

DESCRIPTIVE DECISION THEORY

Jake Chandler

Descriptive decision theory is concerned with characterising and explaining regularities in the choices that people are disposed to make. It is standardly distinguished from a parallel enterprise, normative decision theory, which seeks to provide an account of the choices that people *ought* to be disposed to make. Much of the work in this area has been devoted to the building and testing of formal models that aim to improve on the descriptive adequacy of a framework known as ‘Subjective Expected Utility’ (SEU). This adequacy was first called into question in the middle of the last century and further challenged by a slew of experimental work in psychology and economics from the mid 1960’s onwards.

This entry first sketches out the basic commitments of SEU, before moving on to some of its best-known empirical shortcomings and a small selection of those models that have been proposed to supersede it. The relation between descriptive decision theory and its normative counterpart is then discussed, drawing some connections with a number of related topics in the philosophical literature.¹

1 The standard model: Subjective Expected Utility

- 1.1 Savage’s representation theorem
- 1.2 Savage’s proof
- 1.3 The probability triangle

2 The issue of Independence

- 2.1 Allais’ paradoxes
- 2.2 Theoretical responses
 - 2.2.1 Probabilistic sophistication
 - 2.2.2 Models with Betweenness
 - 2.2.3 Models without Betweenness

3 The issue of probabilistic belief

- 3.1 Ellsberg’s Three Colour paradox
- 3.2 Theoretical responses
 - 3.2.1 Non-additive ‘probabilities’
 - 3.2.2 Multiple priors

4 The issue of Weak Order

- 4.1 Transitivity
- 4.2 Completeness

5 Descriptive vs normative decision theory

6 Further reading

7 Related entries

8 Acknowledgments

1 The standard model: Subjective Expected Utility

The canonical theory of choice—*Subjective Expected Utility* (SEU)—owes its inception to the work of Savage (1954), building on previous contributions by De Finetti (1937), Ramsey (1931) and von Neumann and Morgenstern (1947). It offers a homogenous treatment of both decisions under ‘*risk*’—situations in which the decision maker has knowledge of, or holds firm beliefs regarding, the objective probabilities of all events pertinent to the success of his or her actions—and decisions under ‘*uncertainty*’—in which he or she does not. In its non-normative incarnation, it proposes at the very least that agents can be described *as if*:

- (a) associating with the possible consequences of the acts available to them two numerical quantities:
 - (i) a ‘*utility*’ corresponding to the degree to which they would desire the outcome to occur and
 - (ii) a ‘*subjective probability*’ corresponding to their degree of confidence in the occurrence of the outcome given the performance of the act, a degree of confidence which may or may not be given by a corresponding assessment of objective probabilities;
- (b) being such that their preferences between acts, and hence their dispositions to chose certain acts over others, are determined by these quantities in such a way that acts are ranked by their *subjective expected utility*, i.e. the subjective probability-weighted sum of the utilities of their possible outcomes.

Ontologically bolder incarnations of the view have it that agents are so describable because they *really do* have degrees of belief and desires, introspectively familiar psychological states, that determine their preferences and choices in such a manner.

A number of important formal results, known as ‘representation theorems’, show that this claim about desirability can be derived from a set of *prima facie* plausible general principles, aka ‘postulates’ or ‘axioms’, pertaining to the agents’ preferences over acts. Furthermore, not only are these axioms *collectively sufficient* to derive SEU’s claim, but a significant proper subset of them also turn out to be *individually necessary*. Unsurprisingly then, much of the work on assessing the empirical adequacy of SEU has focused on the testing of the aforementioned axioms. Such tests could, in the best case, undermine a key reason to endorse the claim and, in the worst, provide grounds to reject it. Accordingly, a brief sketch of Savage’s own early result is in order.

1.1 Savage’s representation theorem

In Savage’s framework, *acts* are modelled as functions that map possible *states* of the world to *outcomes*, the consequences, if you wish, of carrying out the relevant act in the relevant state of nature. The set of acts will be denoted by $\mathcal{A} = \{f_1, f_2, \dots, g_1, g_2, \dots\}$, the set of states by $\mathcal{S} = \{s_1, s_2, \dots\}$ and the set of outcomes by $\mathcal{X} = \{x_1, x_2, \dots, x_n\}$. For present purposes, it can be assumed that the acts considered are *simple*, i.e. that their range is finite. An act will be called ‘*constant*’ if and only if it maps all states onto one same outcome. Sets of states, also known as *events*, will be denoted by upper-case letters $A_1, A_2, \dots, B_1, B_2, \dots$ etc. The set of such events will be denoted by \mathcal{E} . E_i^f will denote the set of states that the act f maps onto outcome x_i , i.e. $\{s \in \mathcal{S} : f(s) = x_i\}$. It will also be useful to denote by fAg the act that maps the states in A to the same outcomes that f does and the states outside of A to the same outcomes that g does.

The agent’s choice dispositions at a given point in time are taken to be determined by his or her preferences, in such a way that, from any set of particular acts, the agent is liable to choose all and only those acts to which no other act is strictly preferred. $f \geq g$ will denote the fact that an agent finds act f to be no less desirable than act g . $>$ (strict preference) and \sim (indifference) respectively stand for the asymmetric and symmetric parts of \geq , so that $f > g$ iff $f \geq g$ but not $g \geq f$ and $f \sim g$ iff both $f \geq g$ and $g \geq f$. It is convenient to extend this preference relation to the set of outcomes by setting, for all outcomes x_1 and x_2 , $x_1 \geq x_2$ iff the constant act that yields x_1 in all states is weakly preferred to the one that yields x_2 in all states.

Savage proves that there exists a certain specific set of constraints on preference orderings over acts that will be satisfied if and only if this ordering is representable by a real-valued function U with domain \mathcal{A} (so that $f \geq g$ iff

$U(f) \geq U(g)$, such that

$$U(f) = \sum_{i=1}^n P(E_i^f)u(x_i) \quad (1)$$

where $u : \mathcal{X} \mapsto \mathbb{R}$ is a consequence utility function unique up to positive linear transformation and $P : \mathcal{S} \mapsto [0, 1]$ is a unique subjective probability function, satisfying $P(\emptyset) = 0$, $P(\mathcal{S}) = 1$, and the finite additivity property $P(A \cup B) = P(A) + P(B)$ for all disjoint events A, B . In other words, U returns the sum of the utilities of the possible outcomes, each multiplied by the subjective probability of the set of states which are mapped onto that outcome.

For the case in which \mathcal{X} is finite, Savage's set of axioms numbers six. Only three of these, however, make an appearance in the subsequent discussion. The first requires no comment:

Weak Order \geq is a weak order, that is: it is both transitive (for all acts f, g, h : if $f \geq g$ and $g \geq h$, then $f \geq h$) and complete (for all acts f, g : either $f \geq g$ or $g \geq f$).

The second tells us that, in comparing two acts, one ignores their behaviour on the set of states in which they have identical consequences:

Sure-Thing For all acts f, g, h, h' and any event A : $fAh \geq gAh$ iff $fAh' \geq gAh'$.

The third is given as follows:

Weak Comparative Probability For all outcomes x_1, x_2, x_3, x_4 and events A, B : if $x_1 > x_2$ and $x_3 > x_4$, then $x_1Ax_2 \geq x_1Bx_2$ iff $x_3Ax_4 \geq x_3Bx_4$.

The rationale for its proposal lies in the idea that, if $x_1 > x_2$, then $x_1Ax_2 \geq x_1Bx_2$ reflects a commitment to the claim that A is at least as probable as B , and hence, so too must $x_3Ax_4 \geq x_3Bx_4$, when $x_3 > x_4$.

These three conditions, it should be noted, are individually necessary for SEU representability, so that any SEU maximizer must satisfy them. In addition, Savage proposes two further non-necessary, aka 'structural', conditions—respectively known as 'Non-Degeneracy' and 'Small Event Continuity', as well as a further, necessary, condition of 'Eventwise Monotonicity', which tells us that, under certain mild circumstances, the result of replacing one or more occurrences of a given outcome by another will yield a preferred act if and only if the new outcome is preferred to the original.

1.2 Savage's proof

With all this in hand, Savage's result can be established as follows. First, one introduces a relation of 'subjective comparative probability' \succeq , such that $A \succeq B$ iff for all outcomes x_1 and x_2 such that $x_1 > x_2$, $x_1Ax_2 \geq x_2Ax_1$ iff $x_1Bx_2 \geq x_2Bx_1$. Savage's axioms can then be shown to ensure that \succeq satisfies a number of appropriate properties, with Small Event Continuity ensuring that \succeq is representable by a subjective probability function P that is unique. It is worth noting that, in the presence of Weak Comparative Probability, it is mainly the Sure-Thing principle that allows the derivation of the additivity property of P .

Second, using these axioms again, it can then be established that an agent is indifferent between any two acts that, for each outcome, assign equal probabilities to the respective sets of states that they each map onto that outcome. In other words:

State Neutrality If $P_f = P_g$, then $f \sim g$, where $P_f(x_i) = P(E_i^f)$.

Since it can also be shown that, for every lottery P in \mathcal{P} , there exists an act f such that $P_f = P$, the important upshot of this result is that one can effectively simplify the representation of the agent's preferences over acts, recasting them as preferences over the smaller set \mathcal{P} of so-called *subjective lotteries*, i.e. subjective probability distributions over outcomes. To simplify notation, the preference relation over \mathcal{P} will be denoted by the same symbol, \succeq , allowing context to disambiguate.

