 134, 241–246) contended either that Raiffa's reframing altered his original choice situations, thus making Raiffa's case irrelevant to the problem at issue, or that the preference pattern Raiffa deemed normatively compelling in the reframed situation was not compelling at all.

7 The PhD dissertation and the final academic article

Between 1961 and April 1962, Ellsberg worked on his PhD thesis titled "Risk, Ambiguity and Decision." The thesis, an extensive 389-page document, largely builds upon the *QJE* article, albeit with some significant novelties. Ellsberg did not submit the thesis for publication because he thought it needed substantial editing, but he was unable to complete this because of his commitments as a military analyst (Ellsberg 2011, 224–225).

The thesis opens with an acknowledgment of Ellsberg's intellectual debt to Savage:

> Nearly every page of this study testifies to my intellectual debt to L.J. Savage, whose developments in "Bayesian" statistical decision theory provide the starting-point and main focus for the present work. (xlix)

Ellsberg sent copies to Savage, Raiffa, and possibly other decision theorists but then put the thesis aside for around four decades. Eventually, it was published without modifications in 2001 (Ellsberg 2001).

In this section, we focus on three main aspects in which the thesis diverges from the 1961 article: the proposed model for decision-making under ambiguity, the discussion of the Allais paradox, and Ellsberg's endorsement of the direct approach to subjective probability.[9]

7.1 A third model for decision-making under ambiguity

In the PhD dissertation, Ellsberg once again modified the model he proposed to account for decision-making in situations of ambiguity. One main problem with the model presented in the *QJE* article is that it cannot account for the behavior of individuals who prefer a_2 to a_1 and a_3 to a_4 and whom today we would call ambiguity seekers. The reason is that, by focusing only on min_a, the model of the *QJE* article portrays the individuals as overly pessimistic. Accordingly, whenever they are not fully confident in the information they have, that is, whenever $\rho < 1$, they avoid the ambiguous option.

The model that Ellsberg put forward in the dissertation combines the model *à la* Hurwicz (1951b) used in the St. Louis version of the paper with the model *à la* Hodges and Lehmann (1952) featured in the *QJE* article. Ellsberg called this third model, which subsumes his other two models as special cases, the "restricted Bayes/Hurwicz criterion." The term "restricted" alludes to the assumption that the decision maker restricts her attention to probability distributions within the set Y.

The model states that the decision maker prefers the course of action associated with the highest value of the index $\rho \times est_a + (1-\rho) \times [\alpha max_a + (1-\alpha) min_a]$. In this formula, ρ, est_a, and min_a represent the same variables as in the *QJE* model; max_a denotes the maximum expected utility of course of action a across the probability distributions in Y; and α is Hurwicz's pessimism–optimism parameter.

Note that for $\rho = 1$ (full confidence in the available information), this model coincides with EU; for $\alpha = 0$ (pessimism), it coincides with the model presented in the *QJE* article; and for $\rho = 0$ (no confidence in the available information), it coincides with the model featured in the St. Louis version of the paper.

Ellsberg shows that, by choosing the values of the free parameters ρ and α opportunely, the model can account not only for the behavior of ambiguity-averse individuals who prefer a_1 to a_2 and a_4 to a_3 but also for the behavior of ambiguity seekers who prefer a_2 to a_1, and a_3 to a_4. However, he appears aware that the strategy of enhancing the model's descriptive power by introducing more free parameters carries potential pitfalls. In fact, he is careful to assert that he does not propose "that α and ρ be regarded as 'dummy variables,' useful for achieving pseudo-generality by combining in one portmanteau formula all earlier decision criteria" (2001, 196). In this statement, Ellsberg draws attention to a potential problem of his decision model, which today we might label overfitting, but he does not address it.

[9] Levi (2001) and Zappia (2016) provide further discussions of Ellsberg's Ph.D. thesis.

