

Between Worlds: Daniel Ellsberg (1931–2023)

IVAN MOSCATI

University of Insubria

CARLO ZAPPIA

University of Siena

Abstract: Daniel Ellsberg was a multifaced personality who belonged to multiple worlds: academia, the military, and, in the second part of his life, political activism. This essay reviews Ellsberg's analysis of decision-making under uncertainty, which has been highly influential in economics and reflects his diverse experiences.

Keywords: Decision theory; uncertainty; ambiguity; Ellsberg paradox

JEL Classification: B21, B31, D81

I. INTRODUCTION

There are people who are remembered for indelibly marking the course of historical events and others who are remembered for turning a field of research on its head. Cases in which the same person can be credited with both merits are very rare. Daniel Ellsberg, who died on the 16th of June 2023 at the age of 92, must be remembered in this dual role.

Born in 1931 into a petit-bourgeois American family—his life marked by a car accident caused by his father in which his mother lost her life—Ellsberg was a brilliant young man who studied economics at Harvard, first as an undergraduate and then as a Junior Fellow of the Harvard Society. A scholar of game theory, he was the first, at the age of 23 years, to clarify the peculiar sense in which John von Neumann and Oskar Morgenstern (1944) had used the utility function, a concept that had been the cross and joy of economists at least since the marginalist revolution of the late 19th century (Ellsberg 1954). But he was critical of game theory, almost obsessed with the difficulty with which uncertainty was treated in that context and in economics in general (Ellsberg 1956).

Although he was considered a brilliant figure among economists who studied games and decisions—Nobel Prize winner Thomas Shelling would say of him that he was the most brilliant person he had ever met—Ellsberg did not pursue an academic career. In the 1950s, interaction between academics and applied research centers was commonplace in the US, and Ellsberg was intrigued by the idea of working at the RAND Corporation—a research center set up to provide scientific support to military decision-makers—rather than in academia.

At RAND, he worked on deterrence issues (Ellsberg 1959) and then specialized in the problems of uncertainty associated with the chain of command for the use of nuclear weapons in the face of a possible threat from the Soviet bloc (Ellsberg 1960). Through these studies, his involvement with the US federal administration increased, until he became one of the RAND experts consulted when decisions were made about developments in the Vietnam War.

It was in this context that he became the protagonist of the episode in American history that goes by the name of the Pentagon Papers, referring to the documents leaked by Ellsberg to *The New York Times* in 1971. For Ellsberg, this was a necessary act of civil disobedience: as an insider, he knew that the government had systematically lied to Congress, withholding evidence of strategic analyses that revealed the futility of the choice to continue the war on the ground.

Since then—having escaped the prospect of life imprisonment because the judge hearing his case was made aware that the government had tried to manipulate the evidence and then dismissed the charges (Ellsberg 2002)—he has been an independent analyst and activist. His relevant impact outside academia and in the public debate is evident both in the attention he received from outlets such as *The Guardian*, *The New York Times*, and *The New Yorker*, among others, and in a vast literature on the history of nuclear war planning (Kaplan 1991; Rosenbaum 2011). Several times in recent years he appeared in the international press to support the initiatives of other whistle-blowers such as Julian Assange, Chelsea Manning, and Eric Snowden, and to defend the role of conscientious objection.¹

In economics, however, his name is mainly associated with his analysis of decision-making under uncertainty and, more specifically, with the

¹ The 2009 Academy Award nominated documentary “The most dangerous man in America. Daniel Ellsberg and the Pentagon Papers” is arguably the best source for the story of the quintessential whistle-blower.

so-called Ellsberg Paradox (Ellsberg 1961a). This is a choice situation, designed by Ellsberg, in which decision makers had to express their preferences between gambles with uncertain outcomes. Ellsberg observed that several reasonable decision makers expressed deliberate preferences that violated the then-dominant theory of decision-making under uncertainty, namely the version of expected utility theory advanced some years earlier by Leonard J. Savage (1954). Over the last forty years, while Savage's version has remained the mainstream view, several decision models capable of accounting for choice patterns like those recorded by Ellsberg have been proposed, so making the Ellsberg Paradox the starting point for one of the most thriving research programs in decision theory.

