

EXPECTED UTILITY THEORY, JEFFREY'S DECISION THEORY, AND THE PARADOXES^a

Philippe Mongin* & Jean Baccelli**

May 2020

Abstract. In Richard Bradley's book, *Decision Theory with a Human Face* (2017), we have selected two themes for discussion. The first is the Bolker-Jeffrey (BJ) theory of decision, which the book uses throughout as a tool to reorganize the whole field of decision theory, and in particular to evaluate the extent to which expected utility (EU) theories may be normatively too demanding. The second theme is the redefinition strategy that can be used to defend EU theories against the Allais and Ellsberg paradoxes, a strategy that the book by and large endorses, and even develops in an original way concerning the Ellsberg paradox. We argue that the BJ theory is too specific to fulfil Bradley's foundational project and that the redefinition strategy fails in both the Allais and Ellsberg cases. Although we share Bradley's conclusion that EU theories do not state universal rationality requirements, we reach it not by a comparison with BJ theory, but by a comparison with the non-EU theories that the paradoxes have heuristically suggested.

Keywords: Normative decision theory, Expected utility theory, Bolker-Jeffrey theory, Allais paradox, Ellsberg paradox

1. Introduction

While the theory of decision under risk and uncertainty has long departed from the expected utility (EU) rule of decision in its empirical branch, this rule still looms large in the concerns of its normative branch. This is not to say that it goes unquestioned there; rather, it is

^a The authors thank Orri Stefánsson for inviting this paper and giving them the opportunity to comment on Richard Bradley's thought-provoking work. They also thank two anonymous reviewers for their comments on a first draft.

* GREGHEC, CNRS & HEC Paris; LEMMA, Université de Paris 2-Panthéon-Assas. Email address: mongin@greg-hec.com

** University of Pittsburgh. Email address: jean.baccelli@gmail.com (Corresponding author)

subjected to regular waves of reappraisal, with a fairly wide range of final conclusions, some of which clearly supportive of the claim that EU theory captures the essentials of a theory of rational agency, others clearly dismissive of this claim, whereas still others accept it under diverse qualifications. Generally, these assessments focus on the version of the theory that Savage's canonical treatise, *The Foundations of Statistics* (1954-1972), developed for decision under uncertainty, i.e., the subjective expected utility (SEU) version. But they also often cover the version that von Neumann and Morgenstern (1944-1947) and their followers had developed for decision under risk, as well as alternatives that are less easy to classify. The common feature of most of these versions is to represent an individual's choice as if it resulted from the maximization of the mathematical expectation, for some probability measure that may or may not be specific to that individual, of some quantitative notion of benefit that is specific to that individual.

Richard Bradley's recent book, *Decision Theory with a Human Face* (2017), offers an original variation on this format of assessment. Despite its proclaimed objective of making decision theory psychologically more realistic, it does not call into question the standard division between the positive and the normative when it comes to expounding decision theory, and most of its material can be located on the latter side of this division. At the same time, while laying significant emphasis on the EU rule of decision, and especially its SEU form, it approaches this rule from the perspective of Jeffrey's (1965-1983) theory of decision, which relates to it in a notoriously complex way. Neither does Jeffrey separate his concepts in the way EU theorists normally do, nor does his quantitative rule of decision mathematically belong to the expectational model that these theorists share. For Bradley, these departures mean a methodological advantage. As he writes, "the fact that Jeffrey's theory imposes much weaker requirements on the framing of decision problems is [his] primary reason for preferring his framework to Savage's for developing a theory of decision with a human face" (p. 20). Underlying the whole book is the view that Jeffrey's theory outlines the definition of rational agency that decision theory needs in its normative guise. That the desiderata EU theorists put on individual choice may not be compelling will come out clearly if an encompassing perspective on rationality is adopted, and – here is the major claim, which our first aim will be to evaluate – Jeffrey's theory can provide such a perspective. Bradley devotes a great deal of technical effort to flesh out Jeffrey's theory and its axiomatic complement in Bolker (1966) so as to make this last claim plausible. For he also makes it clear that these pioneering contributions provide no more than an outline that needs to be filled out.

The second aim of the paper is to follow the book in another, quite different theoretical direction. Although Bradley makes generally critical judgments on EU theories when he approaches them from the viewpoint of his enriched Bolker-Jeffrey (BJ) theory, he seems to be unconvinced by the more familiar objections that have been raised against them on the basis of two famous paradoxes of decision theory, i.e., Allais's (1953) and Ellsberg's (1961). Unlike some decision theorists who view these paradoxes as merely empirical refutations, he recognizes that they may have a potential for normative refutation, but his considered view seems to be that they can be deflated by the technique of redefining the decision problems appropriately. Besides being a popular theme in the philosophy of decision theory, this redefinition technique is in and for itself of significant conceptual interest. Bradley's book provides a natural opportunity for dealing with it because he discusses it at more than one place; see p. 55-56 and 172-174 concerning Allais's paradox, and p. 175-177 and 287-288 concerning Ellsberg's. The way he sets up a redefinition strategy against the latter paradox is original, which is another reason for our second thematic choice. (We will also use Bradley's 2016 article, which provides more details on this strategy.)

In a book in which there is so much to praise and approve of, we have selected the two aforementioned themes precisely because we found them to be loci of disagreement. We concur with Bradley on his claim that the normative conditions imposed by EU theories are not universal rationality requirements, but doubt that the Bolker-Jeffrey (BJ) theory, even when expanded as is done in the book, has the resources to reorganize decision theory as a whole, and in fact we also doubt that it can offer a truly enlightening perspective on EU theories. This critical argument is the object of section 2. In another departure, we will call into question the redefinition strategy used to defend these theories against the paradoxes. Section 3 carries the argument for Allais's paradox, and section 4 for Ellsberg's. By investing the paradoxes with normative force and observing how exactly the redescription strategy fails to dissolve them away, we are led to suggest an alternative to Bradley's treatment of these theories: It is the corpus of non-EU theories developed from the paradoxes that offer the most illuminating viewpoint on the limitations of the EU rule of decision.

2. The Bolker-Jeffrey theory and expected utility theory

This section comments on the BJ theory, both in its original form and in Bradley's expanded form, with a view of determining whether it really deserves its overarching role. Although a complete assessment would exceed the limits of this article, we mean to take a position on the

two main issues involved here, i.e., whether the original BJ theory is general enough for a reconstruction of decision theory, and whether, once expanded, it provides a benchmark to evaluate EU theories. We will tackle these two issues in succession. When dealing with the first, we will often use Savage as a point of comparison – a standard practice in discussions concerning Jeffrey.