A further application of the axioms lets us establish that these preferences over lotteries satisfy three important properties: (i) a ‘Mixture Weak Order’ condition, requiring the preferences over lotteries to be transitive and complete, (ii) a ‘Mixture Continuity’ condition, the details of which are not of importance here and finally (iii) an ‘Independence’ condition, which, alongside the ordering condition, will be the focus of considerable discussion in what follows.

To present this last condition, one more definition is required, alongside a piece of notation: For any two lotteries P_f and P_g and $\lambda \in [0, 1]$, one can define a third simple lottery $\lambda P_f + (1 - \lambda)P_g$ in \mathcal{P} , the λ -mixture of P_f and P_g , by setting $(\lambda P_f + (1 - \lambda)P_g)(x)$, the probability assigned to outcome x by the mixture lottery, equal to $\lambda P_f(x) + (1 - \lambda)P_g(x)$. It is heuristically useful to think of $\lambda P_f + (1 - \lambda)P_g$ as a higher-order lottery that yields a probability of λ of playing lottery P_f and a complementary probability of playing P_g . The condition then reads:

Independence For all acts f, g and h and all $\lambda \in (0, 1]$: $P_f \geq P_g$ iff $\lambda P_f + (1 - \lambda)P_h \geq \lambda P_g + (1 - \lambda)P_h$.

The proof is then completed by appealing to a result of von Neumann and Morgenstern (1947), which shows that the aforementioned trio of properties is necessary and sufficient for the representability of \geq by a function U such that $U(P_f) = \sum_{i=1}^n P_f(x_i)u(x_i)$, where $u : \mathcal{X} \mapsto \mathbb{R}$ is a consequence utility function unique up to positive linear transformation.

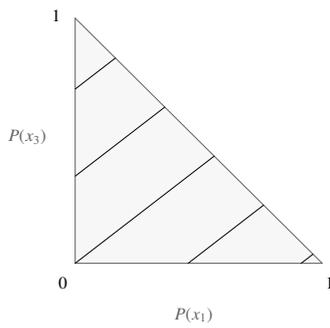
1.3 The probability triangle

The *probability triangle* (aka ‘Marschak-Machina triangle’) offers a helpful visual representation of preferences over the space of lotteries over $\{x_1, x_2, x_3\}$, with $x_3 > x_2 > x_1$. Since, for any $P \in \mathcal{P}$, $P(x_2) = 1 - P(x_1) - P(x_3)$, one can represent the situation two-dimensionally, with lotteries appearing as points in a unit triangle in which the horizontal axis gives us $P(x_1)$ and the vertical one gives us $P(x_3)$. The northwestern, southwestern and southeastern corners respectively correspond to the lotteries yielding x_3, x_2 and x_1 for sure.

Now, as is easily demonstrated, SEU is committed to

Stochastic Dominance For all acts f and g : if, for any outcome x , the probability according to P_f of obtaining an outcome that is weakly preferred to x is at least as great as the corresponding probability according to P_g (in other words: $\sum_{\{y \in \mathcal{X}; y \geq x\}} P_f(y) \geq \sum_{\{y \in \mathcal{X}; y \geq x\}} P_g(y)$), then $P_f \geq P_g$.

Indeed, the above principle follows from Independence and is in fact *equivalent* to Savage’s Eventwise Monotonicity condition, given the other conditions in place (Grant 1995). Therefore lotteries become increasingly preferred both as one moves north and as one moves west, since, in doing either, one shifts probability from a less to a more preferred outcome (from x_2 to x_3 when moving north and from x_1 to x_2 when moving west). The indifference curves are hence upward-sloping. Steeper slopes correspond to greater risk aversion, in the following sense: northeastern movements increase the spread of the distribution, i.e. the degree of risk involved, shifting probabilities from the middle outcome (x_2) to the extremal ones (x_1 and x_3). The steeper the indifference curve, the greater an increase in probability of the best outcome is required in order to compensate for this increased risk. SEU clearly also requires that indifference curves be both linear and parallel.² To illustrate:

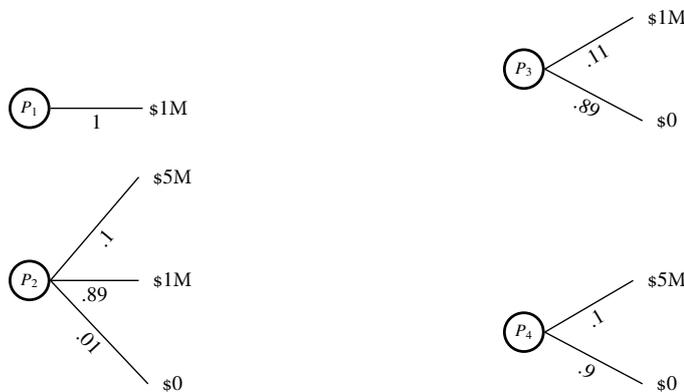


Although SEU continues to enjoy widespread support as a normative model of choice behaviour (though see Section 5 below), it is no longer generally taken to be descriptively adequate. A number of substantial deviations from its predictions were noted as early as the 50's and early 60's by the likes of Allais (1953a, 1953b) and Ellsberg (1961) and further investigated in the 70's. These observations led to the development of alternative models whose own predictive consequences have become the focus of extensive testing in the past three decades or so.³

2 The issue of Independence

2.1 Allais' paradoxes

Allais (1953a, p. 527) considered hypothetical preferences revealed by choices taken from two respective menus of lotteries yielding various increments in wealth with various objective probabilities, one containing P_1 and P_2 below, the other P_3 and P_4 :

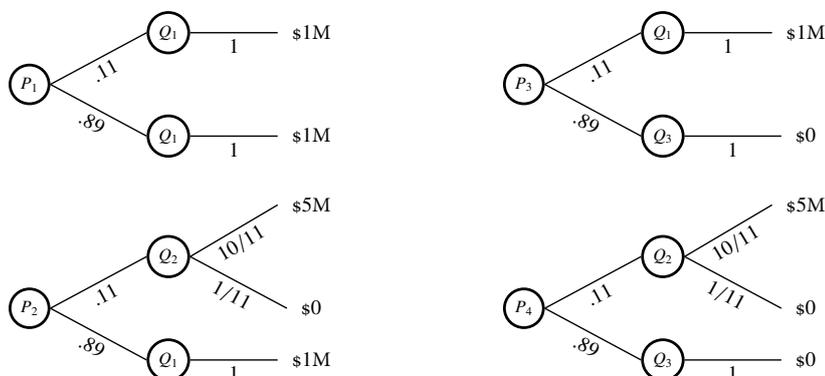


He claimed that, for a substantial proportion of agents, one would find that $P_1 > P_2$ and $P_4 > P_3$ (call these the 'Allais preferences'). However, on the assumptions that (i) the subjects' degrees of belief align themselves with the objective probabilities given and (ii) the outcomes can be adequately characterised entirely in terms of the associated changes in level of wealth, such a combination of preferences runs contrary to Independence. More specifically, it runs counter to the special case of the principle, according to which the substitution of a common 'consequence', i.e. lottery, in a pair of mixtures leave the order of preference unchanged:

Common Consequence For all acts f, g, h, h' and $\lambda \in (0, 1]$: $\lambda P_f + (1 - \lambda)P_h \geq \lambda P_g + (1 - \lambda)P_h$ iff $\lambda P_f + (1 - \lambda)P_{h'} \geq \lambda P_g + (1 - \lambda)P_{h'}$.

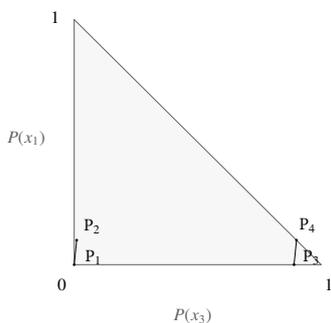
To see why, let $\lambda = 0.11$, Q_1 (the 'consequence' common to P_1 and P_2) be a lottery yielding \$1M for sure, Q_2 be a lottery yielding \$5M with probability 10/11 and \$0 otherwise, and finally Q_3 (the 'consequence' common to P_3

and P_4) a lottery yielding \$0 for sure. P_1 turns out to be a λ -mixture of Q_1 and Q_1 , P_2 one of Q_2 and Q_1 , P_3 one of Q_1 and Q_3 and P_4 one of Q_2 and Q_3 . This is probably best seen by considering the decision trees representing the corresponding compound lotteries:



The upshot of this, by Common Consequence, is then that $P_1 \geq P_2$ iff $P_3 \geq P_4$.⁴

The probability triangle affords a helpful illustration of the incompatibility of the Allais preferences with SEU. Indeed, the segments connecting P_1 and P_2 , on the one hand, and P_3 and P_4 on the other are parallel, so that an EU maximiser, whose indifference curves are also parallel, would be incapable of exhibiting the modal preferences, since no pair of indifference curves could be, as required, such that one crosses the segment $[P_1, P_2]$ from below while the other crosses $[P_3, P_4]$ from above:



In addition to the above, which has come to be known as the *Common Consequence* problem, a further issue, the *Common Ratio* problem, was suggested by Allais (1953a, pp. 529–530). The difficulty this time concerned a further consequence of Independence, which tells us that the order of preference between two identically-weighted mixtures that share a common component lottery is unaffected by a change in the mixture weight:

Common Ratio For all acts f, g, h and $\lambda, \gamma \in (0, 1]$: $\lambda P_f + (1 - \lambda)P_h \geq \lambda P_g + (1 - \lambda)P_h$ iff $\gamma P_f + (1 - \gamma)P_h \geq \gamma P_g + (1 - \gamma)P_h$.