7.2 Ellsberg discovers Allais

Many contemporary readers of the *QJE* article may find it peculiar that neither Allais nor his paradox are mentioned in the paper.[10] In decision-theory literature, in fact, these two paradoxes are often associated: the Allais paradox deals with violations of EU in conditions of risk, whereas the Ellsberg paradox concerns violations of EU under uncertainty. Furthermore, Allais and Ellsberg shared similar normative concerns: both argued that the preferences violating EU recorded in the choice situations they envisaged were clearly defensible from a normative viewpoint.[11] So, why is Allais absent from the *QJE* paper?

The immediate reason is that before 1961 Ellsberg had not read Allais's 1953 article (see Ellsberg 2011, 224). However, Ellsberg almost certainly knew of the choice situation envisaged by Allais even before 1961 because Savage had discussed it in a famous passage of the *Foundations* ([1954] 1972, 101–103). In that passage, Savage argued that preference patterns *à la* Allais violate the Sure-Thing principle and should therefore be corrected.

Arguably, Ellsberg did not pay attention to the Allais paradox in 1961 because it concerned situations of risk, for which he believed that Savage's axioms and EU were normatively valid. In effect, as previously observed, in conditions of risk where ambiguity is absent and $\rho = 1$, the decision model that Ellsberg presented in the *QJE* article coincides with EU, thereby ruling out Allais-like preferences.

At some point between late 1961 and early 1962, Allais sent Ellsberg a copy of his 1953 article and some related works. In studying these materials, Ellsberg came to see the connection between his research questions and Allais's and began to doubt the normative validity of Savage's axioms even in conditions of risk. In a section of his PhD dissertation devoted to Allais, Ellsberg wrote (2001, 256) that:

> before studying carefully certain hypothetical examples proposed by Allais, I would have conjectured that the Savage postulates would prove acceptable to me as normative criteria in any situation providing a clear basis for "definite" probabilities: e.g., urns with known proportions of balls .... Now I am not so sure.

Ellsberg now recognized the normative validity of Allais-like preferences and noted the similarity between the types of violations of EU discussed by Allais and those he himself had described (263).

However, there was a problem: like the decision model featured in the *QJE* article, the model featured in the PhD dissertation rules out preferences *à la* Allais. In fact, in conditions of risk where ambiguity is absent, even Ellsberg's restricted Bayes–Hurwicz criterion reduces to EU. This means that preferences *à la* Allais, which Ellsberg came

---

[10] Allais (1953) imagined two pairs of lotteries L: the first pair consists of L_1, which pays 100 million Francs with $p = 1$, and L_2, which yields 500 million Francs with $p = 0.10$, 100 million Francs with $p = 0.89$, and 0 Francs with $p = 0.01$. The second pair consists of L_3, yielding 100 million Francs with $p = 0.11$ and 0 Francs with $p = 0.89$, and L_4, yielding 500 million Francs with $p = 0.10$ and 0 Francs with $p = 0.90$. Most decision makers prefer L_1 to L_2, and L_4 to L_3. However, this pair of preferences violates EU because no utility function is capable of rationalizing them by using the EU utility formula.

[11] On the normative goals of Allais's paradox, see Mongin (2019).

to judge as normatively valid, violate his own criterion for decision-making. Ellsberg appeared somehow aware of the problem, but neither in the PhD dissertation nor afterwards did he tackle it.

7.3 A direct approach to probability

The third significant innovation in the PhD dissertation is Ellsberg's endorsement of the direct approach to subjective probability advocated by Koopman (1940) and Good (1952, [1960] 1962). According to Ellsberg, Koopman and Good developed a line of thought initiated by John Maynard Keynes in the *Treatise on Probability* (1921), a book Ellsberg only read after completing the *QJE* article.[12]