The formal construction in the original BJ theory is erected by using a single building block - the *proposition*, taken in the Boolean sense of propositional logic, hence subjected to the standard operations of negation, disjunction and conjunction.¹ This unique primitive is meant to express three decision-theoretic concepts that, by contrast, Savage distinguishes formally, i.e., the state of nature, the consequence, and the act. By another contrast, Bradley's formal reconstruction still uses propositions as building blocks, but differentiates them internally so as to express the differences between the preceding three concepts. In the original BJ theory, any proposition has the ability to express states, consequences or acts, depending on which interpretation pleases the decision theorist. Jeffrey's followers have often argued for this semantic versatility on the ground that Savage's alternative approach involves conceptual difficulties; these are reviewed in Joyce (1999, ch. 2) and Bradley (2017, ch. 1), who take roughly similar lines. There is no disputing the fact Savage's definitions do not fully agree with the common sense understanding of decision-theoretic concepts. This may be clearest from the fact that Savage's approach requires the state-independence of utility and the act-independence of probability. But this does not say that the BJ theory is in a position to resolve the familiar difficulties which this raises. To the contrary, it seems that a solution would be forthcoming only if the concepts of states, consequences and acts were kept distinct. To blur their differences is not a way of overcoming the problem of how to formalize them.²

In Bolker's (1966) mathematical treatment, which is the advanced part of BJ theory, the propositions receive a more precise algebraic form and the preference relation, as defined on these objects, satisfy axiomatic conditions in the style of those of decision theory. A representation theorem follows, i.e., from this material Bolker derives the existence of a utility function V , called *desirability* by Jeffrey, and of a probability measure P , both defined

¹ The Boolean operations can of course be replaced by the standard set-theoretic operations on some set. Jeffrey's followers use this handier notation.

² As Bolker (1967, p. 335) himself writes, "we blur the often useful distinction among acts, consequences and events". This critical point is a major reason why mathematical decision theorists are generally unattracted to the BJ theory (see, e.g., Fishburn, 1981, p. 194). Bradley (2017, p. 159) notes this disinclination, in contrast with the high consideration of the theory among philosophers.

on the Boolean algebra of propositions Ω . These items combine as follows: for all disjoint propositions α, β in Ω ,

$$(*) V(\alpha \vee \beta) = [P(\alpha)V(\alpha) + P(\beta)V(\beta)] / [P(\alpha) + P(\beta)],$$

where \vee denotes the Boolean disjunction. For exactness, we reproduce a statement of Bolker's representation theorem in the appendix (see also Bradley, 2017, p. 83). We will now question the suggestion commonly made by Jeffrey's followers that the conditions in this theorem are easily acceptable as rationality conditions, or at least are more so than those of Savage's representation theorem. Here is Bradley (2017, p. 20) again: "the foundational representation theorems for [Jeffrey's theory] require much weaker assumptions about rational preference [than Savage's]".³

Some of Bolker's conditions bear on the objects of preference themselves, i.e., the propositions. He requires them to belong to a *complete atomless* Boolean algebra. This limits the range of applications of the theorem, as he implicitly recognizes (see Bolker, 1967, p. 336-337). Heuristically, there is a tension between the two requisites on the algebra, since completeness would be most easily secured by taking it to be the power set of some set, but this choice would contradict the fact that it is atomless, since the singletons would then enter the algebra. As Bolker mentions, a complete atomless Boolean algebra is an essentially unique object.⁴ One should add in fairness that Savage also imposes strong structural restrictions on his basic sets. In particular, like Bolker's non-atomicity condition, his divisibility postulate P6 entails that the state set has infinite cardinality. However, there are alternative settings and axiom systems that place structural restrictions differently from those of Savage, but deliver equally powerful derivations of the SEU rule, so that users of decision theory can avail themselves of whatever representation theorem best agrees with their intended application. Thus, Anscombe and Aumann's (1963) representation theorem adapts Savage's to a finite state space.⁵ By contrast, Bolker's representation theorem has remained a mostly isolated performance. Some writers have proved a version of it for finite Boolean algebras, but by offering only a partial replication, because they lack the uniqueness statement

³ Note that this claim literally concerns the preference axioms, not the domain assumptions, to be discussed in the next paragraph.

⁴ Up to an isomorphism, it is identical to the set of measurable subsets of the unit interval, when two subsets that differ by a set of measure zero are identified with each other; see Halmos (1974, p. 173). Singletons, which are measurable subsets of the unit interval, disappear from consideration as they are identified with the empty set.

⁵ There are other variants with the same purpose, each being based on a different set of assumptions; a recent example has appeared in Mongin and Pivato (2015).

that a satisfactory representation theorem needs to complement an existence statement – see the appendix for details.

Bolker defines the preference relation on the complete atomless Boolean algebra of propositions less its minimal element, and submits it to two axioms. *Impartiality* does not appear to have a justification independently of the objective of deriving the existence of a probability measure in his representation theorem. Bolker (1967, p. 337) and Jeffrey (1965-1983, p. 147) acknowledge this defect; they may have argued that Savage's corresponding postulate P4 is not in a better position. By contrast, they are willing to defend *Averaging* as a relatively modest rationality axiom, compared with Savage's demanding Sure-Thing Principle (STP), and all their followers – no doubt Bradley among them – endorse this favourable assessment. Indeed, granting the difference in preference objects, *Averaging* is a form of the dominance condition of decision theory, which is widely regarded as the least problematic part of the STP. But this praise must be qualified by the observation that *Averaging* forces the elements of the Boolean algebra to have non-zero probability values; so much becomes clear once the representation theorem is proved. For all its constraining postulates, Savage's system at least avoids this unpalatable implication. To allow for zero probability events matters a great deal not only when there are several interacting agents to consider, as in game theory, but also in decision theory, when one includes the single agent in a dynamic setting.

We have just argued that if Jeffrey's definition of the objects of preference is non-committal, this is a mixed blessing, and that Bolker's conditions for their part are so specific as to be uneasy to apply. But Bradley's book suggests an answer to each worry. For one thing, he undertakes to give Jeffrey's non-descript propositions some internal structure borrowed from the logical theory of conditionals; by this move, he intends to reproduce Savage's distinction between states, consequences and acts within Jeffrey's propositional framework.⁶ For another thing, using a desirability-probability pair (P, V) defined on these new preference objects, he states a number of theses, conditions and properties, thus illustrating the expressive power of the BJ theory *in concreto*. These items specialize the representation of preferences one way or another, and a connection with EU theories arises at this juncture. Precisely, two theorems (2017, p. 164 and 168) are said to recover EU formulas, in the spirit of von Neumann and

⁶ With relevant technical differences, this move is already performed in Bradley (2007). Joyce (1999) also considers the possibility of giving internal structure to Jeffrey's propositions, but his final move consists in partitioning the algebra of these propositions in different ways, with each partition corresponding to a particular interpretation of a decision-theoretic concept (i.e., a state, a consequence or an act). Bradley also uses partitions of the set of propositions with this semantic purpose; see the example in the paragraph following the next.