A presentation of the relevant pairs of options will not be given here. Note simply that, here again, the problematic choices turn out to involve two pairs of options whose respective corresponding segments in the probability triangle run parallel.⁵

A number of experimental studies in the 1960's and 1970's subsequently confirmed the robustness of the effects uncovered by Allais. Slovic & Tversky (1974), for example, report that 17 out of 29 (59%) of subjects in their study exhibit Allais preferences in their investigation of the Common Consequence problem. See MacCrimmon & Larson (1979) for a helpful summary of this and other early work and further data of their own.

Since the late 1970's, a considerable number of generalisations of SEU have been devised to accommodate the problematic preference patterns. A brief survey of these is provided in the following subsection.

2.2 Theoretical responses

2.2.1 Probabilistic sophistication

A substantial proportion of the responses to Allais-type phenomena have involved generalisations of SEU that remain conservative enough to preserve the requirement of what Machina & Schmeidler (1992) call '*probabilistic sophistication*': that preferences over acts reduce to preferences over lotteries and that these in turn obey Mixture Weak Order, Mixture Continuity and Stochastic Dominance, if not Independence.⁶ Machina & Schmeidler offer an axiomatic characterisation of probabilistically sophisticated preferences that gives up Savage's Sure-Thing condition, which plays a critical role in the derivation of Independence, and retains the remainder of his conditions. Since the Sure-Thing principle, however, *also* plays an important role in ensuring the existence of a suitable probability distribution over the set of events, they strengthen the Weak Comparative Probability condition to the following:

Strong Comparative Probability For all outcomes x_1, x_2, x_3, x_4 , acts f, g and disjoint events A, B : if $x_1 > x_2$ and $x_3 > x_4$, then $x_1Ax_2Bf \geq x_2Ax_1Bf$ iff $x_3Ax_4Bg \geq x_4Ax_3Bg$.

where x_1Ax_2Bf denotes the act that yields x_1 for all $s \in A$, outcome x_2 for all $s \in B$ and $f(s)$ for all other s . They then offer a correspondingly amended account of the proposed correspondence between the subjective qualitative probability and preference relations, proposing that, if $x_1 > x_2$, then $A \succeq B$ iff $x_1Ax_2Bf \geq x_2Ax_1Bf$.

2.2.2 Models with Betweenness

Among the models of probabilistically sophisticated preferences that do not satisfy Independence and, more specifically, do not impose the property of *parallelism* of indifference curves, a number still satisfy a weaker principle that imposes *linearity*, namely:

Betweenness For all acts f and g and $\lambda \in [0, 1]$: if $P_f \sim P_g$, then $P_f \sim \lambda P_f + (1 - \lambda)P_g$.

This is notably the case of *Weighted Utility* (WU) (Chew and MacCrimmon 1979, Chew 1983), which proposes that the summands in the expected utility formula each be multiplied by a corresponding weight, so that preferences between lotteries are representable by the more general functional

$$U(f) = \sum_{i=1}^n P_f(x_i)u(x_i) \left(w(x_i) / \sum_{i=1}^n w(x_i)P_f(x_i) \right) \quad (2)$$

where w is a positive real-valued function on \mathcal{X} . If w is constant, one recovers the EU functional. The incorporation of weights accommodates Allais preferences by allowing indifference curves to 'fan out' from a single intersection located in the quadrant to the southwest of the probability triangle. These curves become steeper, and hence represent a greater degree of risk aversion, as one moves northwest, in the direction of increasingly preferred lotteries. A suitably placed intersection allows indifference curves to cross both $[P_1, P_2]$ from below and $[P_3, P_4]$ from above, as required.⁷

2.2.3 Models without Betweenness

There is however substantial evidence that the linearity of indifference curves isn't any more empirically adequate than their parallelism (see Camerer & Ho 1994 for a survey) and a number of models of probabilistically sophisticated preferences give up on Betweenness too. The best known of these is undoubtedly *Rank Dependent Utility* (RDU), a version of which was first proposed by Quiggin (1982).⁸ To present the proposal in functional form, it will be assumed that the subscripts associated with each outcome in \mathcal{X} indicate increasing order of preference, so

that $x_1 \leq x_2 \leq \dots \leq x_n$ and hence $\bigcup_{j=i}^n E_j^f$ is the event given which f yields an outcome at least as preferable as x_i . RDU proposes:

$$U(f) = u(x_1) + \sum_{i=2}^n (u(x_i) - u(x_{i-1}))w\left(P\left(\bigcup_{j=i}^n E_j^f\right)\right) \quad (3)$$

where $w : [0, 1] \mapsto [0, 1]$ is a strictly increasing probability weighting function, such that $w(0) = 0$ and $w(1) = 1$. In other words: the utility of a lottery is equal to the sum of the marginal utility contributions of the outcomes, each multiplied by the weighted probability of obtaining an outcome that is at least as preferable (the marginal contribution of x_1 is $u(x_1)$ and its associated multiplier is $w(P(\{S\})) = w(1) = 1$). If w is the identity function, so that $w \circ P = P$, it turns out that one recovers the expected utility functional. If not, a suitable choice of w enables one to recover the Allais preferences. To see how, assume for simplicity that $u(0) = 0$. One then has $P_1 > P_2$ iff $u(1)w(1) > u(1)w(0.99) + (u(5) - u(1))w(0.1)$ and $P_4 > P_3$ iff $u(5)w(0.1) > u(1)w(0.11)$. This implies that the preferences will be recovered by having w be such that $w(1) - w(0.99) > w(0.11) - w(0.1)$, so that a difference in probability of 0.01 has a greater impact at the higher end of the probability scale than it does towards its relatively lower end.⁹

It should be noted that RDU is itself a special case of what is perhaps the best known alternative to SEU, Kahneman & Tversky's *Cumulative Prospect Theory* (Kahneman & Tversky 1992), which earned Kahneman a Nobel Memorial Prize in Economic Sciences in 2002. This model generalises RDU by introducing a *reference point*, an outcome that partitions the set of outcomes into positive and negative subsets, according to whether these are strictly preferred or strictly dispreferred to it. Two probability transformation functions, w^+ and w^- , are then involved in the preference functional: w^+ in determining the utility contributions of the negative outcomes and w^- playing an analogous role in relation to that of the positive ones. RDU is recovered when w^+ is the dual of w^- .

While RDU does not satisfy Independence, it does satisfy a weakening of this principle known as 'Ordinal Independence' (Green & Jullien 1988). This principle is presented as a constraint on the *cumulative distribution functions* corresponding to various lotteries, which return, for each x_i , the probability of obtaining an outcome that is no better than x_i (i.e. an outcome x_j , with $j \leq i$). The cdf corresponding to P_f shall be denoted by F . We then have

Ordinal Independence For all acts f, f', g and g' and subsets A of \mathcal{X} : If $P_f \geq P_g$, and

(i) for all $x \in A$, $F(x) = G(x)$ and $F'(x) = G'(x)$

(ii) for all $x \notin A$, $F(x) = F'(x)$ and $G'(x) = G'(x)$

then $P_{f'} \geq P_{g'}$.¹⁰

The constraint can more helpfully be put as follows: In comparing two acts, one ignores the values of their respective cdf's on the set of outcomes with respect to which they agree. It is easily verified that the Allais preferences are consistent with this principle. Given probabilistic sophistication, Ordinal Independence can itself be derived from a constraint on preferences over acts known as 'Comonotonic Independence', presented in Subsection 3.2.1 below. Wakker (2010) offers a textbook introduction to RDU and Cumulative Prospect Theory, as well as to related treatments of the issues discussed in the next section.

3 The issue of probabilistic belief

3.1 Ellsberg's Three Colour paradox

In another classic challenge to SEU, Ellsberg (1961) asked subjects to consider a setup in which an urn contains 30 red balls and 60 black or yellow balls in unknown relative proportions and report their preferences between various bets on the colour of a ball drawn at random from the urn. The preferences elicited were the ones holding between f_1 and g_1 below, on the one hand, and f_2 and g_2 , on the other:

	30 balls		60 balls	
	r	b	y	
f_1	\$100	\$0	\$0	
g_1	\$0	\$100	\$0	
f_2	\$100	\$0	\$100	
g_2	\$0	\$100	\$100	

Ellsberg reported that a majority of subjects exhibited the preferences $f_1 > g_1$, but $g_2 > f_2$, an instance of a phenomenon that has come to be known as *ambiguity aversion*: a relative preference for betting on events of known rather than unknown (‘ambiguous’) probability.

If one grants that the outcomes are adequately characterised in sole terms of associated changes in level of wealth, these ‘Ellsberg preferences’ stand in direct contradiction with Savage’s Sure-Thing principle. These preferences also violate Machina & Schmeidler’s Strong Comparative Probability principle, on the natural assumption that the subjects strictly prefer the outcome \$100 to the outcome \$0. And indeed it is easy to see that the Ellsberg preferences are inconsistent with probabilistic sophistication. More specifically, they are incompatible with its being the case that both (i) the decision maker’s preferences over acts are reducible to preferences over corresponding lotteries over outcomes, generated by an assignment of subjective probabilities to the set of events and (ii) he or she partially orders these lotteries by first-order stochastic dominance. To see why, assume that these conditions hold. Note first that P_{g_1} would stochastically dominate P_{f_1} if and only if $P(\{b\}) \geq P(\{r\})$ and that P_{f_2} would stochastically dominate P_{g_2} if and only if $P(\{r\}) \geq P(\{b\})$. $f_1 > g_1$ would entail that P_{g_1} does not stochastically dominate P_{f_1} , and hence that $P(\{r\}) > P(\{b\})$. But $g_2 > f_2$ would entail that P_{f_2} does not stochastically dominate P_{g_2} , and hence that $P(\{b\}) > P(\{r\})$. Contradiction.