In the dissertation, Ellsberg emphasized the difficulty of discussing phenomena related to vague or imprecise likelihood judgments that characterize ambiguous situations using the language of the de Finetti–Ramsey–Savage indirect approach to probability. In fact, this approach admits preference judgements and choice behavior as the only legitimate sources of evidence while rejecting likelihood judgments based on introspection. However, this methodological stance makes it difficult to capture the varying degrees of confidence that decision makers may have in their likelihood judgments. At the beginning of his PhD dissertation, Ellsberg noted (6–7) that:

> the phenomena of vagueness, imprecision or relative lack of confidence in certain of one's subjective judgments are given central importance in the work of I.J. Good, who, in contrast to Savage, approaches the theory of personal probability as a theory of consistent judgment rather than consistent action. … The language of the Koopman/Good theories has one marked advantage for my particular purposes; it is much easier in their approach to talk formally and precisely about vagueness.

In particular, prompted by the works of Koopman, Good, and statistician Cedric Smith (1961) Ellsberg (2001, 65–89, 115–125) suggested adopting an interval approach to probability. In this approach, the decision maker's degree of confidence in a likelihood judgment is represented by two probability numbers $\underline{p}$ and $\overline{p}$, which express the lowest and highest probabilities, respectively, that the decision maker attaches to the event. The stronger the confidence a decision maker has in her likelihood judgment, the smaller the interval $\left[\underline{p}, \overline{p}\right]$; in particular, if she is fully confident in her likelihood judgment, $\underline{p} = \overline{p}$.

Ellsberg discussed the probability-intervals approach extensively and showed that a decision maker can hold a system of beliefs defined in terms of probability intervals without falling victim to a Dutch book, that is, a set of bets involving a certain loss. However, he did not integrate the probability-intervals approach into his restricted Bayes–Hurwicz model for decision-making under ambiguity.

In summary, this section illustrates that Ellsberg's PhD dissertation reveals significant evolution in his views on the theory of decision-making in the months following the completion of the *QJE* paper. The dissertation introduces several new ideas and

[12] More on the relationships between Keynes's and Ellsberg's probability theories in Feduzi (2007) and Zappia (2021b).

insights, although these ideas remained in an embryonic state as Ellsberg left academia shortly after completing the thesis and did not develop them further.

7.4 The final academic article

Ellsberg's final academic article was a brief response to a critical comment on his *QJE* article written by statistician Harry Roberts. Roberts (1963, 330) argued that violations of the Savage axioms are due to "mistakes or misinterpretations" and that the violations would disappear if these mistakes and misinterpretations were avoided or corrected. Ellsberg (1963) countered that Roberts's analysis did not hold for the group of decision makers the *QJE* article was actually about: deliberate violators who do not consider their choices as erroneous or misguided and consciously maintain them.

From a historical perspective, Ellsberg's response to Roberts is valuable because it is the only place, prior to the publication of Ellsberg's PhD dissertation in 2001, where he hinted in print at some of the novel ideas he had advanced in the dissertation, notably his endorsement of the Koopman–Good direct approach to subjective probability (338–339).

8 Savage on Ellsberg

Before overviewing the experimental and theoretical literature that originated in Ellsberg's article, it is in order to shed more light on a coprotagonist of our narrative, Savage, and ask what he thought of Ellsberg's work.

In the *Foundations*, Savage had discussed some criticisms of the theory of personal probabilities developed in the book. Specifically, he had acknowledged that his theory was unable to capture the difference between probability judgments about which "we feel relatively 'sure' as compared with others" ([1954] 1972, 57–58). Savage recognized that his notion of sure and unsure probability judgements was "vague," but argued that no other theory he knew of could managed to render "the notion less vague" in a satisfactory manner (58). Ellsberg cited this passage from the *Foundations* in both the *QJE* article (1961, 660) and the PhD dissertation (2001, 12), asserting that it showed that Savage himself was aware of the difficulties his probability theory faced when dealing with decision-making in conditions of ambiguity.