Between 1961 and April 1962, Ellsberg worked on his Harvard PhD thesis, which largely built upon his 1961 article, albeit with some significant novelties. In particular, in writing the thesis Ellsberg realized that with his examples he has re-proposed a question that had already been addressed by John Maynard Keynes in his *Treatise on Probability* (1921), and that Keynes's approach could contribute to a much wider analysis of uncertainty than that associated with the urn examples. Unfortunately, Ellsberg did not develop the novel insights contained in his dissertation because, after completing it, he became a full-time military analyst for the Defense and State Departments. Ellsberg did not publish his dissertation until many years later, without any changes (Ellsberg 2001). Ellsberg (2011) provides a brief retrospective view on the development of his ideas.

The remainder of this essay is structured as follows. Section II examines the Ellsberg Paradox and its motivations. Section III considers Ellsberg's PhD thesis, and discusses the major novel insights contained in it. Section IV overviews the literature originated by Ellsberg's work. Section V concludes.²

II. THE PARADOX

The article presenting the Ellsberg Paradox should 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 von Neumann and Morgenstern (1944). According to a classification introduced by Frank Knight (1921), in situations of 'risk', the decision

² This essay is based on previous works of the authors: Basili and Zappia (2010), Zappia (2016, 2021) and Moscati (2023, 2024).

maker knows the probabilities of the payoff relevant events, whereas in situations of ‘uncertainty’ she does not.

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 18th century: expected utility theory (EU). They showed that 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 $p(x_i)$.

In his *Foundations of Statistics* (1954), Savage extended EU to situations of uncertainty. He demonstrated that if the decision maker’s preferences between courses of actions with uncertain outcomes satisfy certain postulates, designated as P1–P7, it becomes possible to (1) identify a unique probability measure $\pi(E_i)$ defined over the set of uncertain events E 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 the expected utility $\sum u(x_i)\pi(E_i)$, is calculated by employing the subjective probabilities $\pi(E_i)$ identified at the first step. Because it involves subjective probabilities, to be updated following the Bayes rule as new information becomes available, Savage’s theory is often referred to, also by Ellsberg, as ‘Bayesian decision theory’.

Two aspects of Savage’s version of EU should be recalled here. First, in his axiomatization, a central role is played by postulate P2, which is also called the Sure-Thing principle. P2 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 courses of action 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 .

Second, Savage advocated a normative interpretation of his seven axioms, considering them as maxims of rational behavior rather than descriptions of actual behavior. According to Savage, the normative nature of the axioms 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.

Ellsberg admired Savage’s theoretical construction, embraced Savage’s normative interpretation of the axioms, but doubted that it is always possible to reduce uncertainty to risk by deriving subjective probabilities from preferences between courses of actions. Thus, in the late 1950s Ellsberg began looking for uncertain situations that resist reduction to risk and in which, therefore, one or more of Savage’s axioms are violated and EU does not apply. The Ellsberg Paradox, featured in an article titled “Risk, Ambiguity, and the Savage Axioms” and published in the *Quarterly Journal of Economics* in 1961, provides a compelling example of such a choice situation (Ellsberg 1961a).

Ellsberg imagined a decision maker who is presented with an 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 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. 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.

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 P2, and that this violation makes it impossible to infer probabilities from the decision maker’s preferences. In fact, 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’.³

³ In more detail, the decision maker’s preference for a_1 over a_2 indicates that for her $p(\text{red}) > p(\text{black})$. However, her preference for a_4 over a_3 indicates that she deems $p(\text{black or yellow}) > p(\text{red or yellow})$. Given that the urn only contains red, black, and yellow balls, the additivity principle of probability requires that $p(\text{black or yellow}) + p(\text{red}) = 1$ and $p(\text{red or yellow}) + p(\text{black}) = 1$. Hence, $p(\text{black or yellow}) > p(\text{red or yellow})$ implies that $1 - p(\text{red}) > 1 - p(\text{black})$ or, what was to be shown, $p(\text{red}) < p(\text{black})$.