Considerable empirical evidence has confirmed Ellsberg’s informal observations and related phenomena (beginning with Becker & Brownson 1964 and including classic studies such as Slovic & Tversky 1974 and MacCrimmon & Larsson 1979; see the classic Camerer & Weber 1992, as well as the more up-to-date Trautmann & van de Kuilen 2015, for further details) and the literature now contains a substantial number of generalisations of SEU that can accommodate these.

3.2 Theoretical responses

3.2.1 Non-additive ‘probabilities’

One prominent weakening of SEU that is capable of accommodating the Ellsberg cases is *Choquet Expected Utility* (CEU), initially proposed by Schmeidler (1989). The key concept in its representation of preferences is that of a *capacity*: a function $v : \mathcal{E} \mapsto [0, 1]$, such that $v(\emptyset) = 0$, $v(\mathcal{S}) = 1$ and, for all $A, B \in \mathcal{E}$, $A \subseteq B$ implies $v(A) \leq v(B)$. One can think of this as a kind of non-additive ‘probability’ function, since the additivity property, according to which $v(A \cup B) = v(A) + v(B)$ for disjoint events A and B , does not hold. As with the presentation of RDU, the convention here is that the indices associated with the outcomes indicate increasing preference, so that, again, $\bigcup_{j=i}^n E_j^f$ is the event given which f yields an outcome at least as preferable as x_i . CEU proposes:

$$U(f) = u(x_1) + \sum_{i=2}^n (u(x_i) - u(x_{i-1}))v\left(\bigcup_{j=i}^n E_j^f\right) \quad (4)$$

On this suggestion, then an act is valued by the sum of the marginal utility contributions of the outcomes, each multiplied by the capacity of the event given which that act would yield an outcome that is at least as preferable. There are obvious formal similarities here with RDU and, in fact, the latter can be viewed as the special case of CEU in which the decision maker’s capacities are derived from his or her probabilistic degrees of belief by a probability weighting function ($v = w \circ P$).¹¹

Returning to the Ellsberg preferences in the three colour problem, it is easy to see that $f_1 > g_1$ iff $v(\{r\}) > v(\{b\})$ and $g_2 > f_2$ iff $v(\{b, y\}) > v(\{r, y\})$. These inequalities obviously cannot be simultaneously satisfied in special cases in which c is additive and indeed, in such cases, CEU reduces to SEU. In the more general case, there is no problem: let v , for instance, be such that: $v(\{r\}) = v(\{r, y\}) = v(\{b, y\}) = 1/3$, $v(\{b\}) = v(\{y\}) = 0$ and $v(\{b, y\}) = 2/3$.

Gilboa (1987) and Wakker (1989) have both provided axiomatisations of the proposal in a Savage framework. The key distinguishing feature of these is the effective restriction of Savage's Sure-Thing principle to particular kinds of sets of acts:

Comonotonic Sure-Thing For all acts f, g, h, h' and any event A : if fAh, gAh, fAh' and fAh' are comonotonic, then $fAh \geq gAh$ iff $fAh' \geq gAh'$.

where two acts f and g are *comonotonic* iff, there are no two states s_1 and s_2 , such that $f(s_1) > f(s_2)$ but $g(s_2) > g(s_1)$, or again iff f and g yield orderings of states by desirability of associated consequence that are jointly consistent (Chew & Wakker 1996). Clearly, the Ellsberg preferences are perfectly compatible with this weakening of the Sure-Thing principle, since the acts involved are not comonotonic. For instance, $f_1(r) > f_1(b)$ but $f_2(b) > f_2(r)$.¹²

3.2.2 Multiple priors

The capacity that was used above to illustrate the consistency of CEU with Ellsberg-style preferences has a noteworthy property: it is *convex*, meaning that it is such that, for all $A, B \in \mathcal{E}$, $v(A \cup B) + v(A \cap B) \geq v(A) + v(B)$. It has been shown by Schmeidler (1986) that, if convexity of capacities is imposed, CEU becomes a special case of an approach known as *Maxmin Expected Utility* (MEU) (Gilboa & Schmeidler 1989), which represents the decision maker as maximising minimum expected utility across a non-empty set of probability functions Γ on \mathcal{X} , so that:

$$U(f) = \inf_{P \in \Gamma} \left(\sum_{i=1}^n P(E_i^f) u(x_i) \right) \quad (5)$$

The specific connection is the following: a CEU maximiser with respect to a convex capacity v is an EU maximiner over the so-called *core* of v , defined as the set of probability functions that assign, for every event, a probability that is at least as great as the capacity assigned to that event by v : $\{P \in \mathcal{P} : P(A) \geq v(A), \forall A \in \mathcal{E}\}$.

Now a common, but not mandatory, interpretation of Γ is that it corresponds to the set of objective probability assignments that the decision maker takes to be consistent with his or her evidence. In view of the result just flagged out, this in turn invites an interpretation of capacities as lower estimates of objective probabilities. More specifically, a CEU maximiser whose capacity is convex can be interpreted as considering possible all and only those assignments of objective probabilities that are consistent with the lower estimates given by that capacity. This interpretation of the capacity in the particular example at hand is obviously particularly tempting, as $1/3$ and $2/3$ constitute plausible lower bounds on the decision maker's estimates of the probabilities of $\{r\}$ and $\{b, y\}$, respectively.

If one interprets Γ this way, relaxing CEU with convex capacities to MEU becomes an attractive option, since it allows one to not only model Ellsberg preferences but also accommodate the preferences of decision makers whose views on objective probabilities cannot simply be captured in terms of lower estimates (for example, those involving commitments to certain facts about ratios of probabilities). Due to space considerations, the details of the axiomatic treatment of MEU are omitted here.¹³

Still, MEU remains rather restrictive, as it enforces a fairly radical form of ambiguity aversion. One popular generalisation of the model, α -MEU (Ghirardato *et al* 2004), proposes that the preferences imposed by MEU lie only at one end of a spectrum of possible ambiguity aversion, captured by the following weakening of (5):

$$U(f) = \alpha \inf_{P \in \Gamma} \left(\sum_{i=1}^n P(E_i^f) u(x_i) \right) + (1 - \alpha) \sup_{P \in \Gamma} \left(\sum_{i=1}^n P(E_i^f) u(x_i) \right) \quad (6)$$

where $\alpha \in [0, 1]$. With $\alpha = 1$, one recovers the highly ambiguity-averse MEU. With $\alpha = 0$, we have strongly ambiguity-loving preferences. The parameter α is thus in a sense interpretable as a measure of ambiguity aversion.^{14, 15}

Just as with MEU, however, α -MEU restricts its attention to extremal expected utilities (in this instance best- as well as worst-case). A popular class of proposals allows for the full range of expected utilities across Γ to be factored in, by supplementing the multiple prior model with a *higher order probability distribution* μ . One well-known functional form, that notably features in the ‘*Smooth Model*’ of Klibanoff *et al* (2005), involves taking the expectation, relative to μ , of the *weighted* expected utilities, relative to the members of Γ :

$$U(f) = \sum_{P \in \Gamma} \mu(P) \Phi \left(\sum_{i=1}^n P(E_i^f) u(x_i) \right) \quad (7)$$

A concave Φ will overweigh low expected utilities, resulting in relatively ambiguity-averse preferences.

4 The issue of Weak Order

4.1 Transitivity

While all models mentioned above impose transitivity on preferences, there is a long history of investigating possible violations of the principle, both with respect to choice under certainty and choice under risk. Regarding the latter, in a classic early study, Tversky (1969), suggested significant systematic violations of the transitivity of strict preference, which is entailed by that of weak preference, in relation to a series of lotteries P_1 – P_5 , each offering a chance p_i of receiving a prize x_i and a complementary chance of receiving nothing:

	p_i	x_i
P_1	7/24	\$5
P_2	8/24	\$4.75
P_3	9/24	\$4.5
P_4	10/24	\$4.25
P_5	11/24	\$4

Tversky took his data to suggest that a significant number of subjects were prone to expressing strict preferences for each lottery over its immediate successor, but a strict preference for the last lottery over the first. He proposed that these subjects ranked adjacent lotteries by mere payoff as the differences in the probabilities of winning were barely perceptible, but took probability of winning into consideration in the comparison between P_1 and P_5 , since the difference in values there was large. Although Tversky’s results were later replicated, it should be noted that there is ongoing controversy surrounding the level of empirical support for intransitive preference (see Regenwetter *et al* 2011 for a recent literature review).