As discussed also in the previous sections, Savage was aware of this and other criticisms that Ellsberg directed at his theory. However, he never published any systematic response to them. His only mention of Ellsberg's *QJE* article appeared in the bibliographical supplement to the second edition of the *Foundations* ([1954] 1972, 288). Here, he noted that the papers by Ellsberg, Fellner, and Raiffa featured in the *QJE* symposium were "a key reference for a certain type of departure from the theory of personal probability and utility in this book."

To glean more insight into Savage's views on Ellsberg's work, we can turn to some letters exchanged between Savage, de Finetti, and Fellner, as adeptly discussed by Zappia (2021a). In a letter to de Finetti on March 16, 1962, Savage admitted that "there may be a grain of truth in what he [Ellsberg] is trying to say, but [I] find it very difficult to clear my own head on the subject." In a letter to Fellner dated June 17, 1963, Savage acknowledged that the assumptions made in the *Foundations* "may be unsatisfactory," but added that he did not know how "to make a mathematical theory with more realistic assumptions."[13] These remarks to de Finetti and Fellner echo the previously mentioned passage in the *Foundations* where Savage recognized the limitations of his theory, but argued that no other available model was capable of overcoming them in a satisfactory way.

Combining all these elements, it seems plausible to speculate that Savage held moderate sympathy for Ellsberg's and Fellner's concerns but was not ready to abandon the admittedly imperfect yet systematic and mathematically solid theory he had developed until an alternative theory with comparable systematicity emerged. And neither Fellner's slanted-probabilities theory nor Ellsberg's models, which lacked an axiomatic foundation and had so many free parameters, met his criteria.

9 The influence of Ellsberg's article

As mentioned in the introduction, Ellsberg's *QJE* article has become a seminal contribution in the field of decision theory, garnering over 11,200 Google Scholar citations.[14] Initially, it gained prominence through a descriptive channel rather than the normative arguments in which Ellsberg was interested: the Ellsberg paradox was primarily employed to demonstrate the shortcomings of Savage's EU as a descriptively valid theory of decision-making under uncertainty.

Confidence in the descriptive validity of EU, both in von Neumann–Morgenstern's and Savage's versions of the theory, began to wane in the mid-1960s, when a series of laboratory experiments showed that the choice patterns originally conceived by Allais and Ellsberg were, as Amos Tversky (1969, 40) noted, "systematic, consistent, and predictable." In particular, the experimental studies conducted by Becker and Brownson (1964), MacCrimmon (1968), Slovic and Tversky (1974), and MacCrimmon and Larsson (1979) documented the frequency and systematic nature of Ellsberg-like preference patterns.[15]

Theoretical models capable of accounting for Ellsbergian choice patterns while adhering to the preference-based and axiomatic approach used by Savage began to appear much later, in the 1980s. Broadly speaking, these models replaced one or more of Savage's postulates P1–P7, most frequently P2, with weaker requirements.

The first of these models was proposed by David Schmeidler (1989). He suggested quantifying a decision maker's beliefs in conditions of ambiguity using quasi-probability numbers that do not satisfy the additivity property of probability and are called "capacities." Because the concept of capacity was defined by the French mathematician Gustave Choquet, Schmeidler's model is often referred to as

[13] Both letters are reported in Zappia (2021a, 181).

[14] See https://scholar.google.com. Last accessed January 20, 2024.

[15] For a review of these experimental literature, see Camerer and Weber (1992).

the Choquet expected utility (CEU) model. Itzhak Gilboa, together with Schmeidler, introduced a different model for decision-making under ambiguity, called the maxmin expected utility (MMEU) model (Gilboa and Schmeidler 1989). MMEU is a preference-based, axiomatic model; nevertheless, it can be interpreted as a combination of Hurwicz's generalized Bayes–minimax criterion and Wald's minimax criterion.