Morgenstern and Savage respectively, from relevant lists of specializing conditions. As these conditions are “very demanding” (2017, p. 169), the theorems warrant the judgment that EU theories go beyond what may be expected from a rational agent.⁷ Let us now comment on the two moves just stated.

Bradley’s derivation of EU formulas within his enriched BJ theory goes in terms of a (P, V) representation, not the primitive preference relation. It is obviously more convenient to work with the former than the latter. However, only a representation theorem can tell whether this shortcut is justified. One should first ensure that when putting more structure on the original BJ propositions, one has not lost the completeness and non-atomicity properties of the Boolean algebra that are needed to derive the version of Bolker’s representation theorem that fits the new framework. Thus, Bradley’s (2017, p. 70) definition of a conditional algebra should pass this test.⁸ Second, granting that it does, one should check whether the direct use of a (P, V) representation to state EU-relevant conditions agrees with the uniqueness conclusion of the theorem. Unfortunately, a problem arises in this respect. We illustrate it by commenting on Bradley’s (2017, p. 168) rendering of the SEU value of an act.

What formally represents an act is an element $\wedge_i (S_i \vdash C_i)$ of the conditional algebra, where \wedge is the Boolean conjunction, \vdash an axiomatically defined indicative conditional operator, and for $i = 1, \dots, n$, S_i and C_i are partitioning subsets of propositions (with the interpretations that S_i represent states, and the C_i represent consequences).⁹ For acts so defined, the following defines a SEU representation:

$$(**) \sum_{i=1, \dots, n} V(S_i \wedge C_i) P(S_i),$$

If one replaced (P, V) by (P', V') , this formula would be preserved only if P' and P were equal, and V' were a positive affine transform of V . These are the familiar uniqueness conditions of a SEU representation. However, the uniqueness part of Bolker’s representation theorem is not strong enough to support them (see the appendix). This lack of invariance of $(**)$ would appear to block the ensuing analysis of SEU theory.¹⁰ The problem raised here may affect

⁷ There are more details on these theorems in Bradley and Stefansson’s (2017) article. This article takes an even stronger stand against the view that a rational agent should behave as EU theories prescribe.

⁸ In particular, this definition should handle the role of singletons appropriately; see fn. 4.

⁹ Notice incidentally that this propositional rendering of an act only holds for simple acts, i.e., those which have a finite number of distinct values. Savage’s formal concept of an act is not so restricted.

¹⁰ This invariance problem has a clear conceptual underpinning. When the standard uniqueness property of SEU representations holds, the agent’s beliefs and desires are well identified and moreover separated from each other. But Bolker’s uniqueness conditions are too weak to fulfil this purpose, as, e.g., Joyce (1999, p. 136) notes.

other statements of theses, conditions or properties, since all these statements are in terms of a (P, V) representation.

There is a possible remedy to the looseness of the uniqueness conditions in Bolker's representation theorem. As he has shown, if the V function is unbounded above and below, these conditions become identical with those of standard SEU theory. However, this boundedness restriction appears artificial and furthermore it is stated in terms of the representation itself, not the preference relation as would be desirable.¹¹ An alternative move, initially contemplated by Jeffrey and later implemented by Joyce (1999, p.138-145), is to complement Bolker's preference relation with a qualitative probability relation, and connect the two by a suitable coherence condition. When this is done, the uniqueness part of Bolker's representation theorem involves the same uniqueness conditions as in standard SEU theories. However, this strengthening in the conclusion is the counterpart of a strengthening in the assumptions that may be resisted not simply on the grounds of mathematical elegance, but also for a theoretical reason. The established analysis of subjective probability in decision theory consists in deriving it from the preference relation alone, which thus serves as the common primitive of the representation of beliefs and desires. To depart from this scheme is tantamount to departing from the main tradition of decision theory. Joyce fully endorses this departure as a consequence of his "non-pragmatist" position in the philosophy of probability, a position he formulates as the rejection of the tenet that "the laws of rational belief are underwritten by the laws of rational desire" (1999, p. 90; see also Joyce, 1998). Bradley can evade the predicament of his non-invariant formulas by replacing Bolker's original representation theorem by Joyce's version, but it remains to be seen whether he would be willing to pay the philosophical price that this substitution involves.¹²

To summarize this section, we have begun by arguing that the original BJ theory could not constitute an appropriate starting point for a reconstruction of decision theory. Bradley in effect grants this point since he proposes enriching the BJ theory, but in another step, we have questioned the possibility of recovering EU theories via this construction.

¹¹ Jeffrey's (1965-1983, p. 142) suggestion to match the unboundedness of desirability with a preference condition is obscure and usually omitted from the ensuing literature.

¹² Bradley (2017, p. 84-85) states Joyce's version of Bolker's representation theorem and adds that he will draw on it. Whether he effectively does in the sequel would need to be clarified.

3. The redefinition strategy in the Allais paradox

Moving to the second theme of this paper, this section explores the redefinition strategy that has sometimes been proposed to deflate the Allais paradox, a strategy that is echoed, and as it seems, by and large approved, by Bradley (2017, p. 55-56 and 172-174). A brief reminder concerning the paradox is in order.¹³ It challenges EU theory in the form given to it by von Neumann and Morgenstern (1944-1947) and followers. Their setting drastically simplifies Savage’s arrangement of states, consequences and acts. There is a consequence set, no state set, and the objects of preference are *lotteries*, which are formalized as probability measures on the consequence set; these objects implicitly play the role of acts since an act of choice can be identified with what it chooses.¹⁴

Allais (1953 and 1979) staged an ideal agent who is asked to make two choices in succession, first between lotteries p_1 and q_1 , and then between lotteries p_2 and q_2 (the numbers are million French francs):

p_1 : 100 with prob 1	p_2 : 100 with prob 0.11; 0 with prob 0.89
q_1 : 500 with prob 0.10; 100 with prob 0.89; 0 with prob 0.01	q_2 : 500 with prob 0.10; 0 with prob 0.90

Table 1

Under the assumption that the agent’s choices comply with the von Neumann-Morgenstern (VNM) theory, the following equivalence holds:

(***) The agent chooses p_1 over q_1 if and only if he chooses p_2 over q_2 .

