1 Introduction

Expected utility theory is a model for the analysis of individual decision-making in situations of risk or uncertainty. These are situations where the outcome of a course of action depends on whether some event occurs. For instance, the outcome of the decision to buy a lottery ticket depends on whether the purchased ticket is drawn, and the outcome of the decision to undergo surgery depends on whether the operation goes well.

In order to state expected utility theory (henceforth EU) with sufficient precision, we need some notation. Assume that N uncertain events E_i are possible, with $i = 1, \ldots, N$, and that they are mutually exclusive (no more than one event can occur at the same time) and jointly exhaustive (at least one of the events must occur). A course of actions yielding outcome x_1 if event E_1 occurs, outcome x_2 if event E_2 occurs, and so on can be represented as $[x_1, E_1;\ x_2, E_2; \ldots;\ x_N, E_N]$. For instance, the lottery-ticket situation can be modeled by stating that there are two possible events: "the purchased ticket is drawn" (event E_1) and "the purchased ticket is not drawn" (event E_2). If event E_1 occurs, the decision maker wins a vacation in Greece and is even reimbursed for the lottery ticket (outcome x_1); if event E_2 occurs, the decision maker loses the five euros he paid for the ticket (outcome x_2).

Assume also that, for each event E_i, it is possible to identify the probability $p(E_i)$ that the event occurs. Such probability can be an "objective" probability that the decision maker knows. For instance, if there are 10,000 lottery tickets, the objective probability that the purchased ticket is drawn is 1 in 10,000, or 0.0001. In this case, decision theorists talk of decision-making "under risk." If objective probabilities are not available or are not known by the decision maker, decision theorists talk of decision-making "under uncertainty." In this case, $p(E_i)$ can be seen as a "subjective" probability expressing the decision maker's degree of belief that event E_i will occur. For instance, if the decision maker believes that the surgery goes well nine times out of ten, the subjective probability $p(E_i)$ is 0.9.[1]

Finally, suppose that there exists a real-valued function $u(\cdot)$ that assigns a number $u(x_i)$ to each outcome x_i. For instance, $u(\textit{holiday in Greece}) = 60{,}000$ and $u(-5\ \textit{Euros}) = -7$. This function is called a "utility function" and, as we will discuss, the meaning of the numbers $u(x_i)$ is a controversial issue.

[1] From a philosophical perspective, and in particular in the light of the enormous literature generated by Lewis's ([1980] 1987) paper on chance and credence, the distinction between "objective" and "subjective" probability used in decision theory may appear simplistic. However, as will become clear in the following sections, it is sufficient for the purposes of this Element.

For the moment, however, it suffices to say that EU states that the value a decision maker assigns to a course of action $[x_1, E_1; \ldots; \ x_N, E_N]$ is expressed by its "expected utility" – in other words, by the average of the utility values $u(x_i)$, each weighted by its probability $p(E_i)$; that is, the value of a course of action is expressed by $\sum_{i=1}^{N} u(x_i)p(E_i)$. In order to lighten the notation, when the indexing of variables is not needed, I shall write this formula as just $\sum u(x_i)p(E_i)$.

In our example, the expected utility of the course of action "buy the lottery ticket" is $60{,}000 \times \frac{1}{10{,}000} + (-7) \times \frac{9{,}999}{10{,}000}$, which is around -1.

1.1 The Two Faces of EU: Normative and Descriptive

Expected utility theory has a double character: It is both "normative" and "descriptive." As a normative theory, EU states what a sensible or, as economists and philosophers tend to say, a "rational" decision maker ought to do: select the course of action associated with the highest expected utility. For instance, if the expected utility attached to the course of action "don't buy the lottery ticket" is higher than the expected utility attached to the course of action "buy the ticket," the decision maker ought not buy the ticket. Intended as a descriptive theory, EU aims at describing what actual decision makers do, even if their behavior might not appear rational.