Intransitivities of a somewhat different kind are also predicted by Loomes & Sugden’s (1982, 1987) *Regret Theory*.¹⁶ The guiding idea behind this proposal is that the appreciation of a given outcome in a given state is an essentially comparative matter. It is determined by the regret (or the rejoicing) associated with the thought that the alternatively available acts would have led, in the same circumstances, to a particular set of alternative outcomes. In the special case of binary alternatives, this intuition translates into the following *menu-dependent* preference functional:

$$U_{\{f,g\}}(f) = \sum_{s \in S} P(\{s\}) M(f(s), g(s)) \quad (8)$$

where $M : X \times X \mapsto \mathbb{R}$ is a comparative utility function that is increasing in its first argument and non-decreasing

in its second. In their discussion of the framework, Loomes & Sugden present things equivalently as follows:

$$f \succeq g \text{ iff } \sum_{s \in \mathcal{S}} P(\{s\}) \Psi(f(s), g(s)) \geq 0$$

where $\Psi(f(s), g(s))$ is defined as $M(f(s), g(s)) - M(g(s), f(s))$. This quantity thus corresponds to the net balance of regret/rejoicement associated with choosing f over g in states s . ‘Depending on the properties of Ψ , decision makers can be characterised as being ‘regret-neutral’, ‘regret-averse’ or even ‘regret-seeking’. Regret neutrality corresponds to the case in which, for all $x_1, x_2, x_3 \in \mathcal{X}$, $\Psi(x_1, x_3) = \Psi(x_1, x_2) + \Psi(x_2, x_3)$. Under these conditions, choice behaviour is consistent with SEU. Regret aversion corresponds to the situation in which Ψ satisfies the following convexity requirement: for $x_1 > x_2 > x_3$, $\Psi(x_1, x_3) > \Psi(x_1, x_2) + \Psi(x_2, x_3)$. Loomes & Sugden (1982) have shown that, at least under the assumption of probabilistic independence of the lotteries involved, this type of disposition can predict both the Common Consequence and the Common Ratio effects: Regret Theory does not entail Independence.¹⁷

To obtain a sense of the violations of transitivity predicted by Regret Theory, here is an example due to Loomes & Sugden 1987. Assume convexity of Ψ and consider the following decision problem, where $x_1 < x_2 < x_3$ and $P(A_i) = 1/3$:

	A_1	A_2	A_3
f	x_1	x_2	x_3
g	x_3	x_1	x_2
h	x_2	x_3	x_1

According to Regret Theory, $f > g$ iff $\Psi(x_1, x_3) + \Psi(x_2, x_1) + \Psi(x_3, x_2) > 0$. Convexity of Ψ will ensure that this inequality holds. By similar reasoning, it can then be established that $g > h$ and $h > f$.¹⁸

The above example also clearly demonstrates that Regret Theory permits violations of State Neutrality, since the different acts yield the same probability distributions over outcomes. Loomes & Sugden (*ibid*) further show that violations of Stochastic Dominance are licensed by their model. However, in spite of these departures from orthodoxy, it should be noted that Regret Theory retains a number of other strong consequences of SEU, including the Sure-Thing principle, as well as Betweenness for probabilistically independent distributions. An instructive axiomatisation of a generalisation of (8) to finite menus is offered in Sugden 1993. See Bleichrodt & Wakker 2015 for a clear overview of the framework, and its relation to the experimental data.

4.2 Completeness

Although the issue comes last in this catalogue of empirical challenges to SEU, early doubts regarding the empirical adequacy of the completeness assumption were aired by the very architects of the framework, including von Neumann & Morgenstern (1947, p. 630) and Savage (1954, p. 21). For instance, von Neumann & Morgenstern write: ‘It is very dubious, whether the idealization of reality which treats this postulate as a valid one, is appropriate or even convenient’.

Failure of completeness has been claimed to stem both from either (i) incompleteness in judgments of comparative probability or (ii) incompleteness in preferences between outcomes. Both sources of incompleteness can be handled in ‘*multi-prior expected multi utility*’ models, which offer what one might call a ‘supervaluationist’ representation of preferences over acts, as follows:

$$f \succeq g \text{ iff, for all } \langle P, u \rangle \in \Phi, \sum_{i=1}^n P(E_i^f) u(x_i) \geq \sum_{i=1}^n P(E_i^g) u(x_i)$$

where Φ is a set of pairs of probability and utility functions. Due to space considerations, axiomatic details are left out here. The interested reader is referred to the recent general treatment given by Galaabaatar & Karni (2013),

who relate their results to important earlier work by the likes of Bewley (1986), Seidenfeld *et al* (1995), Ok *et al* (2012), and Nau (2006), among others.

5 Descriptive vs normative decision theory

While it was fairly immediately recognised that Allais had demonstrated an empirical shortcoming of SEU, it is important to note that his ambitions somewhat outstripped this achievement. He further suggested that his findings also give reason to doubt the *normative* adequacy of the theory. On his view, two types of consideration can be brought to the table in the assessment of a theory of rational choice. The first is a demonstration that the theory deductively follows from, or lies in logical conflict with, various general principles of secure epistemic standing. The second is a body of experimental evidence regarding ‘the conduct of persons who, one has reason in other respects [(‘that is on criteria that are free of all reference to any consideration of random choice.’)] to believe, act rationally.’ (Allais 1953b, p. 34).¹⁹ However, he found no adequate evidence of the first kind that could be marshalled to support anything quite as strong as SEU. He rejected, for instance, Marshak’s (1951) ‘long-run success’ argument for expected utility maximisation in situations of risk (Allais *ibid*, pp. 70–73). He did grant the existence of a ‘consistency’ requirement according to which ‘a man will be deemed to act rationally (a) if he pursues ends that are mutually consistent (i.e. not contradictory), (b) if he employs means that are appropriate to these ends.’ (Allais *ibid*, p. 78). But, this requirement, he claimed, simply entailed that preferences over lotteries be weakly ordered and satisfy Stochastic Dominance. This left data on choice behaviour to adjudicate on the further commitments of SEU. This data, in his view, clearly supported the rational permissibility of violating Independence.

Savage did not explicitly discuss the probative force of the collective preferences of his peers in relation to Allais’ cases. He did however comment on the bearing of *his own* personal preferences, which Allais had famously elicited from him at a Paris symposium 1952 and which found themselves in violation of the recommendations of SEU. Granting that it would have been irrational for him to maintain both these preferences and a commitment to the normative adequacy of his axioms, he reported that further ‘reflection’ inclined him to revise the former, judging these to have been in error, on par with a logical inconsistency in beliefs. This fact, he claimed, entitled him to retain his normative commitments (see Savage 1952, pp. 101–103).²⁰ Since it is easy to surmise that Savage took his own inclinations to be representative of those of the population at large, his comments have been widely taken to implicitly suggest an alternative experimental route to the testing of theories of rational choice. (See Slovic & Tversky 1974, and Jallais & Pradier 2005. This is also the view of Ellsberg, who offers, in Ch. 1 of his 1961 doctoral dissertation, reprinted as Ellsberg 2001, a worthwhile discussion of the issues of present interest, with Zappia 2016 providing a recent philosophically-oriented discussion.). This procedure would involve determining, not whether certain decision makers exhibit patterns of preference proscribed by the theory, but whether they still exhibit such patterns *after* reflection on their conflict with the theory’s basic axioms.

A number of studies set out to test the normative adequacy of SEU along the proposed lines. MacCrimmon (1968) reported violations, in a sample of experienced business executives, of a wide range of consequences of SEU, a number of which persisted after subjects were notably provided with considerations both supporting and undermining these principles. Those principles with respect to which offending preferences were later corrected included most notably Transitivity and Stochastic Dominance. Allais- or Ellsberg-style preferences were substantially more resilient, however, a fact confirmed in a later study by Slovic & Tversky (1974). Another type of resilience of preferences, not considered by Savage, was more recently investigated by van de Kuilen & Wakker (2006). They studied the effects of providing feedback on decision outcomes on the prevalence of common consequence effects in sequences of choices, finding, however, a significant reduction in SEU violations.

In spite of a long-standing tradition of bringing to bear theories of rational choice on various philosophical problems,²¹ the issue of the potential relevance of descriptive decision theory to its normative counterpart does not appear to have sparked much interest in the philosophical community. Allais’ challenge to Savage has largely been ignored in the philosophical literature.²²

Having said this, a fair amount of philosophical attention has been devoted to the related issue of the connection between norms of *reasoning* and observed patterns of *inference*. One influential line of thought to be found there, which seems pertinent to Allais' claims, originates in Goodman's discussion of the justification of inductive reasoning. On his view, '[t]he task of formulating rules that define the difference between valid and invalid inductive inferences is much like the task of defining any term with an established usage' (Goodman 1965, p. 66). Just as semantic analyses can be endorsed on the basis of providing good systematisations of a set of intuitions regarding the applicability of particular terms in particular situations, Goodman claims, normative theories of reasoning can similarly be justified by their good fit with 'the particular...inferences we actually make and sanction' (Goodman *ibid.*, p. 63): no further considerations are required in order to be able to endorse a particular principle as rationally binding.

Goodman's discussion is a brief one and, on our reading at least, leaves open a number of questions. Should we admit as relevant any considerations beyond observed patterns of inference, such as properties of long-run convergence to the truth, and so on? To whom does 'we' refer to when Goodman speaks of 'the particular...inferences we actually make and sanction'? Experts? The human population at large? Should we be circumscribing the class of relevant inferences to those judgments that one might want to call 'considered'? These are important matters to settle. Indeed, a certain combination of answers to these—entailing that the justification of normative theories of reasoning hinges entirely on their ability to systematise 'immediate and untutored' inferential dispositions observed in the general population—notoriously led Cohen (1981) to endorse the startling claim that, since normative and descriptive models are borne of the very same data set, behavioural evidence is in principle incapable of establishing human irrationality. For further discussion of this general topic, see for instance Stich (1990, Ch. 4), Stein (1996, Ch. 5), Stanovich (1999, Ch. 1), and Thagard (1982).²³

Although neither Allais nor Goodman draw the connection, a potential justification for the evidential relevance of experimental data in normative theory building can perhaps be sought in the literature on the Condorcet Jury Theorem and related results.²⁴ This theorem tells us that, under certain conditions, the probability that a majority verdict, with regards to a particular matter, in a group of n minimally reliable people casting yes/no votes on a particular question converges to 1 as n tends to infinity, converging more quickly the greater the individual reliabilities. Furthermore, majority reliability reaches significant levels, even given very limited individual reliability, for fairly modest group sizes. Of course, the issue of interest does not quite fit that specific model: while the expression of Allais preferences can be arguably interpreted as a 'vote' against the normative adequacy of Independence, the expression of preferences consonant with this principle can hardly be interpreted as a vote in favour of it.