This is readily checked by noting that the following two inequalities are equivalent:

$$u(100) > 0.10 u(500) + 0.01 u(0) + 0.89 u(100)$$

¹³ The paradox has given rise to a large literature, which is covered in part by Mongin (2019).

¹⁴ “Outcome” is a more common term than “consequence” in the VNM context, but we use the latter for uniformity.

and

$$0.11 u(100) + 0.89 u(0) > 0.10 u(500) + 0.90 u(0).$$

Allais's proposed resolution of the two choice problems is (p_1, q_2) , which thus violates the VNM theory. The large experimental literature spawned by his paradox has led many, even among specialists, to overlook the fact that he intended it as a *normative* argument against this theory. Mongin (2019) has stressed this underestimated facet of the paradox, and distinguished between a weak normative reading, to the effect that the (p_1, q_2) pair of choices cannot be dismissed as being irrational, while the VNM-abiding pairs cannot be dismissed on this ground either, and a strong normative reading, to the effect that for at least some individuals, (p_1, q_2) is the unique rational pair. Allais sketches two reasons for his resolution. Even on the weaker reading of the paradox, they are an embarrassment for EU theorists.

According to Allais's *certainty argument*, in the first choice problem, the absolute certainty of p_1 , given the relatively high gain it secures, is a reason for the agent to choose it over q_1 despite the fact that the chance of getting nothing in this lottery is very small. In the other choice problem, the same agent could very well choose q_2 over p_2 , i.e., the riskier of the two lotteries, because the chances of getting nothing are nearly equal whereas the possible gains are substantially different. According to the *complementarity argument* (well identified by Bradley, 2017, p. 172-173), the 1% chance of getting nothing in lottery q_1 of the first choice problem reappears in the second one as a difference between the 90% chances of getting nothing in q_2 and the 89% chances of getting nothing in p_2 . However, this 1% chance has a different psychological effect in the two situations: being isolated in the first, it has a weight it does not have in the second, where other chances of getting nothing occur in both lotteries.

There is a possible move to salvage VNM theory from the paradox, and this move is not entirely unrelated to Allais's statement of reasons for the paradoxical choices.¹⁵ Roughly speaking, it consists in saying that the agent's choices are influenced by his anticipation of the negative feelings that these choices could bring about after the fact. This broad idea can be cashed in terms of *regret* or *disappointment*. The agent will fear to experience these feelings if the 0 consequence in q_1 realizes, because there is a small unit chance of getting this consequence, but he will not associate the same fear with the realization of 0 in p_2 or q_2 ,

¹⁵ While the explanation suggested by Allais may thus be used as a motivation, Allais himself would not have condoned the move described next. He advocated instead the development of a proper alternative to VNM theory. That being clarified, the first elaborate occurrence of the move may be in Raiffa (1968, p. 85-86), who actually contemplates it without endorsing it. Like Savage (1954-1972, p. 101-103), Raiffa interprets the paradoxical pair of choices as an irrationality that needs to be corrected.

because, this time, there are several chances of getting it and together they make it the likely consequence. This psychological reasoning suggests that the two choice problems may be rewritten in the following way:

p_1 : 100 with prob 1	p_2 : 100 with prob 0.11; 0 with prob 0.89
q_1^* : 500 with prob 0.10; 100 with prob 0.89; 0 and regret or disappointment with prob 0.01	q_2 : 500 with prob 0.10; 0 with prob 0.90

Table 2

Let us now return to (***) . This equivalence is not falsified anymore if the second lottery in the first choice problem is q_1^* rather than q_1 . However, there is much to be said against such a resolution of the paradox.

The first and quite obvious objection is that the theory has been salvaged only in a logically trivial sense. What is problematic in (***) is the implication that if the agent chooses p_1 over q_1 , then he chooses p_2 over q_2 . Now with Table 2, this implication holds, but only because its antecedent clause is not satisfied. One would have rather hoped a reformulation that would permit *inferring* - and thus tentatively explaining - the paradoxical choices. As these contradict the VNM theory, this more ambitious objective would have required that this theory be reformulated. Theories of regret or disappointment avoidance become relevant at this juncture.¹⁶ The laziness of the present revision, which by contrast merely consists in rewriting the choice problems, is obvious. Notice that the logical triviality which we highlight here is much more specific than the methodological triviality discussed elsewhere in the redescription literature, i.e., the concern that seemingly any consideration can be brought to bear on the redescription of an outcome. The propositions currently available for alleviating

¹⁶ See, among others, the theories developed for regret avoidance in Bell (1982) and Loomes and Sugden (1982), and for disappointment avoidance in Bell (1985) and Loomes and Sugden (1986). All these theories significantly depart from the VNM one. On regret theory, see also the retrospective by Bleichrodt and Wakker (2015).

methodological triviality (see for instance the recent survey in Buchak, 2013, ch. 4) leave logical triviality entirely unaddressed.

There is a second objection, which relates to the particular way choice problems are rewritten. If one is willing to include feelings in the definition of consequences, this should be done anywhere that is necessary, and not simply at one place. We illustrate this with disappointment. Unlike regret, which is prompted by comparing the chosen action with a non-chosen one in terms of their respective consequences, disappointment follows from comparing the consequences of one and the same action. There cannot be any disappointment associated with p_1 , but q_1 contains two occasions for disappointment, i.e., the realization of 0 instead of 100 or 500, and that of 100 instead of 500. Each of p_2 and q_2 contains an occasion for disappointment, i.e., the realization of 0 instead of 100, and that of 0 instead of 500, respectively. It would be unnecessary to fill Table 2 with the missing disappointment data if the point were only to preserve VNM theory in a logically trivial sense; but we have just said that this was too mediocre an objective. These data become relevant precisely if one is after a genuine explanation of the Allais paradox in terms of disappointment, and the theories of disappointment avoidance developed in non-EU theory indeed require them to be stated exhaustively.

Lastly and perhaps most importantly, one may question a step that the previous paragraph took for granted. It is questionable that when the choice of a lottery brings about negative feelings, they need to be captured by redefining the consequences. A more natural way of taking them into account is to include them in the *evaluation* of consequences, while leaving the definition of the latter unchanged, and this is indeed how the non-EU theories of regret or disappointment avoidance proceed. They capture the psychological element by properties of the utility function, the way it combines with the probabilities, or both aspects. Admittedly, the issue touched upon here is intricate, because it relates to the principle of *consequentialism*, which claims that as much content as possible should be packed in the description of consequences, and there are relevant arguments for supporting this principle in decision theory. To discuss consequentialism would require a separate paper, but we may briefly say that it is only a methodological principle, and that like any such principle, it should be judged by its fruitfulness. It seems clear that it would be utterly unpractical to develop a theory of regret or disappointment avoidance in terms of a utility function defined on items such as (x with/without regret/disappointment), rather than on consequences as usually defined. Here

again, the testimony of the full-fledged theories of regret or disappointment avoidance available in non-EU theory is unequivocal.

