Decision Theory

Decision theory provides a general, mathematically rigorous account of decision-making under uncertainty. The subject includes *rational choice theory*, which seeks to formulate and justify the normative principles that govern optimal decision-making, and *descriptive choice theory*, which aims to explain how human beings actually make decisions. Within both these areas one may distinguish *individual decision theory*, which concerns the choices of a single agent with specific goals and knowledge, *game theory*, which deals with interactions among individuals. This entry will focus on rational choice theory for the single agent, but some descriptive results will be mentioned in passing.

Decision Problems. It is standard to portray decision makers as facing choices among *acts* that cause desirable or undesirable *consequences* when performed in various *states of the world*. Acts characterize those aspects of the world that an agent can directly control. States specify contingencies beyond her control that might influence the consequences of acts. Each combination of an act A and state S fixes a unique consequence $A(S)$ that describes the result of doing A in S. When there are only finitely many acts and states the decision situation can be represented as a matrix

	S_1	$S_2\ldots$	S_n
A_1	$A_1(S_1)$	$A_1(S_2)\ldots$	$A_1(S_n)$
A_2	$A_2(S_1)$	$A_2(S_2)\ldots$	$A_2(S_n)$
A_3	$A_3(S_1)$	$A_3(S_2)\ldots$	$A_3(S_n)$
$\vdots$	$\vdots$	$\vdots$	$\vdots$
A_m	$A_m(S_1)$	$A_m(S_2)\ldots$	$A_m(S_n)$

The agent decides the row, the world decides the column, and these together determine the consequence.

In any well-formed decision problem (a) the value of each consequence is independent of the act and state that bring it about, (b) each consequence is sufficiently detailed to settle every matter about which the agent intrinsically cares, (c) neither acts nor states have any value except as a means for producing consequences, and (d) the agent will not believe that she has the ability to causally influence which state obtains. When these conditions are met, the agent's goals and values affect her decision only via her desires for consequences, and her beliefs influence her choice via her uncertainty about which state obtains. The agent will use her beliefs about states to select an act that provides the best means for securing a desirable consequence.

For theoretical purposes, it is useful to idealize the decision setting by assuming that the repertoire of actions is quite rich. Specifically, for each consequence *c* there is a *constant act* [*c*] that produces *c* in every state of the world, and, for any acts *A* and *B*, and any disjunction of states *E*, there is a "mixed act" $A_E \cup B_{\sim E}$ that produces *A*'s consequence when *E* holds and *B*'s consequence when ~*E* holds. While real agents will typically be unable to realize such recherché prospects as these, imagining that decision makers have attitudes toward them often helps one determine which realistic acts should be performed.

This model applies to "one-choice" decisions made at a specific time. Early decision theorists believed that sequences of decisions made over time could be reduced to one-shot decisions among contingency plans, or *strategies*, but this view now has few adherents. The topic of dynamic decision-making lies beyond the scope of this essay. For relevant discussions see Hammond [1988], McClennen [1990], and Joyce [1995].

Subjective Expected Utility. The central goal of rational choice theory is to identify the conditions under which a decision maker's beliefs and desires *rationalize* the choice of an action. According to the 'standard model' of decision-theoretic rationality, an action is rational just in case, relative to the agent's beliefs and desires, it has the highest subjective expected utility of any available option. This subjective expected utility theory (*SEU*) has its roots in the work of Blaise Pascal, Daniel Bernoulli, Vilfredo Pareto, and F.P. Ramsey, and finds its fullest expression in Leonard Savage's *Foundations of Statistics* (1954). According to *SEU*, a rational agent's basic desires can be represented by a *utility function* ***u*** that assigns a real number $\boldsymbol{u}(c)$ to each consequence c. The value of $\boldsymbol{u}(c)$ measures the degree to which c would satisfy the agent's desires and promote her aims. Likewise, the agent's beliefs can be characterized by a *subjective probability function* ***P*** whose values express the agent's subjective degrees of confidence, or *credences*, in the states of the world. ***P*** is assumed to be unique, and ***u*** is unique once the choice of a unit and a zero for measuring utilities are fixed. Given ***P*** and ***u***, the expected utility of each act A is a weighted average of the utilities of its consequences, so that $\boldsymbol{Exp}_{\boldsymbol{P},\boldsymbol{u}}(A) = \Sigma_{i=1}\, \boldsymbol{P}(S_i)\boldsymbol{u}(A(S_i))$. According to the core doctrine of *SEU*, the choice of an act is rational only if it maximizes the chooser's *subjective expected utility*, so that $\boldsymbol{Exp}_{\boldsymbol{P},\boldsymbol{u}}(A) \geq \boldsymbol{Exp}_{\boldsymbol{P},\boldsymbol{u}}(B)$ for all acts B. This should not be taken to suggest that the agent sees herself as maximizing expected utility, or even that she has the concept of expected utility. *SEU* does not propose expected utility maximization as a decision procedure, but as a way of assessing the results of such procedures. Rational decision makers merely act *as if* they maximize subjective expected utility; they need not explicitly do so.