Importantly, 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.

As mentioned above, Ellsberg embraced Savage's normative interpretation of postulates P1-P7 and therefore of EU. From a normative point of view, the fact that most subjects violate the axioms is not sufficient to reject them: maybe the violations are due to careless or random choices, and a reasonable decision maker, upon reflection and deliberation, is willing to correct her choices so that they conform to the axioms. Therefore, Ellsberg was looking for 'deliberate violators', that is, for ostensibly reasonable decision makers who, after having well understood the meaning of Savage's axioms and the fact that their preferences violated them, consciously decide to uphold their preferences.

He proposed the choice situation involving the urn to the participants to seminars at Harvard, Chicago, Yale, RAND, and other research institutions where he presented his paper between 1959 and 1960. These participants included some of the major decision theorists of the period, that is, people who could be safely assumed to be reasonable: Paul Samuelson, Jacob Marschak, Howard Raiffa, Gérard Debreu, Robert Schlaifer, Norman Dalkey, and, last but not least, Savage himself.

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 corrected their choices. Yet others violated the axioms and maintained their choices. He reported that the group of deliberate violators included, besides Ellsberg himself, Marschak, Dalkey, and even Savage.

For Ellsberg, the deliberate choices of these individuals show that there exists a class of situations of uncertainty—which he called situations of "ambiguity"—in which Savage's theory loses its descriptive and normative validity. In the last part of his 1961 article, Ellsberg advanced an alternative decision model that can account for choices in conditions of ambiguity. However, Ellsberg's model did not catch up in decision theory, probably because it lacks an axiomatic foundation in terms of preferences.

III. BEYOND THE PARADOX: UNCERTAINTY IN THE REAL WORLD

As noted in the introduction, after graduating Ellsberg seemed destined for an academic career. Perhaps with an eye to a return to academia that never materialized, although he was already a permanent researcher at

RAND, he continued to work on the issues raised in the 1961 article to obtain a PhD from Harvard. It can be argued that his PhD thesis is a testimony to what Ellsberg might have done had he continued to study decision theory.

First, the thesis outlines the characteristics of a normative approach to the decision problem that generalizes Savage’s Bayesian approach and could be therefore described as ‘generalized Bayesian’ (Levi 2001). Ellsberg endorsed Savage’s subjectivist approach—which crowned the probabilistic revolution independently launched by Frank Ramsey ([1926] 1931) and Bruno de Finetti’s ([1931] 1993)—but found a strict Bayesian view too extreme. In his view, the ability to identify precise *a priori* probabilities could indeed be extended beyond decisions in risky contexts, as Savage had shown, but this possibility could not be taken for granted in all contexts. Ellsberg emphasized that, as in Knight (1921), the claim to reduce all uncertainties to risks is unjustified at the normative level, regardless of any descriptive violations. To defend his position, he referred in the thesis not only to Knight but also to Keynes’s *Treatise on Probability* (1921). He credited Keynes, on the one hand, for anticipating the notion of ambiguity through his considerations on the weight of arguments and, on the other hand, for not being afraid to explore the possibility of devising a theory of imprecise probabilities. Ellsberg’s generalized Bayesian approach—which he associated in particular to the mathematicians and statisticians who continued Keynes’s work, such as Bernard Koopman (1940), Irving Good (1952), and Cedric Smith (1961)—is “a theoretical approach that admits vagueness as an explicit factor without apology and provides a formal vocabulary for discussing it” (Ellsberg 2001, 10).

Ellsberg acknowledged that the introduction of subjective probabilities expands the field of rational behavior to include the influence on economic behavior of individual circumstances, personality, and information. For him, however, Savage’s Bayesian decision theory unjustifiably abstracts from the possibility that there may be decision contexts where opinions are not well defined. And since undefined opinions can be represented by probability intervals, although a complete ordering of events in terms of probabilities is no longer possible because the intervals may overlap, one must start from a theory in which the ordering is partial. Hence the conclusion that “the attractiveness of such a system of partially ordered intervals as a mathematical model for a subjective system of de-

degrees of belief some of which are 'indefinite' (as opposed to a model consisting of the real numbers between 0 and 1) may be immediately evident" (Ellsberg 2001, 72; italics in the original).