Finally, while this section has focused on the issue of the bearing of descriptive decision theory on its normative counterpart, it should be noted that there has been some discussion of the converse direction of influence. Both Guala (2000) and Starmer (2005) have argued that the development of descriptive theories of choice has been guided by a bias towards retaining a core of principles taken to be normatively adequate. In the case of decision making under risk these are essentially the transitivity component of Weak Order and Stochastic Dominance, which are satisfied according to the vast majority of non-SEU theories that have been developed to date.²⁵ Starmer claims to find an argument justifying this practise in a well-known paper by Friedman and Savage (1952). This line of thought, which Starmer takes issue with, proceeds from the assumption that bona fide principles of rationality would be evident as such to most subjects and that decision makers will accordingly behave in line with them.

6 Further reading

While the philosophical literature on the topic remains rather sparse, there is no shortage of first-rate summaries in the economics and psychology literatures. For thorough presentations of the technical results referred to in Section 1, see Fishburn (1970, Ch. 14) or the slightly less detailed Kreps (1988, Ch. 9). Ch. 3 of Joyce (1999) is also helpful here. Regarding the literature on Independence specifically, discussed in Section 2, see Machina (1987), Starmer (2000) and Weber & Camerer (1987). Regarding the issue of probabilistic belief specifically,

discussed in Section 3, see Camerer & Weber (1992), Etner *et al* (2012), Gilboa & Marinacci (2013), Machina & Siniscalchi (2014), and Trautmann & van de Kuilen (2015). A number of broader surveys cover both the above issues, and some. These include most notably Camerer (1995) and the excellent Sugden (2004).

Finally, for a clear and detailed historical account of the development of the experimental literature on decision-making, see Heukelom (2014).

7 Related entries

belief, formal representations of | decision theory | economics, philosophy of | preferences | probabilities, imprecise | rational choice, normative: expected utility

8 Acknowledgments

I am extremely grateful to John Quiggin, Peter Wakker, Paul Weirich and Greg Wheeler for helpful comments on an earlier draft of this entry. Any remaining omissions, errors or inaccuracies remain, of course, my own responsibility.

¹Disclaimer: The literature on the topic is absolutely vast, comprising thousands of papers spread across a number of academic disciplines. Accordingly, only an extremely selective glimpse at some of the relevant work can be provided. One particularly noteworthy omission is the issue of sequential, aka ‘dynamic’, choice, which would have required an entire further entry to present in a satisfactory manner. Space considerations have also precluded a discussion of a number of important empirical phenomena associated with choice behaviour, including so-called framing effects and preference reversals. For further details regarding these issues, as well as the points touched upon in this entry, the interested reader is referred to the excellent surveys listed below in Section 6.

²Each will consist in those points $(P(x_1), P(x_3))$ such that $u(x_1)P(x_1) + u(x_2)(1 - P(x_1) - P(x_3)) + u(x_3)P(x_3)$ is equal to some constant expected utility value c . Rearranging this, one obtains the linear equation $P(x_3) = [u(x_2) - u(x_1)]/[u(x_3) - u(x_4)]P(x_1) + c - u(x_2)$, the slope of the line being independent of $P(x_1)$ and $P(x_2)$.

³In what follows, the various methodological and statistical issues pertaining to the assessment of the models on offer have been glossed over. Hey (2014) provides a useful discussion, from this angle, of a number of important studies. Salient issues include the choice of methods of preference elicitation, the stochastic specification of the models, and a set of more fundamental problems connected with hypothesis testing and model selection.

⁴It is well known that Independence, a principle governing *single* choices, can be recovered from a number of conditions on temporal *sequences* of choices. Violations of the principle therefore immediately raise the question of where further to pin the blame. Due to space considerations, discussion of these issues has been omitted. The interested reader is referred to Section 4 of Sugden’s (2004) excellent survey article.

⁵This presentation of Allais’ paradoxes as respectively pertaining to the Common Consequence and Common Ratio principles is owed to Mongin (2009).

⁶Regarding the empirical adequacy of Stochastic Dominance, Tversky and Kahneman (1986) observe, perhaps not entirely surprisingly, significant violations of the principle when (but only when) the dominance relation is not immediately evident. It should be noted that a number of early Mixture-Independence-free models did not satisfy this requirement. These included most famously the weighted probability model of Handa (1977) and Kahneman & Tversky’s closely related *Prospect Theory* (1979), which essentially restricted the violations tolerated by Handa’s model to those cases that are ‘transparent’. Stochastic Dominance *is* however a feature of the successors of these proposals, discussed below.

⁷It should be noted that this general pattern of fanning out that must be imposed by WU to accommodate the local pattern of fanning out characteristic of Allais preferences lands the model in trouble elsewhere. Indeed further data has been collected that is consistent with an additional local pattern of *fanning in* towards the northwest of the triangle (see for instance Camerer 1989). Other Betweenness-satisfying models of probabilistically sophisticated preferences can better accommodate this data by allowing for *mixed fanning* patterns. These include Chew’s *Implicit Utility* (Chew 1989), which generalises WU, and Gul’s *Disappointment Aversion* (Gul 1991).

⁸Other models in this category include *Quadratic Utility* (Chew, Epstein and Segal 1991) and *Lottery Dependent Utility* (Becker & Sarin 1987), which generalises WU.

⁹In fn. 7 above, it was pointed out that WU is unable to accommodate certain mixed fanning patterns suggested by the data. This is not the case with RDU, which shares in this respect the flexibility noted in relation with Implicit Utility and Disappointment Aversion.

¹⁰This formulation is due to Quiggin (1992).

¹¹For more on this, see Sections 10.1 to 10.5 of Wakker (2010).

¹²It should be noted that Cumulative Prospect Theory can also be formulated as a sign-dependent generalisation of CEU. In this formulation, the functional form involves both a positive and a negative capacity. Wakker & Tversky (1993) offer an axiomatisation that notably restricts

the Sure-Thing principle to sets of acts that are not only comonotonic but also *cosigned*, i.e. such that for all acts f, g in that set, there exists no state s such that $f(s) > 0$ but $g(s) < 0$, where 0 designates the reference point.

¹³Gilboa & Schmeidler (1989) provided their initial axiomatic characterisation in a special case of the Savage setup, the Anscombe-Aumann (AA) framework (Anscombe & Aumann 1963), which has not been discussed in this entry. In this setting, additional structure is imposed on the set of outcomes, which are modeled as assignments of objective probabilities to the members of a set of further alternatives. Gilboa & Schmeidler's axioms were more recently translated into the classic Savage model by Ghirardato *et al* (2003), who also present an analogue of the AA-style axiomatisation of CEU due to Schmeidler (1989).

¹⁴It is not clear, however, what value of this parameter would correspond to ambiguity-neutrality. As Etner *et al* (2012) note, the most likely candidate, the value $1/2$, does not lead one to recover SEU, which is a plausible model for ambiguity neutral preferences.

¹⁵MEU subsumes Wald's (1950) much earlier *Minimax* approach to decision making under *complete* suspension of judgment regarding objective probabilities (the approach is so called, because it involves minimisation of maximum loss). Wald's model corresponds to the special case in which $\Gamma = \mathcal{P}$. α -MEU subsumes Hurwicz's (1951) *Generalized Bayes-Minimax* approach in the same manner. See Ch. 13 of Luce & Raiffa (1957) for an overview of these precursors and their axiomatic foundations.

¹⁶This proposal is equivalent, for pairwise choices, to Fishburn's (1989) *Skew-Symmetric Additive Utility theory*.

¹⁷As Loomes & Sugden (1987) later note, in the special case of preferences over independent distributions over three outcomes, Regret Theory yields indifference curves that are consistent with those generated by WU, given certain parameter values, fanning out from a point located in the southwestern quadrant, outside the probability triangle.

¹⁸Regret Theory proposes to replace SEU with a more general model that accommodates certain behaviours that can be plausibly interpreted as violating Transitivity or Independence. It is suggestive, however, of an alternative position which interprets these behaviours in a manner that renders them *consistent* with the standard model. This position involves '*fine-graining*' the description of the outcomes involved, so as to distinguish

- (a) two menu-*dependent* outcomes (e.g. its being the case that x_1 rather than x_3 versus its being the case that x_1 rather than x_2 ; denote these respectively by $\langle x_1, x_3 \rangle$ and $\langle x_1, x_2 \rangle$) with corresponding menu-*independent* utilities (e.g. $u(\langle x_1, x_3 \rangle)$ versus $u(\langle x_1, x_2 \rangle)$),

where Regret Theory only recognises

- (b) one menu-*independent* outcome (e.g. its being the case that x_1) with two menu-*dependent* utilities (e.g. $M(x_1, x_3)$ versus $M(x_1, x_2)$).

Note that the general strategy of fine-graining outcome descriptions to defuse apparent counterexamples to expected utility maximisation has been carried out with respect to a variety of cases and in a variety of manners. For more on this, see Buchak (2013, Ch. 4) and works referenced therein, including Broome (1991), Weber (1998) and Weirich (1986).

¹⁹This kind of line of thinking, it appears, may have precedents in the decision theoretic literature. Jallais *et al* (2008) claim that the appeal to 'expert behaviour' was also pivotal in the emergence of SEU as a popular normative theory, finding such a reliance in both Bernoulli's writings on the St Petersburg game and Condorcet's critique thereof.