In principle, the two faces of EU are disconnected, in the sense that EU could be normatively valid but descriptively invalid. Thus, the decision maker may agree that it is not sensible for him to buy the lottery ticket but then, for some reason, he might buy it anyway.[2] However, as we will see, the history of EU shows that the normative and descriptive dimensions of the theory are strictly interrelated.

1.2 EU in Economics and Philosophy

EU for decision-making under risk was originally advanced by Daniel Bernoulli ([1738] 1982) in the eighteenth century but entered economics much later, in the 1870s. Since then, EU has undergone changing fortunes in the discipline. Between the 1870s and 1910s, most economists accepted it, although with some reservations. In the 1920s and 1930s, in the context of the so-called ordinal turn in utility analysis, further criticisms against the theory were raised, and by the early 1940s the supporters of EU in economics were few. The fortunes of EU began to recover in the mid-1940s, when John von Neumann and Oskar

[2] In order to maintain gender equilibrium in the use of third-person singular personal pronouns without impairing the readability of the text, I shall use masculine pronouns in odd sections and feminine pronouns in even sections.

Morgenstern ([1944] 1953) advanced a novel, preference-based version of the theory.

In the von Neumann–Morgenstern version, and with the extension to decision-making under uncertainty provided by Leonard Jimmie Savage ([1954] 1972), EU became the dominant economic model of individual decision-making under risk and uncertainty, a position that it retained at least until the 1990s.

Beginning in the 1970s, the accumulation of robust experimental evidence against EU prompted decision theorists to advance a number of models alternative to EU, such as prospect theory (Kahneman & Tversky 1979; Tversky & Kahneman 1992). However, none of these alternative models has yet achieved the level of consensus that EU once enjoyed. For this reason, and also thanks to its simplicity and adaptability, EU remains the primary model in numerous areas of economics dealing with decisions under risk or uncertainty, such as finance, the theory of asymmetric information, and game theory.

Expected utility theory has also played an important role in philosophy. Since the mid-1950s, EU has been used as a normative theory of rational choice (Davidson, McKinsey, & Suppes 1955). With some simplification, we may say that in philosophy EU has maintained this normative status until the present. Philosophers have paid relatively little attention to the non-EU models advanced in the last forty years or so, and this is primarily because they do not perceive them as normatively valid (for a discussion, see Buchak 2013; Okasha 2016; Bradley 2017).

In the 1970s, David Lewis (1974) and other philosophers began to attach a further meaning to EU, namely that it offers a formalized version of the common-sense, or folk-psychological, understanding of decision-making. According to this interpretation, common sense suggests that our decisions result from the combination of our desires and beliefs. Expected utility theory would capture desires through the utility function $u(x)$ and beliefs through the probability function $p(E)$, and it would indicate a simple way to combine them, via the formula $\sum u(x_i)p(E_i)$. However, as we will discuss, this interpretation is debatable. In particular, it is controversial whether the utility function $u(x)$ can be actually interpreted as capturing desires and whether the summations and multiplications needed to calculate the value $\sum u(x_i)p(E_i)$ are really simple.

1.3 This Element

In this Element, I offer an accessible but technically detailed review of EU. My approach falls between the history of ideas and economic methodology. At the historical level, I review EU by following its conceptual evolution from its original formulation in the eighteenth century through its transformations and extensions in the mid-twentieth century to more recent supersession by post-EU theories such as prospect theory. In reconstructing the history of EU, I focus on the methodological issues that have accompanied its evolution: Is EU descriptively and/or normatively valid? Can the utility function featured in EU be interpreted as expressing how intensely a decision maker desires an outcome? More generally, do the utility function and the other components of EU correspond to entities that actually exist in the minds of decision makers, or are they best understood as fictional constructs that may lack any mental reference? On many of these issues, no consensus has yet been reached, and in this Element I offer my view on them.