Actually, when the Allais paradox is subjected to the redefinition strategy, the typical objective is not to preserve VNM theory as was implied in our discussion. Instead, most users of this strategy – Bradley among them – apply it to Savage’s STP. They reformulate the paradox by introducing a framework of states of nature in which this principle makes sense. A typical example is as follows. Here lotteries are represented by assuming that the decision theorist makes a drawing from an urn with 100 equiprobable balls:¹⁷

	Balls 1-10	Ball 11	Balls 12-100
p_1	100	100	100
q^*_1	500	0 and regret or disappointment	100
p_2	100	100	0
q_2	500	0	0

Table 3

Given that the 100 ball urn defines a set of states of the world, the initial lotteries of the Allais paradox can be turned into acts in Savage's sense, i.e., mappings from states to consequences. The resulting framework really is intermediate between the VNM one, in which probabilities are attributed to consequences without states being defined, and the Savage one, in which states are explicit and probabilities appear only in the representation theorem.

At any rate, within this framework, the STP can become the focus of attention. As a brief reminder, this principle says that if a preference comparison holds between two acts f and g that share the same consequences over some subset P of the state space, the same preference comparison holds when f and g coincide again on P with different shared consequences and remain otherwise unchanged. In Table 3, f and g are p_1 and q^*_1 and they share a common consequence of 100 million on a set P , which is “Balls 12-100”. Then, f and g become p_2 and q_2 , which have a new common consequence of 0 on P . By the STP, p_1 is preferred to q^*_1 if

¹⁷ This representation is popular among philosophers of decision theory when they discuss the Allais paradox; see Buchak (2013, ch. 4) for an example and a list of previous references. It is actually inherited from Savage (1954-1972, p. 101-103), who used it in his rebuttal of the Allais paradox. However, Savage’s rebuttal did not amount to using the redefinition strategy, which is these writers’ focus of attention.

and only if p_2 is preferred to q_2 , an equivalence that the pair (p_1, q_2) contradicts. The received argument is that Table 3 salvages the STP. For this principle to apply, it must be the case that p_1 is equal to p_2 , and q^*_1 equal to q_2 , outside the subset P of common values, but this is not the case on the state “Ball 11”.

The three objections previously raised against the redefinition strategy can be repeated here with equal force. But there is still another way of disposing of the present use of the strategy, which is simply to point out that a discussion of the Allais paradox should deal with the VNM theory, not with the Savage one. The difference in the objects of the two theories, which we have emphasized, precludes any straightforward identification of their respective axioms. In the standard use of the redefinition strategy, there is a faulty amalgamation of *the VNM independence condition*, which is the prominent axiom of von Neumann and Morgenstern, with the STP, which is the prominent postulate of Savage.¹⁸

We now complete our rebuttal with an argument that is relative to another part of VNM theory, namely the *principle of compound lotteries*. In essence, the principle says that agent makes no difference between a compound lottery – i.e., a lottery having some lotteries among its consequences – and its reduction to an ordinary lottery, as obtained by applying the multiplication rule of probabilities.¹⁹ To see how this works on the Allais paradox, one needs to introduce the following auxiliary lottery:

l : 500 with prob 10/11,

0 with prob 1/11.

Then, Allais's two choice problems are restated as follows:

¹⁸ Heuristically, there is more to VNM independence than to the STP, and this is confirmed by a step in Savage's proof of his representation theorem, in which he derives the VNM lottery framework and axiom system with a view of using the VNM representation theorem as a lemma (see Savage, 1954-1972, p. 73-76, and Fishburn, 1970, p. 203-206). This derivation requires Savage's full set of postulates P1-P6.

¹⁹ The principle is more commonly said of “reduction of compound lotteries”, a slightly misleading phrase because it operates in both directions, from the compound form to the reduced one, and vice-versa. It is actually contained in the mathematical representation of a lottery as a probability measure on the set of consequences X , because this representation automatically identifies a convex combination of probability measures on X (hence a compound lottery) with a probability measure on X (hence a reduced form lottery).

p'_1 : 100 with prob 0.11; 100 with prob 0.89	p_2 : 100 with prob 0.11; 0 with prob 0.89
q'_1 : l with prob 0.11; 100 with prob 0.89	q'_2 : l with prob 0.11; 0 with prob 0.89

Table 4

With this restatement, Allais's pair of choices become (p'_1, q'_2) . Let us now return to Table 2 and see how the principle of compound lotteries would operate on it. Given the redefinition of consequences in the first choice problem, we need to introduce a new auxiliary lottery:

l^+ : 500 with prob 10/11,

0 and regret or disappointment with prob 1/11;

Now from the principle of compound lotteries, Table 2 is equivalent to:

p'_1 : 100 with prob 0.11; 100 with prob 0.89	p_2 : 100 with prob 0.11; 0 with prob 0.89
q'_1 : l^+ with prob 0.11; 100 with prob 0.89	q'_2 : l with prob 0.11; 0 with prob 0.89

Table 5

We claim that this restatement is psychologically dubious. Suppose that a drawing of q'_1 results in the lottery l^+ and the drawing made in this lottery results in the bad final consequence, i.e., *0 and regret or disappointment*. Now let us similarly suppose that a drawing of q'_2 results in l , and the subsequent drawing in the bad final consequence, i.e., 0. Why should the agent have a negative feeling in one sequence of drawings and not in the other? If 0 is disappointing, this is by comparison with 500, and this comparison is the same in both cases. The crucial point is that bygones are bygones. When reaching the stage of a second drawing, the agent does not have to consider the unrealized consequences 100 or 0. A more roundabout argument would take care of the regret interpretation of the negative feeling. From this psychological discussion, it follows that there is no reason to distinguish between l and l^+ as is done in Table 5. Nevertheless, it appears that, if the principle of compound lotteries holds, this table is equivalent to Table 2. Thus, one of two things has to give, either the principle, or the redefinition strategy that led to Table 2. It would be incoherent to deny the principle in order to salvage the strategy, since the strategy itself is intended to salvage VNM theory, which includes the principle as a component part. Accordingly, it is the strategy that should go. Accordingly, one way or another, one cannot avoid venturing into non-EU territory.