On this basis, the thesis investigates decision rules to be applied when the domain of subjective probabilities is made more realistic to a much greater extent than in the 1961 article. Among other things, an optimism parameter *à la* Hurwicz (1951) is considered, together with a parameter indicating the extent to which a given probability measure is considered reliable. As noted above, Ellsberg was unable to provide a proper axiomatic structure to justify the alternative decision rules he advocated, but it is noteworthy that both the issue of how to multiply utilities and probabilities when these are non-additive probabilistic weights and the issue of how to consider multiple probability priors are identified by Ellsberg as crucial. Indeed, his most general suggested decision rule, which he called 'restricted Bayes-Hurwicz', bears a striking resemblance to the recent attempts to move on from Savage's decision rules (see section IV).

There is a methodological argument worth noting here. Ellsberg's aim was to reject Savage's theory *and* to provide heuristic advice on how to develop a new normative theory that is arguably less likely to be confronted with deliberate violation when tested, and therefore more general. Ellsberg was aware that while a counterexample is a necessary condition for falsifying of a normative theory, it is not also sufficient for its definitive rejection (Guala 2000). Ellsberg's conjecture about the origin of the anomaly was therefore intended to offer a sketch of a new, rival normative theory effectively capable of accounting for the phenomenon of uncertainty.

This leads us to a second aspect of Ellsberg's dissertation that deserves to be highlighted: in it, the need to deal with situations of 'extreme ambiguity' appears with accentuated insistence. The axiomatic structure proposed by Savage, Ellsberg (2001, 28) admitted, can be considered "eminently reasonable" in many decision-making situations; but it is not necessarily "*uniquely* reasonable", that is, rational and applicable to every possible decision context. For Ellsberg, the weakness of Savage's approach remained that he had not looked beyond "small words", that is, situations where the list of possible states of the world can be considered as exhaustive. In doing so, Savage had neglected the relevance of crucial features that often characterize the real world.

Ellsberg emphasized that this is not the case in a whole range of areas, and in particular for military research and development, where the level

of investments is huge and the commitment to pursue national objectives extends into the distant future. Here “uncertainty is the central fact: uncertainty about the future environment, about the possible performance of alternative programs launched today, even about the operational national objectives in that environment and time, and hence about the future evaluation of performance” (Ellsberg 2001, 16). The models that are useful for evaluating the costs and benefits of alternative programs with respect to different objectives, to which Savage had given the much needed axiomatic structure, are not sufficiently developed to deal with this kind of uncertainty.

Strangely enough, after a careful analysis of Keynes’s *Treatise*, Ellsberg did not dwell on the *General Theory* (1936) or the 1937 *Quarterly Journal of Economics* article in which Keynes (1937) puts uncertainty at the heart of his critique of the ‘Benthamite calculus’. But it is significant that Ellsberg referred to the even more extreme position of George Shackle (1955). Shackle’s “somber reflections” on the ineffectiveness of any probabilistic consideration in the face of the uncertainty of the real world, which had not been valued in the 1961 article, are now considered “too ominously relevant [...] to be dismissed” (Ellsberg 2001, 17).

The research Ellsberg was doing at the same time as an analyst at RAND must have had an impact on what Ellsberg considered relevant for decision-making in real world contexts. Ellsberg’s perspective at RAND is reflected in a series of studies on conflict and bargaining, including an analysis of the deterrence policy a superpower can enforce to prevent ‘attack’ from becoming a dominant strategy for the opponent (Ellsberg 1959). The approach was not to use abstract game-theoretic notions, and he focused more on decision processes than on equilibrium outcomes, much like Schelling (1960). The informational conundrum that military decision-makers are plunged into when acting in a real strategic context analyzed in one of these papers (Ellsberg 1961b) is quoted by Ellsberg in the thesis as an obvious example of uncertainty that cannot be dealt with as if risk.