After 1990, preference-based, axiomatic models capable of accounting for Ellsberg-like choice patterns multiplied. A number of these models were motivated by normative concerns, that is, by the goal of overcoming the notion of economic rationality as defined by Savage's axioms, also known as "Bayesian rationality" or the "Bayesian paradigm".[16] However, it is fair to say that, at the normative level, the Bayesian paradigm has continued to dominate economics until the present day.

10 Conclusion

The research stream on decision-making under ambiguity originated by Ellsberg *QJE* article has been one of the most prolific in economic theory over the last 30 years.[17] In a sense, however, it has become a victim of its own success: currently, a plethora of models for decision-making under ambiguity and a corresponding plethora of preference-based axiomatic systems supporting these models are available. However, as two eminent contributors to this research stream, Gilboa and Massimo Marinacci, noted in 2013, none of these decision models has reached a level of consensus comparable to the consensus that Savage's theory once enjoyed, and "even if a single paradigm will eventually emerge, it is probably too soon to tell which one it will be" (2013, 232). It is fair to say that the state of the art in the theory of decision-making under ambiguity has not changed over the last decade.

Thus, the trajectory of the field remains uncertain. Nevertheless, one thing is clear: Ellsberg's (1961a) classic article ignited a development process in decision theory that is still ongoing.

Acknowledgements I am grateful to Marco LiCalzi and Carlo Zappia for helpful comments on earlier drafts. I owe a special thanks to Massimo Marinacci for several conversations and feedback on the topics addressed here. The usual disclaimer applies. Part of the work for this paper was conducted while visiting the Department of Economics, Ca' Foscari University Venice, and the Laboratoire PHARE, University of Paris 1 Panthéon–Sorbonne; I thank these institutions for their kind hospitality and the Italian Ministry for University and Research, project "Department of Excellence 2023–2027", CUP J37G22000330001, for financial support.

Funding Open access funding provided by Università degli Studi dell'Insubria within the CRUI-CARE Agreement.

Declarations

Conflict of interest The author declares that he has no conflict of interest.

[16] See Gilboa and Marinacci (2013) for a review of this more recent and normatively oriented literature, and Gilboa et al. (2009) and Gilboa (2015) for a critical discussion of the Bayesian paradigm.

[17] Machina (2001), Wakker (2010), Etner et al. (2012), Bühren et al. (2023) provide extensive reviews of this literature.