4. The redefinition strategy in the Ellsberg paradox

Ellsberg's (1961) paradox challenges SEU theory and more generally the principle of *probabilistic sophistication*, to the effect that the agent's decision rule is based on a subjective probability measure.²⁰ In a conveniently simplified variant, the paradox goes as follows. A three-ball urn is known to contain one red ball and two other balls that are either black or yellow in unknown proportions. One ball is to be drawn from the urn and the agent is faced with two successive choice problems, each of which involves two bets on the colour of the ball. The first choice is between the bet that the colour is red (f_1) and the bet that it is black (g_1); the second choice is between the bet that the colour is red or yellow (f_2) and the bet that it is black or yellow (g_2). Ellsberg's own solution is (f_1, g_2). To see that this pair of choices contradicts probabilistic sophistication, assume in standard fashion that, when faced with two bets, which put the same stake on two different events, the agent prefers to bet on the event he regards as being the more likely. Then, the agent takes "red" to be more likely than "black" (first choice), but also takes "black or yellow" to be more likely than "red or yellow" (second choice), which contradicts the additivity of probability. That Ellsberg's pair of choices (f_1, g_2) also contradicts Savage's core postulate for SEU theory, i.e., the STP, can be seen from the next table, in which bets are listed as rows and states of the world as columns. A \$100 stake has been fixed for concreteness.

	R	B	Y
f_1	100	0	0
g_1	0	100	0
f_2	100	0	100
g_2	0	100	100

Table 6

Since f_2 and g_2 can be obtained from f_1 and g_1 simply by changing their common consequence on the state Yellow from 0 to 100, preferring f_1 to g_1 and g_2 to f_2 violates the STP.

Ellsberg gave a psychological explanation for these choices, which – like Allais for his recommended choices – he also meant to be an argument for their reasonableness, and as an

²⁰ The difference between a SEU formula and the more recent formulas of probabilistic sophistication hinges on the fact that the latter may be non-linear in the probabilities; e.g., they may involve distorting the latter. For a review of the Ellsberg paradox and the surrounding literature, see Machina and Siniscalchi (2014).

invitation to explore non-EU theories. As this goes, the agent prefers betting on events that can be given precise probability values, such as Red in the first choice and Black or Yellow in the second, to betting on events that have imprecise probability values, such as Black in the first choice and Red or Yellow in the second. Exactly 1/3 of the balls are red, 2/3 of the balls are black or yellow, but the proportion of black balls is 0, 1/3 or 2/3, and that of red or yellow balls is 1/3, 2/3 or 1. This account in terms of *ambiguity aversion* has been explored thoroughly in the literature; but we will follow a different line here.

As some mathematical decision theorists have observed, the agent's uncertainty in the Ellsberg paradox can be redescribed in a way that casts a different light on the conflict between this paradox and the SEU conditions.²¹ Specifically, the states of the world can be redefined by specifying not only the colour of the ball drawn, as in the previous table, but also *which set of coloured balls it is drawn from*. In the next table, both pieces of information appear; e.g., R(RBB) means that Red is the colour of the ball drawn and that the urn has one red and two black balls. The latter information only concerns the number of coloured balls, irrespective of the order in which colours may be listed; i.e., instead of RBY, one could have written RYB, YRB, and so on.

	R(RBB)	R(RYY)	R(RBY)	B(RBB)	B(RBY)	Y(RYY)	Y(RBY)
f_1	100	100	100	0	0	0	0
g_1	0	0	0	100	100	0	0
f_2	100	100	100	0	0	100	100
g_2	0	0	0	100	100	100	100

Table 7

This move cannot yet dissolve the conflict with the STP and probabilistic sophistication since it has simply consisted in refining the initial states. However, now comes another move, which is to redefine *the consequences*. This can be done in several ways, and we first follow Bradley's (2017, p. 175-177) before defending an alternative way. Like other writers on Ellsberg's paradox, Bradley relies on the natural symmetry assumption that for any composition of the urn, each ball has an equal chance of being drawn. Then, if the agent bets on B, his chances of winning \$100 are not the same whether B(RBB) or B(RBY) realizes – they are 2/3 in the former and 1/3 in the latter. Bradley argues that consequences should be

²¹ This redescription is due to Ergin and Gul (2009) and further elaborated by Machina (2011).

redefined specifically so as to capture this difference in the chances of winning the bet, and thus replaces them with vectors $(x; \alpha)$, where x is the amount of money that results from the bet and the (redefined) state, and α is the chance of winning the bet that is associated with that state. The next table implements this suggestion.²²

	R(RBB)	R(RYY)	R(RBY)	B(RBB)	B(RBY)	Y(RYY)	Y(RBY)
f_1	(100; 1/3)	(100; 1/3)	(100; 1/3)	(0; 1/3)	(0; 1/3)	(0; 1/3)	(0; 1/3)
g_1	(0; 2/3)	(0; 0)	(0; 1/3)	(100; 2/3)	(100; 1/3)	(0; 0)	(0; 1/3)
f_2	(100; 1/3)	(100; 1)	(100; 2/3)	(0; 1/3)	(0; 2/3)	(100; 1)	(100; 2/3)
g_2	(0; 2/3)	(0; 2/3)	(0; 2/3)	(100; 2/3)	(100; 2/3)	(100; 2/3)	(100; 2/3)

Table 8

After this joint operation on states and consequences, the Ellsberg pair of choices does not contradict the STP anymore. While f_1 and g_1 , as well as f_2 and g_2 , retain common consequence values on Y(RBY), they do not on Y(RYY), and moreover, the discrepancies between f_1 and f_2 , and between g_1 and g_2 , are such that the antecedent condition of the STP is not satisfied; hence this principle holds vacuously. Similarly, the argument used to show that the Ellsberg paradox contradicts probabilistic sophistication vanishes. For instance, it is not anymore possible to conclude from the choice of f_1 over g_1 that the agent regards the event R as being more probable than the event B; this is simply because the stakes are not anymore the same across the two bets. Of course, the point made against logically trivial resolutions applies here no less strongly than it did in our discussion of the Allais paradox. This is however not all there is to say on Bradley's proposal.