But this is even more evident when considering documents that were classified at the time, either because they were intended for RAND use only or because they were provided directly to the US Defense Department (Ellsberg 2017). When asked to report on the ability of a chain of command to cope with the event of a nuclear attack by US adversaries, Ellsberg stressed that the typical post-attack situation would consist of an unprecedented kind of war, and that there was strong reason to expect

that “important aspects of it will not have been anticipated, nor appropriate responses pre-planned” (Ellsberg 1960, 33). In his advice as a military analyst, he suggested a strategy based on adaptive behavior because the general problem to be faced would surely be that of dealing with “information, events, and evaluations which were unanticipated in the programming phase and for which no plans exist”. The real decision-making environment he described was one in which a list of “potential surprises” was to be recorded (Ellsberg 1960, 47-49). The issue of how to adapt the US strategic response to “circumstances which might be unforeseen, ambiguous or both”, is also examined in some unsigned notes dated May 1961. Here Ellsberg (1961c, 16) concluded that since “the list of possibilities is not exhaustive” surprises in the form of “wholly unforeseen circumstances are likely”.

In essence, Ellsberg did not use the event-uncertainty framework that his urn examples share with Savage’s small world. Instead, based on his experience as a military consultant, he was led to emphasize the state-space uncertainty associated with unprecedented situations.

IV. INFLUENCE

As already mentioned, in economics Ellsberg’s name is mainly associated with the decision paradox featured in his *Quarterly Journal of Economics* article (1961a). This article has become a seminal contribution in the field of decision theory and garnered over 11,000 Google Scholar citations. 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 ambiguity.

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 violating EU were systematic 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.

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 axioms, 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 probability weights that do not satisfy the additivity property 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 the Choquet expected utility (CEU) model. To have an idea about how capacities $c(E)$ work, consider a patient facing the decision of whether to undergo surgery, which can succeed (event E_1) or fail (event E_2). In Schmeidler's model, the patient may believe that the surgery succeeds 8 times out of 10 and fails 1 time out of 10; in this case, $c(E_1) = 0.8$, $c(E_2) = 0.1$, and $c(E_1) + c(E_2) = 0.9$.

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). In this model, the decision maker's beliefs in conditions of ambiguity are captured by assuming that she considers possible a set of probability distributions π_1, \dots, π_Z over the set of events. Returning to our medical example, we can imagine that the patient consults two different doctors. According to doctor 1, the surgery goes well 9 times out of 10, while according to doctor 2, it goes well only 5 times out of 10. In this ambiguous situation, the patient may consider two possible probability distributions, π_1 and π_2 , where $\pi_1(E_1) = 0.9$ and $\pi_2(E_1) = 0.5$.

After 1990, preference-based, axiomatic models capable of accounting for the Ellsberg paradox and, more generally, for decision-making under ambiguity multiplied, often following either the capacity approach used in the CEU model, or the multi-probability approach featured in the MMEU model. Ellsberg is also mentioned in a growing literature in philosophical decision theory that exploits the normative content of his argument with the aim of providing a normative theory of rationality tailored to agents who do not obey to the standard requirements of Bayesian rationality (see in particular Bradley 2017). It is fair to say that the research stream originated by the Ellsberg Paradox has been one of the most prolific in economic theory over the last 30 years.⁴

⁴ Gilboa (2009, 2024), Wakker (2010), Etner, Jeleva, and Tallon (2012), Gilboa and Marinacci (2013), and Machina and Siniscalchi (2014) provide extensive reviews of this literature.

V. CONCLUSION

As we have tried to illustrate in this essay, Ellsberg's multifaceted analysis of decision-making under uncertainty was influenced by the diverse contexts he was involved in.

First, of course, there was the academic context. We have seen that his seminal 1961 article launched a theoretical and experimental literature that uses the notion of ambiguity to study certain kinds of uncertainty that are not reducible to risks.

Secondly, there was Ellsberg's involvement with applied research for the military, which prompted him to also consider those aspects of decision-making associated with situations where the list of possible states of the world cannot be considered exhaustive.