In Section 2, I first reconstruct how EU originated in the discussions that some mathematicians of the late seventeenth and early eighteenth centuries had about the likelihood of certain aleatory events occurring in betting; I then illustrate Bernoulli's version of EU and discuss its explanatory structure from a methodological viewpoint. In Section 3, I explain why most economists accepted Bernoulli's EU between the 1870s and 1910s and how the "ordinal turn" in utility analysis, which was completed by the late 1930s, led to a generalized rejection of the theory. Section 4 presents von Neumann and Morgenstern's novel, preference-based version of EU. In Section 5, I discuss a number of theoretical and methodological issues related to von Neumann and Morgenstern's EU, such as its relationship with Bernoulli's EU, its descriptive and normative validity, including a discussion of the "Allais paradox," and the appropriate interpretation of the utility function $\widetilde{u}(x)$ featured in von Neumann and Morgenstern's EU. In Section 6, I move to Savage's extension of EU to situations of uncertainty; among other things, I discuss the descriptive and normative validity of Savage's EU, including an illustration of the "Ellsberg paradox," and the theoretical status of the utility function $\widetilde{u}(x)$ and the probability measure $p(E)$ featured in Savage's theory. Finally, Section 7 presents a quick overview of the theories that go beyond EU and have been advanced since the mid-1970s. I am convinced that certain features of EU become fully clear only when it is contrasted with theories alternative to it. In particular, I focus on prospect theory, which arguably is the most influential of the post-EU theories, and compare the ways in which EU and prospect theory model risk attitudes; at the methodological level, I argue that, contrary to what its advocates typically claim, prospect theory is not psychologically more realistic than EU and, like EU, is best understood as an "as-if" model of decision-making.

A few final remarks about the scope and features of the present Element are in order. First, although technically detailed, my review of EU attempts to be mathematically accessible to any student of economics and philosophy. Readers seeking a mathematically more advanced presentation of EU may consult other works, such as Kreps (1988) or Gilboa (2009). Second, the overview of the post-EU theories I offer in Section 7 is extremely limited. Several comprehensive works reviewing this literature are available: Schoemaker (1982), Starmer (2000), Wakker (2010), Gilboa and Marinacci (2013), and Lipman and Pesendorfer (2013) focus on the economic literature, while Peterson (2017) and Steele and Stefánsson (2020) concentrate on the philosophical literature. Finally, some parts of this Element draw on related work of mine: The more historical parts draw on Moscati (2016) and Moscati (2018), and the more methodological sections are based on Moscati (in press).

2 Bernoulli's EU

2.1 A New Field of Study

The scientific disciplines that we today call probability theory and decision theory came into being as a single field of study around 1650, when the French mathematicians Blaise Pascal and Pierre de Fermat, as well as the Dutch mathematician and astronomer Christiaan Huygens, used a principled approach to solve issues related to gaming and betting that had been occasionally discussed since the late fifteenth century. These issues deal with the likelihood of certain aleatory events related to dice throwing or coin tossing, the fair price to pay in order to participate in such events, and the so-called problem of points, which concerns the equitable division of a monetary prize in a game of chance interrupted before completion.

The new field of study displayed contrasting dimensions that have accompanied its evolution, and later the evolution of probability theory and decision theory, until the present. First, the idea of probability was dual from its very emergence. As the historian Ian Hacking (1975, 1) stressed, the probability notion was originally connected, on the one hand, "with the degree of belief warranted by evidence" and, on the other, "with the tendency, displayed by some chance devices, to produce stable relative frequencies."

The analysis of decisions was also dual from the start, as it was both normative and descriptive. For early decision theorists, the normative dimension was the dominant one, and their analysis mostly concerned the maximum price that people of good sense should pay to participate in certain games of chance, or the equitable division of money in the problem of points. However, early decision theorists also aimed at describing what people actually do and therefore tested the normative recommendations obtained by mathematical reasoning against introspective psychological evidence, or against the ordinary behavior of competent individuals, such as merchants or skilled players of games of chance.