On reflection, this proposal seems less natural and principled than the following alternative. Instead of associating with each monetary consequence the chance of winning the bet, given the state, let us associate with it *the chance of drawing a ball of a given colour, from an urn of a given composition*. Comparing with Table 8, this is, first, a more natural progression starting from the refinement of Table 6 into Table 7. Second, this tracks more rigorously the cause of all complications in the Ellsberg scenario, namely that two forms of uncertainty interact, the “subjective” uncertainty pertaining to the composition of the urn, and the “objective”

²² This corresponds to Table 5 in Bradley's (2016) article, which gives more details on his restatement of the Ellsberg paradox. The analysis reported next would also follow, had Bradley replaced each monetary consequence x by $(x; \beta)$, with β the chance, given the redefined state, of receiving x (β generally differing from α , the associated chance of winning the bet).

uncertainty pertaining to the draw of a ball of a given colour from an urn of a given composition. Finally, generally speaking, it is more satisfactory to redefine a consequence (a cell of the initial table) by adding information that is specific to that consequence (like the chance of drawing a given ball from a given urn, which refers only to the column of the cell), than by adding information that pertains to the bet taken as a whole (such as the chance of winning the bet associated to that cell, that generically refers to another cell on the same row). However the strategies of redefining consequences should be finally evaluated, this appears to be more in line with what they intend to achieve.²³ The next table implements this alternative redefinition of consequences.

	R(RBB)	R(RYY)	R(RBY)	B(RBB)	B(RBY)	Y(RYY)	Y(RBY)
f_1	(100; 1/3)	(100; 1/3)	(100; 1/3)	(0; 2/3)	(0; 1/3)	(0; 2/3)	(0; 1/3)
g_1	(0; 1/3)	(0; 1/3)	(0; 1/3)	(100; 2/3)	(100; 1/3)	(0; 2/3)	(0; 1/3)
f_2	(100; 1/3)	(100; 1/3)	(100; 1/3)	(0; 2/3)	(0; 1/3)	(100; 2/3)	(100; 1/3)
g_2	(0; 1/3)	(0; 1/3)	(0; 1/3)	(100; 2/3)	(100; 1/3)	(100; 2/3)	(100; 1/3)

Table 9

With this table, the argument about probabilistic sophistication is still blocked (the stakes differing too much across the bets), but the conflict with the STP reappears. Indeed, f_1 and g_1 , and f_2 and g_2 , have common parts on the Y event, exactly as in the initial setting; so Ellsberg’s pair of choices (f_1 , g_2) violates the STP. The analysis of the paradox could be pursued to further clarify the distinction between the two forms of uncertainty that have been showed to underlie the paradox. The conclusion would be that the EU rule can be maintained on each component of uncertainty when it is taken separately, but still fail when they are considered jointly. The recently developed *second-order EU theories* exploit a similar contrast by introducing EU representations for both first- and second-order uncertainty, while maintaining a divide between these two forms; there is no way of reducing the second-order uncertainty to the first. These theories were in part heuristically motivated by the Ellsberg paradox and arguably provide an explanation for it.²⁴ However, we may stop before this

²³ Buchak (2013, p. 121-122) distinguishes between “local” and “global” redefinitions of consequences. Bradley’s suggestion is of the latter type, while ours is of the former. Buchak is generally critical of the redescription literature, but perhaps more so when it takes the "global" form.

²⁴ See Klibanoff, Marinacci and Mukerji (2005), Nau (2006) and Ergin and Gul (2009). The general principle is to have first-order uncertainty represented by a EU functional whose values serve as arguments for a non-linear transformation of the EU functional that represents second-order uncertainty. There are of course many alternative explanations for the Ellsberg paradox in the ambiguity literature, and we do not touch on them here.

advanced stage, as the last table already shows that the redefinition process fails to salvage SEU theory. A major lesson from this detailed examination is that – be it through second-order EU theory or some other non-SEU model – it is unavoidable to venture into non-EU territory.

5. Conclusion

We have seen that the paradoxes revisited here pose normative challenges to the EU rule, and that redefinition strategies have been devised to meet these challenges. The strategies are notably more complex in the Ellsberg case than in the Allais case, which may explain why the philosophy of decision theory is less profuse on the former than the latter, Bradley being an exception. The major observation in the Ellsberg case is that there are two different kinds of uncertainty at work. A redefinition of the consequences also occurs, but as a by-product of that of the states of nature. With our proposed redefinition, unlike Bradley's, SEU theory is still violated, so the strategy fails here for a logical reason. By contrast, in the Allais case, it fails for a compound of methodological and semantic reasons.

Regardless of these differences, both paradoxes share the important feature that they irresistibly point towards explanations in terms of non-EU theories. By connecting the Allais choices with the feelings of regret or disappointment, the redefinition school opens the way to theories that have been developed to capture the avoidance of these feelings, and these theories depart from the EU rule of decision and its underlying axioms in various ways. By undertaking an explanation of the Ellsberg paradox in terms of a distinction between two forms of uncertainty, the redefinition school opens another Pandora box of non-EU theories. A general feature of either sets of theories is that they contradict the EU rule of decision for given values of their parameters, but also entail it for other values of these parameters, so that they should be more accurately described as *generalizations* of EU theories. As such, they offer a perspective from which it is possible to appreciate what makes EU theories special, and indeed more restrictive than rationality alone would require. Connecting here the analysis of paradoxes with the objections initially raised against the BJ theory, we submit that this non-EU perspective may be more enlightening than the alleged generalization provided by the latter.

APPENDIX

The set of objects in Bolker's (1966) representation theorem as well as in the BJ theory generally speaking (see Jeffrey, 1965-1983) is a Boolean algebra Ω , whose elements will be termed *propositions*, using the analogy with propositional logic that is central to Jeffrey's interpretation. (Bradley says "prospects".) The usual ingredients of a Boolean algebra will be designated here in standard logical notation: \vee and \wedge for the sup and inf operations, \vDash for the induced relation, \top and \perp for the maximum and minimum elements. By a *preference relation*, we mean a weak ordering \succeq on Ω , with \succ and \sim denoting its asymmetric (or strict preference) and symmetric (or indifference) parts respectively.

Bolker's representation theorem. Let Ω be a complete atomless Boolean algebra, less the proposition \perp , and let \succeq be a preference relation on Ω that is continuous and satisfies the following two conditions:

(i) *Averaging condition.* If α and β are disjoint propositions of Ω , then

$$\alpha \succ \beta \implies \alpha \succ \alpha \vee \beta \succ \beta, \text{ and}$$

$$\alpha \sim \beta \implies \alpha \sim \alpha \vee \beta \sim \beta;$$

(ii) *Impartiality.* If α, β, γ are pairwise disjoint propositions of Ω , $\alpha \sim \beta$, not $\alpha \sim \gamma$, and $\alpha \vee \gamma \sim \beta \vee \gamma$, then for all γ' in Ω that is disjoint from α and β , $\alpha \vee \gamma' \sim \beta \vee \gamma'$.