In particular, his PhD dissertation shows that his strategy of using urn examples to criticize Savage and introduce the notion of ambiguity did not mean that he was uninterested in a more general kind of uncertainty than that usually subsumed under the heading of ambiguity in the current literature.

REFERENCES

- Basili, Marcello, and Carlo Zappia. 2010. "Ambiguity and uncertainty in Ellsberg and Shackle." *Cambridge Journal of Economics* 34 (3): 449–474
- Becker, Selwyn W., and Fred O. Brownson. 1964. "What price ambiguity? Or the role of ambiguity in decision-making." *Journal of Political Economy* 72 (1): 62–73.
- Bradley, Richard. 2017. *Decision Theory with a Human Face*. Cambridge: Cambridge University Press.
- de Finetti, Bruno. [1931] 1993. "On the subjective meaning of probability." In *Induction and Probability*, edited by Paola Monari and Daniela Cocchi, 291–321. Bologna: Clueb.
- Ellsberg, Daniel. 1954. "Classic and current notions of 'measurable utility'." *Economic Journal* 64 (255): 528–556.
- Ellsberg, Daniel. 1956. "Theory of the reluctant duelist." *American Economic Review* 46 (5): 909–923.
- Ellsberg, Daniel. 1959. "The theory and practice of blackmail." RAND Corporation, mimeo. Available at <https://www.rand.org/pubs/papers/P3883.html>.
- Ellsberg, Daniel. 1960. "Strategic objectives and command and control problems." RAND Corporation, mimeo. Available at <https://archive.org/details/StrategicObjectivesAndCommandControlProblems/mode/2up>.
- Ellsberg, Daniel. 1961a. "Risk, ambiguity, and the Savage axioms." *Quarterly Journal of Economics* 75 (4): 643–669.
- Ellsberg, Daniel. 1961b. "The crude analysis of strategy choices." *American Economic Review* 51 (2): 472–478.

- Ellsberg, Daniel. 1961c. "Basic national security policy: Central War section." Unsigned note for the U.S. Department of Defense. Available at <https://archive.org/details/BNSPDraftCentralWarSection/page/n1/mode/2up>.
- Ellsberg, Daniel. 2001. *Risk, Ambiguity and Decision* (Ph.D. Thesis, Harvard University, 1962). New York: Routledge.
- Ellsberg, Daniel. 2002. *Secrets: A Memoir of Vietnam and the Pentagon Papers*. New York: Viking.
- Ellsberg, Daniel. 2011. "Introduction to the symposium issue: notes on the origins of the Ellsberg urns." *Economic Theory*, 48 (2): 221–227.
- Ellsberg, Daniel. 2017. *The Doomsday Machine: Confessions of a Nuclear War Planner*. New York: Bloomsbury Publishing.
- Etner, Johanna, Meglena Jeleva, and Jean-Marc Tallon, J.-M. 2012. "Decision theory under ambiguity." *Journal of Economic Surveys* 26 (2): 234–270.
- Gilboa, Itzhak. 2009. *Theory of Decision under Uncertainty*. Cambridge: Cambridge University Press.
- Gilboa, Itzhak. 2024. "Decision under Uncertainty—State of the Science." *Annual Reviews in Economics* 16, forthcoming.
- Gilboa, Itzhak, and Massimo Marinacci. 2013. "Ambiguity and the Bayesian paradigm." In *Advances in Economics and Econometrics: Theory and Applications*, edited by Daron Acemoglu, Manuel Arellano, and Eddie Dekel, Vol. 1, 179–242. New York: Cambridge University Press.
- Gilboa, Itzhak, and David Schmeidler. 1989. "Maxmin expected utility with a non-unique prior." *Journal of Mathematical Economics* 18 (2): 141–153.
- Good, Irving J. 1952. "Rational decisions." *Journal of the Royal Statistical Society, Series B* 14 (1): 107–114.
- Guala, Francesco. 2000. "The logic of normative falsification: rationality and experiments in decision theory." *Journal of Economic Methodology* 7 (1): 59–93.
- Hurwicz, Leonid. 1951. "Optimality criteria for decision making under ignorance." Cowles Commission, Discussion Paper: Statistics No. 370.
- Kaplan, Fred. 1991. *The Wizards of Armageddon*. Stanford: Stanford University Press.
- Keynes, John Maynard. 1921. *A Treatise on Probability*. London: Macmillan.
- Keynes, John Maynard. 1936. *The General Theory of Employment, Interest and Money*. London: Macmillan.
- Keynes, John Maynard. 1937. "The general theory of employment." *Quarterly Journal of Economics* 51 (2): 209–223.
- Knight, Frank H. 1921. *Risk, Uncertainty, and Profit*. Boston: Houghton Mifflin.
- Koopman, Bernard O. 1940. "The axioms and algebra of intuitive probability." *Annals of Mathematics* 41 (2): 269–292.
- Levi, Isaac. 2001. "Introduction." In D. Ellsberg, *Risk, Ambiguity and Decision*, ix–xxxvii. New York: Routledge.
- MacCrimmon, Kenneth R. 1968. "Descriptive and normative implications of the decision-theory postulates." In *Risk and Uncertainty*, edited by Karl Borch and Jan Mossin, 3–32. London: Macmillan.
- MacCrimmon, Kenneth R., and Stig Larsson. 1979. "Utility theory: axioms versus 'paradoxes'." In *Expected Utility Hypotheses and the Allais Paradox*, edited by Maurice Allais and Ole Hagen, 333–409. Dordrecht: Reidel.