²⁰Savage did not elaborate on the rationale underpinning this test, but one view prevalent among some of his contemporaries was that bona fide principles of rationality are self-evident. Morgenstern (1979, p. 180) provides a case in point, claiming that: '[SEU] is normative in the sense that the theory is 'absolutely convincing' which implies that men will act accordingly. If they deviate from the theory, an explanation of the theory and of their deviation will cause them to readjust their behavior.'

²¹For illustrations, see Section 4 of the entry titled '[Normative Theories of Rational Choice: Expected Utility](#)'.

²²This is not to say, of course, that SEU is uncontroversial in philosophy. The completeness requirement has been a long-time suspect, in terms of normative, in addition to empirical, adequacy. Anand (2009) offers a rather critical review of the case for transitivity of rational preferences, as does Maher (1993, Ch. 2), who remains more optimistic. The normative adequacy of the assumption of probabilistic belief has also been under scrutiny for many years. It has been claimed for instance that probabilistic beliefs have insufficient structure (i) to be able to model a doxastic state that rationality requires to reflect not only the balance, but also what Keynes (1921) called the 'weight', of evidence (see e.g. Joyce 2005) or again (ii) to be able to model certain types of rationally permissible agnosticism, as suggested by van Fraassen's 'Water/Wine' problem (van Fraassen 1989). See the entry titled '[Imprecise probabilities](#)' for further various endorsements of multiple prior models in the philosophical literature, as well as the entry titled '[Formal representations of belief](#)'. For cases in which probabilistic belief holds, in cases of decision making under risk, philosophical thinking has tended to take a more orthodox line. One notable recent exception here is Buchak (2013), who has defended the normative adequacy of RDU. See Chs. 5–7 of her book. Doubts regarding Independence are also voiced in an accessible handbook entry by McClennen (2009). See Maher (1993, Ch. 3) for a less pessimistic survey of the considerations.

²³Unsurprisingly, analogous questions crop up in the recent experimental philosophy literature, in which it is suggested that empirical frequencies of intuitions regarding the applicability of a term may serve as defeasible evidence in favour of, or against, a particular theory of how the term ought to be applied. There, for instance, one also finds the recommendation that attention be focused on the judgments of the subset of the population that one would independently identify as 'experts' (see e.g. Hales 2006).

²⁴The connection has however been noted in the experimental philosophy literature, by Sytma & Livengood 2014 and Talbot 2014. See List & Pettit (2011 Ch. 4) for accessible presentation of the theorem, which is also discussed in the entry titled '[Social Choice Theory](#)'.

²⁵Though see Mongin (2009) for an alternative perspective on the development of alternatives to SEU.

REFERENCES

- Allais, M., 1953a, “Le Comportement de l’Homme Rationnel devant le Risque: Critique des Postulats et Axiomes de l’Ecole Américaine”, *Econometrica* 21: 503–546.
- , 1953b, “Fondements d’une Théorie Positive des Choix Comportant un Risque et Critique des Postulats et Axiomes de L’Ecole Américaine”, *Econométrie, Colloques Internationaux du CNRS XL*: 257–332; page reference is to the translation titled “The Foundations of a Positive Theory of Choice Involving Risk and a Criticism of the Postulates and Axioms of the American School” in M. Allais & O. Hagen (ed.) 1979, pp. 27–145.
- Allais, M., and O. Hagen (ed.), 1979, *Expected Utility Hypotheses and the Allais Paradox*, Dordrecht: Reidel.
- Anscombe, F.J., and R.J. Aumann, 1963, “A Definition of Subjective Probability”, *Annals of Mathematics and Statistics* 34: 199–205.
- Anand, P., 2009, “Rationality and Intransitive Preference Foundations for The Modern View”, in P. Anand, P. K. Pattanaik and C. Puppe (ed.), *The Handbook of Rational and Social Choice*, Oxford: Oxford University Press, pp. 156–172.
- Becker, S., and F. Brownson, 1964, “What Price Ambiguity? Or the Role of Ambiguity in Decision-Making”, *Journal of Political Economy* 72: 62–73.
- Becker, J.L., and R.K. Sarin, 1987, “Lottery Dependent Utility”, *Management Science* 33: 1367–1382.
- Bewley, T.F., 1986, “Knightian Decision Theory: Part I”, Cowles Foundation discussion paper no. 807.
- Bleichrodt, H., and P.P. Wakker, 2015, “Regret Theory: A Bold Alternative to the Alternatives”, *The Economic Journal* 125: 493–532.
- Broome, J. (1991). *Weighing Goods: Equality, Uncertainty, and Time*. Oxford: Basil Blackwell.
- Buchak, L., 2013, *Risk and Rationality*, Oxford: Oxford University Press.
- Camerer, C., 1989, “An Experimental Test of Several Generalized Utility Theories”, *Journal of Risk and Uncertainty* 2: 61–104.
- , 1995, “Individual Decision Making”, in J. Kagel and A.E. Roth (ed.), *Handbook of Experimental Economics*, Princeton, NJ: Princeton University Press, pp. 587–703.
- Camerer, C., and T.-H. Ho, 1994, “Violations of the Betweenness Axiom and Nonlinearity in Probability”, *Journal of Risk and Uncertainty* 8: 167–96.
- Camerer, C., and M. Weber, 1992, “Recent Developments in Modeling Preferences: Uncertainty and Ambiguity”, *Journal of Risk and Uncertainty* 5: 325–370.
- Chew, S.H., 1983, “A Generalization of the Quasilinear Mean with Applications to the Measurement of Income Inequality and Decision Theory Resolving the Allais Paradox”, *Econometrica* 51: 1065–1092.
- , 1989, “Axiomatic Utility Theories with the Betweenness Property”, *Annals of Operations Research* 19: 273–298.
- Chew, S.H., L.G. Epstein and U. Segal, 1991, “Mixture Symmetry and Quadratic Utility”, *Econometrica* 59 (1): 139–163.
- Chew, S.H., and K. MacCrimmon, 1979, “Alpha-Nu Choice Theory: A Generalization of Expected Utility Theory”, Working Paper 669, University of British Columbia.
- Chew, S.H., and P. Wakker, 1996, “The Comonotonic Sure-Thing Principle”, *Journal of Risk and Uncertainty* 12: 5–27.
- Cohen, L.J., 1981, “Can Human Irrationality be Experimentally Demonstrated?”, *Behavioral and Brain Sciences* 4: 317–370.
- de Finetti, B., 1937, “La Prévision: Ses Lois Logiques, ses Sources Subjectives”, *Annales de l’Institut Henri Poincaré* 7: 1–68.
- Ellsberg, D., 1961, “Risk, Ambiguity, and the Savage Axioms”, *Quarterly Journal of Economics* 75 (4): 643–669.
- , 2001, *Risk, Ambiguity, and Decision*, New York & London: Garland.
- Etner, J., M. Jeleva and J.M. Tallon, 2012, “Decision Theory under Ambiguity”, *Journal of Economic Surveys* 26: 234–270.

- , 1989, “Non-Transitive Measurable Utility for Decision under Uncertainty”, *Journal of Mathematical Economics* 18: 187–207.
- , 1970, *Utility Theory for Decision Making* (Publications in Operations Research, No. 18), New York: John Wiley and Sons.
- Friedman, M., and L.J. Savage, 1952, “The Expected-Utility Hypothesis and the Measurability of Utility”, *Journal of Political Economy* 60 (6): 463–474.
- Galaabaatar, T., and E. Karni, 2013, “Subjective Expected Utility with Incomplete Preferences”, *Econometrica* 81: 255–284.
- Ghirardato, P., F. Maccheroni, M. Marinacci and M. Siniscalchi, 2003, “A Subjective Spin on Roulette Wheels”, *Econometrica* 71 (6): 1897–1908.
- Gilboa, I., 1987, “Expected Utility with Purely Subjective Non-Additive Probabilities”, *Journal of Mathematical Economics* 16 (1): 65–88.
- Gilboa, I., and M. Marinacci, 2013, “Ambiguity and the Bayesian Paradigm”, in D. Acemoglu, M. Arellano, and E. Dekel (ed.), *Advances in Economics and Econometrics: Theory and Applications* (Tenth World Congress of the Econometric Society), New York: Cambridge University Press.
- Gilboa, I., and D. Schmeidler, 1989, “Maxmin Expected Utility with a Non-Unique Prior”, *Journal of Mathematical Economics* 18: 141–153.
- Goodman, N., 1965, *Fact, Fiction and Forecast*, Indianapolis, Ind.: Bobbs-Merrill.
- Grant, S., 1995, “Subjective Probability without Monotonicity: Or How Machina’s Mom May also Be Probabilistically Sophisticated”, *Econometrica* 63 (1): 159–189.
- Green, J.R., and B. Jullien, 1988, “Ordinal Independence in Nonlinear Utility Theory”, *Journal of Risk and Uncertainty* 1: 355–387.
- Guala, F., 2000, “The Logic of Normative Falsification: Rationality and Experiments in Decision Theory”, *Journal of Economic Methodology* 7 (1): 59–93.
- Gul, F., 1991, “A Theory of Disappointment Aversion”, *Econometrica* 59: 667–686.
- Hales, S.D., 2006, *Relativism and the Foundations of Philosophy*, Cambridge, MA: MIT Press.
- Handa, J., 1977, “Risk, Probabilities and a New Theory of Cardinal Utility”, *Journal of Political Economy* 85: 97–122.
- Harless, D.W., and C. Camerer, 1994, “The Predictive Utility of Generalized Expected Utility Theories”, *Econometrica* 62: 1251–1289.
- Heukelom, J., 2014, *Behavioral Economics: A History*, Cambridge: Cambridge University Press.
- Hey, J.D., 2014, “Choice Under Uncertainty: Empirical Methods and Experimental Results”, in M. Machina and K. Viscusi (ed.) *Handbook of the Economics of Risk and Uncertainty, Volume 1*, Amsterdam: Elsevier, pp. 809–850.
- Hurwicz, L., 1951, “Some Specification Problems and Applications to Econometric Models”, *Econometrica* 19: 343–344.
- Jallais, S., and P.-C. Pradier, 2005, “The Allais Paradox and its Immediate Consequences for Expected Utility Theory”, in P. Fontaine and R. Leonard (ed.) *The Experiment in the History of Economics*, London: Routledge, pp. 25–49.
- Jallais, S., P.-C. Pradier, and D. Teira, 2008, “Facts, Norms and Expected Utility Functions”, *History of the Human Sciences* 21 (2): 45–62.
- Joyce, J.M., 1999, *The Foundations of Causal Decision Theory*, Cambridge: Cambridge University Press.
- , 2005, “How Probabilities Reflect Evidence”, *Philosophical Perspectives* 19: 153–178.
- Kahneman, D., and A. Tversky, 1979, “Prospect Theory: An Analysis of Decision under Risk”, *Econometrica* 47: 263–291.
- Keynes, J.M., 1921, *A Treatise on Probability*, London: Macmillan.
- Klibanoff, P., M. Marinacci, and S. Mukerji, 2005, “A Smooth Model of Decision Making Under Uncertainty”, *Econometrica* 6: 1849–1892.
- Kreps, D.M., 1988, *Notes on the Theory of Choice*, Boulder, CO: Westview Press.