References

Allais, M.: Le comportement de l'homme rationnel devant le risque: Critique des postulats et axiomes de l'ecole americaine. Econometrica **21**, 503–546 (1953)
Becker, S.W., Brownson, F.O.: What price ambiguity? Or the role of ambiguity in decision-making. J. Polit. Econ. **72**, 62–73 (1964)
Bühren, C., Meier, F., Pleßner, M.: Ambiguity aversion: bibliometric analysis and literature review of the last 60 years. Manag. Rev. Quarter. **73**, 495–525 (2023)
Camerer, C.F., Weber, M.: Recent developments in modelling preferences: uncertainty and ambiguity. J. Risk Uncertain. **5**, 325–370 (1992)
Chipman, J.S.: Stochastic choice and subjective probability. In: Willner, D. (ed.) Decisions, Values, and Groups, pp. 70–95. Pergamon Press, New York (1960)
Davidson, D., Suppes, P., Siegel, S.: Decision Making: An Experimental Approach. Stanford University Press, Stanford (1957)
Econometric Society: Report of the Atlantic City meeting. Econometrica **26**, 314–328 (1958)
Econometric Society: Report of the St. Louis meeting, December 28–30, 1960. Econometrica **29**, 442–482 (1961)
Ellsberg, D.: Classic and current notions of 'measurable utility.' Econ. J. **64**, 528–556 (1954)
Ellsberg, D.: Theory of the reluctant duelist. Am. Econ. Rev. **46**, 909–923 (1956)
Ellsberg, D.: Winning at Russian roulette [abstract]. Econometrica **26**, 325 (1958)
Ellsberg, D.: The art of coercion: a study of threats in economic conflict and war. Lectures at the Lowell Institute, Boston, March 1959. Available at https://www.ellsberg.net. Accessed 2 Aug 2023 (1959)
Ellsberg, D.: Risk, ambiguity, and the Savage axioms. Quart. J. Econ. **75**, 643–669 (1961a)
Ellsberg, D.: The crude analysis of strategy choices. American Economic Review **51**, Papers and Proceedings, 472–478 (1961b)
Ellsberg, D.: Risk, ambiguity, and the Savage axioms: reply. Quart. J. Econ. **77**, 336–342 (1963)
Ellsberg, D.: Risk, Ambiguity and Decision (Ph.D. Thesis, Harvard University, 1962). Routledge, New York (2001)
Ellsberg, D.: Secrets: A Memoir of Vietnam and the Pentagon Papers. Viking, New York (2002)
Ellsberg, D.: Biographical statement. Available at https://www.ellsberg.net. Accessed 2 Aug 2023 (2006)
Ellsberg, D.: Notes on the origins of the Ellsberg urns. Econ. Theor. **48**, 221–227 (2011)
Etner, J., Jeleva, M., Tallon, J.-M.: Decision theory under ambiguity. J. Econ. Surv. **26**, 234–270 (2012)
Feduzi, A.: On the relationship between Keynes's conception of evidential weight and the Ellsberg paradox. J. Econ. Psychol. **28**, 545–565 (2007)
Fellner, W.: Distortion of subjective probabilities as a reaction to uncertainty. Quart. J. Econ. **75**, 670–689 (1961)
Fienberg, S.E.: When did Bayesian inference became 'Bayesian'? Bayesian Anal. **1**, 1–40 (2006)
de Finetti, B.: On the subjective meaning of probability. In: Monari, P., Cocchi, D. (eds.) Induction and Probability, pp. 291–321. Clueb, Bologna ([1931] 1993)
Frenkel-Brunswik, E.: Intolerance of ambiguity as an emotional and perceptual personality variable. J. Pers. **18**, 108–143 (1949)
Gilboa, I.: Rationality and the Bayesian paradigm. J. Econ. Methodol. **22**, 312–334 (2015)
Gilboa, I., Marinacci, M.: Ambiguity and the Bayesian paradigm. In: Acemoglu, D., Arellano, M., Dekel, E. (eds.) Advances in Economics and Econometrics: Theory and Applications, vol. 1, pp. 179–242. Cambridge University Press, New York (2013)