- Machina, Mark J., and Marciano Siniscalchi. 2014. "Ambiguity and ambiguity aversion." In *Handbook of the Economics of Risk and Uncertainty*, edited by Mark J. Machina and W. Kip Viscusi, Vol. 1, 729–807. Amsterdam: North-Holland.
- Moscatti, Ivan. 2023. *The History and Methodology of Expected Utility*. Cambridge: Cambridge University Press.
- Moscatti, Ivan, 2024. "Ellsberg 1961: text, context, influence." *Decisions in Economics and Finance*, forthcoming.
- Ramsey, Frank P. [1926] 1931. "Truth and probability." In Frank P. Ramsey, *Foundations of Mathematics and Other Logical Essays*, edited by Richard B. Braithwaite, 156–198. London: Kegan Paul.
- Rosenbaum, Ron. 2011. *How the End Begins. The Road to a Nuclear World War III*. New York: Simon and Schuster.
- Savage, Leonard J. 1954. *The Foundations of Statistics*. New York: Wiley.
- Schelling, Thomas. 1960. *The Strategy of Conflict*. Cambridge, MA: Harvard University Press.
- Schmeidler, David. 1989. "Subjective probability and expected utility without additivity." *Econometrica* 57 (3): 571–587.
- Shackle, George L.S. 1955. *Uncertainty in Economics*. Cambridge: Cambridge University Press.
- Slovic, Paul, and Amos Tversky. 1974. "Who accepts Savage's axiom?" *Behavioral Science* 19 (6): 368–373.
- Smith, Cedric A.B. 1961. "Consistency in statistical inference." *Journal of the Royal Statistical Society, Series B* 23 (1): 1–25.
- von Neumann, John, and Oskar Morgenstern. 1944. *Theory of Games and Economic Behavior*. Princeton: Princeton University Press.
- Wakker, Peter. 2010. *Prospect Theory*. Cambridge: Cambridge University Press.
- Zappia, Carlo. 2016. "Daniel Ellsberg and the validation of normative propositions." *Æconomia* 6 (1): 57–79.
- Zappia, Carlo. 2021. "From Knightian to Keynesian uncertainty: contextualising Ellsberg's ambiguity." *Cambridge Journal of Economics* 45 (5): 1027–1046.

Ivan Moscatti is Professor of Economics at the University of Insubria. His research focuses on the history and philosophy of decision theory.
Contact e-mail: <ivan.moscatti@uninsubria.it>

Carlo Zappia is Associate Professor of Economics at the University of Siena. His research focuses on the history of uncertainty in economics.
Contact e-mail: <zappia@unisi.it>