- List, C., and P. Pettit, 2011, *Group Agency: The Possibility, Design, and Status of Corporate Agents*, Oxford: Oxford University Press.
- Loomes, G., and R. Sugden, 1982, "Regret Theory: An Alternative Theory of Rational Choice under Uncertainty", *Economic Journal* 92: 805–824.
- , 1987, "Some implications of a More General Form of Regret Theory", *Journal of Economic Theory* 41: 270–287.
- Luce, R.D., and H. Raiffa, 1957, *Games and Decisions: Introduction and Critical Survey*, New York: Wiley.
- Machina, M.J., 1987, "Choice under Uncertainty: Problems Solved and Unsolved", *Journal of Economic Perspectives* 1 (1): 121–154.
- Machina, M.J., and D. Schmeidler, 1992, "A more Robust Definition of Subjective Probability", *Econometrica* 60 (4): 745–780.
- Machina, M.J., and M. Siniscalchi, 2014, "Ambiguity and Ambiguity Aversion", in M. Machina and K. Viscusi (ed.) *Handbook of the Economics of Risk and Uncertainty, Volume 1*, Amsterdam: Elsevier, pp. 729–807.
- MacCrimmon, K.R., 1968, "Descriptive and Normative Implications of the Decision-Theory Postulates", in K. Borch and J. Mossin (ed.), *Risk and Uncertainty*, New York: St. Martins Press, pp. 3–32.
- MacCrimmon, K.R., and S. Larsson, 1979, "Utility Theory: Axioms versus 'Paradoxes'", in M. Allais and O. Hagen 1979, pp. 333–409.
- Maher, P., 1993, *Betting on Theories*, Cambridge: Cambridge University Press.
- Marschak, J., 1951, "Why 'Should' Statisticians and Businessmen Maximize 'Moral Expectation'", *Proceedings of the Second Berkeley Symposium on Mathematical Statistics and Probability*, Berkeley: University of California Press, pp. 493–506.
- May, K.O., 1954, "Intransitivity, Utility, and the Aggregation of Preference Patterns", *Econometrica* 22: 1–13.
- McClennen, E.F., 2009, "The Normative Status of the Independence Principle", in P. Anand, P. K. Pattanaik and C. Puppe (ed.) *The Handbook of Rational and Social Choice*, Oxford: Oxford University Press, pp. 140–155.
- Mongin, P., 2009, "Duhemian Themes in Expected Utility Theory", *French Studies In The Philosophy Of Science* (Volume 276 of *Boston Studies In The Philosophy Of Science*), pp. 303–357.
- , 2014, "Le Paradoxe d'Allais. Comment Lui Rendre sa Signification Perdue?", *Revue Économique* 65: 743–779.
- Morgenstern, O., 1979, "Some Reflections on Utility", in M. Allais & O. Hagen 1979, pp. 175–184.
- Nau, R., 2006, "The Shape of Incomplete Preferences", *Annals of Statistics* 34: 2430–2448.
- Ok, E.A., P. Ortoleva, and G. Riella, 2008, "Incomplete Preferences under Uncertainty: Indecisiveness in Beliefs vs. Tastes", *Econometrica* 80 (4): 1791–1808.
- Quiggin, J., 1982, "A theory of anticipated utility", *Journal of Economic Behavior and Organization* 3: 323–343.
- , 1992, *Generalized Expected Utility Theory: The Rank Dependent Model*, Dordrecht: Kluwer.
- Ramsey, F.P., 1931, "Truth and Probability", in R.B. Braithwaite (ed.) *The Foundations of Mathematics and other Logical Essays*, New York: Harcourt and Brace, pp. 156–198.
- Regenwetter, M., J. Dana, and C.P. Davis-Stober, 2011, "Transitivity of Preferences", *Psychological Review* 118 (1): 42–56.
- Savage, L.J., 1954, *The Foundations of Statistics*, New York: Wiley, Second Edition.
- Schmeidler, D., 1986, "Integral Representation without Additivity", *Proceedings of the American Mathematical Society* 97: 255–261.
- , 1989, "Subjective Probability and Expected Utility without Additivity", *Econometrica* 57: 571–587.
- Seidenfeld, T., M.J. Schervish, and J.B. Kadane, 1995, "A Representation of Partially Ordered Preferences", *Annals of Statistics* 23: 2168–2217.
- Slovic, P., and A. Tversky, 1974, "Who Accepts Savage's Axiom?", *Systems Research and Behavioral Science* 19: 368–373.
- Stanovich, K.E., 1999, *Who Is Rational? Studies of Individual Differences in Reasoning*, Mahwah, NJ: Lawrence Erlbaum Associates.
- Starmer, C., 2005, "Normative Notions in Descriptive Dialogues", *Journal of Economic Methodology* 12 (2): 277–289.

- , 2000, “Developments in Non-Expected Utility Theory: The Hunt for a Descriptive Theory of Choice under Risk”, *Journal of Economic Literature* 38: 332–382.
- Stein, E., 1996, *Without Good Reason: The Rationality Debate in Philosophy and Cognitive Science*, Oxford: Clarendon Press.
- Stich, S., 1990, *The Fragmentation of Reason*, Cambridge, Mass.: MIT Press.
- Sugden, R., 1993, “An Axiomatic Foundation For Regret Theory”, *Journal of Economic Theory* 60: 159–180.
- , 2004, “Alternatives to Expected Utility: Foundations”, in Barberà, S., P.J. Hammond and C. Seidl (ed.), *Handbook of Utility Theory: Volume 2 Extensions*, Boston, MA: Springer, pp. 685–755.
- Sytsma, J., and J. Livengood, 2014, *The Theory and Practice of Experimental Philosophy*, Peterborough, ON: Broadview Press.
- Talbot, B., 2014, “Why so Negative? Evidence Aggregation and Armchair Philosophy”, *Synthese* 191 (16): 3865–3896.
- Thagard, P., 1982, From the Descriptive to the Normative in Psychology and Logic, *Philosophy of Science* 49 (1): 24–42.
- Trautmann, S.T., and G. van de Kuilen, 2015, “Ambiguity Attitudes”, in G. Keren & G. Wu (ed.), *The Wiley Blackwell Handbook of Judgment and Decision Making*, Oxford: Blackwell, 89–116.
- Tversky, A., 1969, “Intransitivity of Preferences”, *Psychological Review* 76: 31–48.
- Tversky, A., and D. Kahneman, 1986, “Rational Choice and the Framing of Decisions”, *The Journal of Business* 59 (4): 251–278.
- , 1992, “Cumulative Prospect Theory: An Analysis of Attitudes towards Uncertainty and Value”, *Journal of Risk and Uncertainty* 5: 297–323.
- van de Kuilen, G., and P. Wakker, 2006, “Learning in the Allais Paradox”, *Journal of Risk and Uncertainty* 33: 155–164.
- van Fraassen, B., 1989, *Laws and Symmetry*, Oxford: Oxford University Press.
- von Neumann, J., and O. Morgenstern, 1947, *Theory of Games and Economic Behavior*, Princeton: Princeton University Press, 2nd Edition.
- Wald, A., 1950, *Statistical Decision Functions*. New York: John Wiley and Sons.
- Wakker, P.P., 2010, *Prospect Theory: For Risk and Ambiguity*, Cambridge: Cambridge University Press.
- , 1989, “Continuous Subjective Expected Utility with Nonadditive Probabilities”, *Journal of Mathematical Economics* 18: 1–27.
- Wakker, P.P. and A. Tversky, 1993, “An Axiomatization of Cumulative Prospect Theory”, *Journal of Risk and Uncertainty* 7 (2): 147–175.
- Weber, M., 1998, “The Resilience of the Allais Paradox”, *Ethics*, 109 (1): 94–118.
- Weber, M., and C. Camerer, 1987, “Recent Developments in Modelling Preferences under Risk”, *OR Spektrum* 9 (3): 129–151.
- Weirich, P., 1986, “Expected Utility and Risk”, *British Journal for the Philosophy of Science*, 37: 419–442.
- Zappia, C., 2016, “Daniel Ellsberg and the Validation of Normative Propositions”, *Oeconomia* 6 (1): 57–79.