2.2 Expected Payoff

One common tenet of the theories put forward by Pascal, Fermat, and Huygens was the identification of what today we call "mathematical expectation," "expected value," or "expected payoff" as the parameter indicating the fair price of a monetary gamble. Using the notation introduced in Section 1, the expected payoff of a course of action $[x_1, E_1; \ldots;\ x_N, E_N]$ that yields monetary payoff x_i if event E_i occurs is given by $\sum x_i p(E_i)$. For brevity, I call the "expected-payoff hypothesis" the tenet that the fair price of a monetary gamble corresponds to its expected payoff.

The most explicit supporter of this hypothesis was Huygens, whose book *De ratiociniis in ludo aleae* (*On Calculation in Games of Chance*, 1657) was the first published treatise in the new field of study. To illustrate the hypothesis, Huygens considered a game in which an individual hides three shillings in one hand and seven shillings in the other, while another individual chooses one of the hands. For Huygens, to the latter individual this gamble is worth $\frac{1}{2} \times 3 + \frac{1}{2} \times 7 = 5$ shillings. More generally, Huygens stated that "if I have the same chance to get *a* or *b* [where *a* and *b* are amounts of money], the game is worth to me as much as $\frac{a+b}{2}$" (Huygens [1657] 1920, 62–63).

During the period 1660–1710, the expected-payoff hypothesis was accepted by most scholars working in the field, including, among others, two members of the Bernoulli family of Swiss mathematicians, Jakob Bernoulli and his nephew Nicolaus Bernoulli. In his influential book *Ars conjectandi* (*The Art of Conjecturing*, completed in 1705 but published posthumously in 1713), Jakob argued that the expected-payoff hypothesis is "the fundamental principle of the whole art [of conjecturing]" (Bernoulli [1713] 2006, 134). In his doctoral dissertation *De usu artis conjectandi in iure* (*On the Use of the Art of Conjecturing in Law*, 1709), Nicolaus applied probability and decision theory to a series of practical issues, ranging from insurance theory to life expectancy. Following Huygens and his uncle Jakob, Nicolaus also maintained that mathematical expectation is the fundamental parameter on which the art of conjecturing should be based (Bernoulli [1709] 1975, 290–291). However, it was Nicolaus himself who, some years later, conceived a game situation that contradicted the expected-payoff hypothesis.

2.3 Nicolaus Bernoulli's Game and Moral Impossibility

In a letter to the French mathematician Pierre Rémond de Montmort dated September 9, 1713 (see Spiess 1975, 557), Nicolaus Bernoulli imagined a game of dice in which individual A pays one *écu* (a French coin used before the Revolution of 1789) to individual B if, by rolling a die, B obtains a six on the first throw, two *écus* if B obtains a six on the second throw, four *écus* if B obtains a six on the third throw, and so on. The expected payoff of this game is $\left(1 \times \frac{1}{6}\right)+\left(2 \times \frac{5}{36}\right)+\left(4 \times \frac{25}{216}\right)+\ldots$, that is, $\sum_{i=0}^{\infty} 2^i \frac{5^i}{6^{i+1}}$, which is a positive infinite number. Therefore, according to the expected-payoff hypothesis, B ought to pay an infinite amount of money to participate in this game. This appeared to Nicolaus not only descriptively implausible but also normatively wrong.

In further correspondence with de Montmort (Spiess 1975, 558–560), Nicolaus proposed to overcome the problem by introducing some "moral," that is, psychological, element into the purely mathematical expected-payoff formula. In particular, Nicolaus made use of the notion of "morally impossible" (*moraliter impossibile*) events put forward by his uncle Jakob in *Ars conjectandi* (Bernoulli [1713] 2006, 316). Jakob had argued that people of good sense consider events with a very small probability as impossible and therefore treat them as if the probability of these events were zero. Nicolaus applied this idea to his dice game and argued that people are willing to pay only a limited amount of money to participate in it because they consider the events associated with high gains as morally impossible.

Considered from a contemporary viewpoint, Jakob's and Nicolaus's idea that individuals may attach to an event a moral weight that is different from its objective probability has some common traits with the idea of "probability weighting" that will be discussed in Section 7 as a key feature of prospect theory.