Gilboa, I., Postlewaite, A., Schmeidler, D.: Is it always rational to satisfy Savage's axioms? Econ. Philos. **25**, 285–296 (2009)
Gilboa, I., Schmeidler, D.: Maxmin expected utility with a non-unique prior. J. Math. Econ. **18**, 141–153 (1989)
Good, I.J.: Rational decisions. J. R. Stat. Soc. B **14**, 107–114 (1952)
Good, I.J.: Subjective probability as a measure of a nonmeasurable set. In: Nagel, E., Suppes, P., Tarski, A. (eds.) Logic, Methodology and Philosophy of Science, pp. 319–329. Stanford University Press, Stanford (1962)
Guala, F.: The logic of normative falsification: rationality and experiments in decision theory. J. Econ. Methodol. **7**, 59–93 (2000)
Hamilton, V.: Perceptual and personality dynamics in reactions to ambiguity. Br. J. Psychol. **48**, 200–215 (1957)
Hodges, J.L. Jr., Lehmann, E.L.: The use of previous experience in reaching statistical decisions. Ann. Math. Stat. **23**, 396–407 (1952)
Hurwicz, L.: Optimality criteria for decision making under ignorance. Cowles Commission, Discussion Paper: Statistics No. 370 (1951a)
Hurwicz, L.: Some specification problems and applications to econometric models [abstract]. Econometrica **19**, 343–344 (1951b)
Keynes, J.M.: A Treatise on Probability. Macmillan, London (1921)
Knight, F.H.: Risk, uncertainty, and profit. Houghton Mifflin, Boston (1921)
Koopman, B.O.: The axioms and algebra of intuitive probability. Ann. Math. **41**, 269–292 (1940)
Levi, I.: Introduction. In: Ellsberg, D. Risk, Ambiguity and Decision. Routledge, New York (2001)
Luce, R.D., Raiffa, H.: Games and Decisions. Wiley, New York (1957)
MacCrimmon, K.R.: Descriptive and normative implications of the decision-theory postulates. In: Borch, K., Mossin, J. (eds.) Risk and Uncertainty, pp. 3–32. Macmillan, London (1968)
MacCrimmon, K.R., Larsson, S.: Utility theory: axioms versus 'paradoxes.' In: Allais, M., Hagen, O. (eds.) Expected Utility Hypotheses and the Allais Paradox, pp. 333–409. Reidel, Dordrecht (1979)
Machina, M.J.: Further readings on choice under uncertainty, beliefs and the Ellsberg Paradox. In: Ellsberg, D. Risk, ambiguity and decision. Routledge, New York (2001)
Marinacci, M.: Model uncertainty. J. Eur. Econ. Assoc. **13**, 1022–1100 (2015)
Mongin, P.: The Allais paradox: what it became, what it really was, what it now suggests to us. Econ. Philos. **35**, 423–459 (2019)
Moscati, I.: How economists came to accept expected utility theory: the case of Samuelson and Savage. J. Econ. Persp. **30**, 219–236 (2016)
Moscati, I.: Measuring Utility: From the Marginal Revolution to Behavioral Economics. Oxford University Press, New York (2018)
Mosteller, F., Nogee, P.: An experimental measurement of utility. J. Polit. Econ. **59**, 371–404 (1951)
von Neumann, J., Morgenstern, O.: Theory of Games and Economic Behavior. Princeton University Press, Princeton ([1944] 1953)
Raiffa, H.: Risk, ambiguity, and the Savage axioms: a comment. Quart. J. Econ. **75**, 690–694 (1961)
Ramsey, F.P.: Truth and probability. In: Braithwaite, R.B. (ed.) Foundations of Mathematics and Other Logical Essays, pp. 156–198. Kegan Paul, London (1931)
Roberts, H.: Risk, ambiguity, and the Savage axioms: a comment. Quart. J. Econ. **77**, 327–336 (1963)
Savage, L.J.: The theory of statistical decision. J. Am. Stat. Assoc. **46**, 55–67 (1951)
Savage, L.J.: The Foundations of Statistics. Dover, New York ([1954] 1972)
Schmeidler, D.: Subjective probability and expected utility without additivity. Econometrica **57**, 571–587 (1989)
Slovic, P., Tversky, A.: Who accepts Savage's axiom? Behav. Sci. **19**, 368–373 (1974)
Smith, C.A.B.: Consistency in statistical inference. J. Roy. Stat. Soc. B **23**, 1–25 (1961)
Tversky, A.: Intransitivity of preferences. Psychol. Rev. **76**, 31–48 (1969)
Wakker, P.P.: Prospect Theory. Cambridge University Press, Cambridge (2010)
Wald, A.: Statistical Decision Functions. Wiley, New York (1950)
Zappia, C.: Daniel Ellsberg and the validation of normative propositions. Oeconomia **6**, 57–79 (2016)
Zappia, C.: Rationality under uncertainty: classic and current criticisms of the Bayesian viewpoint. Eur. J. Hist. Econ. Thought **25**, 1387–1419 (2018)
Zappia, C.: Leonard Savage, the Ellsberg paradox and the debate on subjective probabilities: evidence from the archives. J. History Econ. Thought **41**, 1–24 (2021a)

Zappia, C.: From Knightian to Keynesian uncertainty: contextualising Ellsberg's ambiguity. Camb. J. Econ. **45**, 1027–1046 (2021b)