Then, there exists a probability measure P and function V (a "desirability function") such that for all propositions α, β of Ω ,

$$\alpha \succeq \beta \iff V(\alpha) \geq V(\beta),$$

and if α and β are disjoint,

$$V(\alpha \vee \beta) = [P(\alpha)V(\alpha) + P(\beta)V(\beta)] / [P(\alpha) + P(\beta)].$$

Moreover, given the normalization $V(\top) = 0$, the probability measure P' and the function V' can replace P and V in the above equations if and only if there exist $a > 0$ and c such that for all propositions α' of Ω , $cV(\alpha') + 1 > 0$, with the following properties: for all propositions α of Ω ,

$$P'(\alpha) = P(\alpha) [cV(\alpha) + 1] \text{ and } V'(\alpha) = aV(\alpha) / [cV(\alpha) + 1].$$

For finite Boolean algebras, Domotor (1978), and recently, Gravel, Marchant and Sen (2018) derive the existence part of the theorem from rather similar axioms, plus richness or continuity conditions put on the preference relation. These variants have no uniqueness result, which is unsurprising given the techniques of proof appropriate for finite sets. Ahn's (2008) version departs from Bolker's and the above writers' reliance on Boolean algebras, as he

defines his preference objects as sets of lotteries and, crucially, he exploits the topological structure of these sets. The axiom set includes a form of Averaging (there called Disjoint Set Betweenness) and a form of Impartiality (called Balancedness). The representation theorem has existence and uniqueness conclusions that are closely related to those of Bolker, whose mathematical contribution is put to work in the proof, but the utility representation is more precise as it takes the form of a conditional SEU formula.

REFERENCES

- Ahn, D. (2008). Ambiguity without a State Space. *The Review of Economic Studies* 75 (1), 3–28.
- Allais, M. (1953). Le comportement de l'homme rationnel devant le risque : critique des postulats et axiomes de l'école américaine. *Econometrica* 21 (4), 503–546.
- Allais, M. (1979). The So-Called Allais Paradox and Rational Decisions under Uncertainty. In M. Allais and O. Hagen (Eds), *Expected Utility and the Allais Paradox*, pp. 437–683. Dordrecht: D. Reidel,
- Anscombe, F. and R. Aumann (1963). A Definition of Subjective Probability. *The Annals of Mathematical Statistics* 34 (1), 199–205.
- Bell, D. (1982). Regret in Decision Making under Uncertainty. *Operations Research* 30 (5), 961–981.
- Bell, D. (1985). Disappointment in Decision Making under Uncertainty. *Operations Research* 33 (1), 1-27.
- Bleichrodt, H. and P.P. Wakker (2015), "Regret theory: a Bold alternative to the alternatives", *Economic Journal*, 125, p. 493–532,
- Bolker, E. (1966). Functions Resembling Quotients of Measures. *Transactions of the American Mathematical Society* 124 (2), 292–312.
- Bolker, E. (1967). A Simultaneous Axiomatization of Utility and Subjective Probability. *Philosophy of Science* 34 (4), 333–340.
- Bradley, R. (2007). A Unified Bayesian Decision Theory. *Theory and Decision* 63 (3), 233–263.
- Bradley, R. (2016). Ellsberg's Paradox and the Value of Chances. *Economics & Philosophy* 32 (2), 231–248.
- Bradley, R. (2017). *Decision Theory with a Human Face*. Cambridge: Cambridge University Press.
- Bradley, R. and O. Stefánsson (2017). Counterfactual Desirability. *The British Journal for the Philosophy of Science* 68 (2), 485–533.
- Buchak, L. (2013). *Risk and Rationality*. Oxford: Oxford University Press.
- Domotor, Z. (1978). Axiomatization of Jeffrey Utilities. *Synthese* 39 (2), 165–210.

- Ellsberg, D. (1961). Risk, Ambiguity, and the Savage Axioms. *The Quarterly Journal of Economics* 75 (4), 643–669.
- Ergin, H. and F. Gul (2009). A Theory of Subjective Compound Lotteries. *Journal of Economic Theory* 144 (3), 899–929.
- Fishburn, P. (1970). *Utility Theory for Decision Making*. New York, Wiley.
- Fishburn, P. (1981). Subjective Expected Utility: A Review of Normative Theories. *Theory and Decision* 13 (2), 139–199.
- Gravel, N., T. Marchant, and A. Sen (2018). Conditional Expected Utility Criteria for Decision Making under Ignorance or Objective Ambiguity. *Journal of Mathematical Economics* 78, 79–95.
- Halmos, P. (1974). *Measure Theory*. New York: Springer.
- Jeffrey, R. (1965). *The Logic of Decision*. Chicago: Chicago University Press (2nd edition, 1983).
- Joyce, J. (1998). A Nonpragmatic Vindication of Probabilism. *Philosophy of Science* 65 (4), 575–603.
- Joyce, J. (1999). *The Foundations of Causal Decision Theory*. New York: Cambridge University Press.
- Klibanoff, P., M. Marinacci, and S. Mukerji (2005). A Smooth Model of Decision Making under Ambiguity. *Econometrica* 73 (6), 1849–1892.
- Loomes, G. and R. Sugden (1982). Regret Theory: An Alternative Theory of Rational Choice under Uncertainty. *The Economic Journal* 92 (368), 805–824.
- Loomes, G. and R. Sugden (1986). Disappointment and Dynamic Consistency in Choice under Uncertainty. *The Review of Economic Studies* 53 (2), 271–282.
- Machina, M. (2011). Event-Separability in the Ellsberg Urn. *Economic Theory* 48 (2-3), 425–436.
- Machina, M. and M. Siniscalchi (2014). Ambiguity and Ambiguity Aversion. In M. Machina and W. Viscusi (Eds), *Handbook of the Economics of Risk and Uncertainty*, Volume 1, pp. 729–807. Oxford: North-Holland.
- Mongin, P. (2019). The Allais Paradox: What It Became, What It Really Was, What It Now Suggests to Us. *Economics & Philosophy* 35 (3), 423–459.
- Mongin, P. and M. Pivato (2015). Ranking Multidimensional Alternatives and Uncertain Prospects. *Journal of Economic Theory* 157, 146–171.
- Nau, R. (2006). Uncertainty Aversion with Second-Order Utilities and Probabilities. *Management Science* 52 (1), 136–145.
- Raiffa, H. (1968). *Decision Analysis: Introductory Lectures on Choices under Uncertainty*. Reading: Addison-Wesley.
- Savage, L. (1954). *The Foundations of Statistics*. New York: Dover (2nd edition, 1972).

von Neumann, J. and O. Morgenstern (1944). *The Theory of Games and Economic Behavior*. Princeton: Princeton University Press (2nd edition, 1947).