2.4 Cramer's Game and Moral Value

Some years later, in 1728, Nicolaus resumed the discussion about the game and its conflict with the expected-payoff hypothesis in correspondence with Gabriel Cramer, another eminent Swiss mathematician. In a letter to Nicolaus dated May 21, 1728 (Spiess 1975, 560–561), Cramer put forward a simplified version of the game in which a coin rather than a die was used. In Cramer's setting, A pays one *écu* to B if, by tossing a coin, B obtains tails on the first throw, two *écus* if B obtains tails on the second throw, four *écus* if B obtains tails on the third throw, and so on. Cramer calculated that the expected payoff of this coin game is $\left(1 \times \frac{1}{2}\right)+\left(2 \times \frac{1}{4}\right)+\left(4 \times \frac{1}{8}\right)+\ldots=\frac{1}{2}+\frac{1}{2}+\frac{1}{2}+\ldots$, that is, infinite, just like the expected payoff of Nicolaus's game. It is Cramer's coin game, and not Nicolaus's dice game, that was later labeled the "St. Petersburg game" (see Section 2.6.2).

2.4.1 Moral Value

Cramer agreed with Nicolaus that "no person of good sense" (*de bon sens*) (Spiess 1975, 560) would be willing to pay an infinite amount of money to participate in games such as those he and Nicolaus had conceived of. In order to solve the problem, Cramer also introduced a "moral," meaning a psychological, element in the expected-payoff formula. But while Nicolaus focused on morally impossible events, that is, on the psychological evaluation of probabilities, Cramer concentrated on "moral values," that is, on the psychological evaluation of money.

In his letter to Nicolaus, Cramer argued that the reason for the discrepancy between the expected-payoff hypothesis and the behavior of reasonable people stems from the fact that "mathematicians evaluate money in proportion to its quantity while people of good sense evaluate money in proportion to the use they can make of it" (560). Cramer called people's evaluation of money the *"moral value*" (*valeur morale*) of money.

Cramer submitted two distinct hypotheses about the moral value of money. First, people may consider all amounts of money above a certain level as equivalent. For instance, Cramer suggested, if an individual considers all amounts of money above 2^{24} = 16,777,216 *écus* as equivalent, the price she is willing to pay to participate in the St. Petersburg game is not infinite, and more precisely is around 13 *écus*. Second, Cramer observed that "100 million yield more satisfaction than 10 million, but not 10 times as much" (Spiess 1975, 561); in other words, the moral value of money increases less than proportionally to the increase of money. For Cramer, a possible way of capturing this circumstance is to assume that the moral value of a sum of money x is given by its square root $\sqrt{x}$. Under this latter assumption, Cramer continued, the "moral expectation" (*esperance morale*) of the coin game is finite (as opposed to its infinite mathematical expectation), and the price a person of good sense ought to pay for the game is less than three *écus*.

2.4.2 Priority Issues

Considered from a contemporary standpoint, Cramer's two hypotheses are already instantiations of EU. What Cramer called the moral value of money is akin to the utility of money $u(x)$, and both of his hypotheses state that the value of the coin game is given by $\sum u(x_i)p(E_i)$. In the first hypothesis $u(x_i) = x_i$ for all $x_i \leq 2^{24}$, and $u(x_i) = 2^{24}$ for all $x_i > 2^{24}$, while in the second hypothesis $u(x_i) = \sqrt{x_i}$ for all x_i. Nonetheless, I think that the standard practice of associating the birth of EU with Daniel Bernoulli (introduced in Section 2.5) is legitimate, and this is for various reasons.

First, as we will see in a moment, Bernoulli was not aware of Cramer's hypotheses when he proposed his own version of EU. Second, Bernoulli's arguments in favor of EU are much more extended and systematic than the arguments cursorily suggested by Cramer in his letter to Nicolaus. Third, the two theories are similar but not identical, and in particular Cramer did not take