Representation of Rational Preference. A central challenge for *SEU* is to find a principled way of characterizing credences and utilities. Following the lead of Ramsey [1931], the standard solution involves proving a *representation theorem* that shows how an agent's beliefs about states and desires for outcomes are related to her all-things-considered *preferences* for acts. The agent is assumed to make three sorts of comparative evaluations between acts: she might strictly prefer *A* to *B*, written $A > B$, weakly prefer *A* to *B*, $A \succsim B$, or be indifferent between them, $A \approx B$. These relations hold, respectively, just in case the agent judges that, on balance, *A* will do more than, at least as much as, or exactly as much as, *B* will to satisfy her desires and promote her aims. The totality of such evaluations is the agent's *preference ranking*.

Early decision theorists, motivated by a misguided scientific methodology, thought of preferences as operationally defined in terms of overt choices, so that, by definition, an agent prefers *A* to *B* if and only if she will incur a cost to choose *A* over *B*. Even though this sort of behaviorism remains firmly ensconced in some areas of economics, it has been widely and effectively criticized, Sen [1977] and Joyce [1999]. In the end, preferences are best thought of as subjective judgments of the comparative merits of actions as promoters of desirable outcomes. While such judgments are closely tied to overt choice behavior, the relationship between the two is nowhere near as direct and unsophisticated as behaviorism suggests.

The representation theorem approach seeks to justify *SEU* by (i) imposing a system of axiomatic constraints on preference rankings, (ii) arguing that these express requirements of rationality, and then (iii) proving that any preference ranking that satisfies the axioms can be associated with a probability ***P*** and a utility ***u*** such that each of $A >, \succsim, \approx B$ hold if and only if, respectively, $\boldsymbol{Exp}_{\boldsymbol{P},\boldsymbol{u}}(A) >, \geq, = \boldsymbol{Exp}_{\boldsymbol{P},\boldsymbol{u}}(B)$. An agent whose preferences can be 'represented' in this way evaluates acts *as if* she were aiming to maximize expected utility relative to ***P*** and ***u***.

Frame Invariance. All versions of *SEU* share a common set of core principles. The first says that *logically* equivalent redescriptions of prospects should not alter preferences.

SEU_1 *Frame Invariance*. The evaluation of an act should not depend on how its consequences happen to be described.

People often violate this constraint. Consider the following two decision 'framings' due to Shafir and Tversky [1995]:

- You receive $300, and are then given a choice between getting another $100 for sure or getting $200 or $0 depending on the toss of a fair coin.

- You receive $500, but are then forced to choose between returning $100 for sure or returning $200 or $0 depending on the toss of a fair coin.

Since both decisions offer a sure $400 or a fifty-fifty chance of $300 or $500, SEU_1 requires agents to make the same choice in each case (though it does not tell them which choice to make). As it turns out, most people make the "safe" choice in the first case and take the sure $400, but they make “risky” choice in the second case by taking the fifty-fifty gamble. Cognitive psychologists attribute this violation of SEU_1 to the following two irrational tendencies of human decision makers.

> *Divergence from Status Quo*. People are more concerned with incremental *gains* and *losses*, seen as changes in the status quo, than with total well-being or overall happiness.

> *Asymmetrical Risk Aversion*. People eschew risk when pursuing *gains*, but to seek risk when avoiding losses.

Under the first description, where the status quo is $300, people see themselves as trying to secure an additional gain, and so opt for the ‘safe’ alternative. Under the second description, where the status quo is $500, people see themselves avoiding losses, and so incline toward the “risky” choice. These divergent attitudes are quite irrational given that the options are effectively identical.

Value Independence. The second principle requires each act to have a value that depends only on the values and probabilities of the outcomes it might cause.

SEU_2 *Value Independence*. If the agent prefers *A* to *B* in a decision where *C* is not an option, then she should still prefer *A* to *B* even if *C* is an option, provided that *C*’s inclusion does not provide any information about state probabilities.

Apparent counterexamples to SEU_2 as a requirement of rationality always involve violations of the proviso. For example, Luce and Raiffa [19xx] discuss a diner who, thinking he is in a greasy spoon, prefers salmon to steak, but then orders steak when told that snails are on the menu. SEU_2 is vindicated by the observation that the availability of snails provides the diner with evidence for thinking that he is in fine restaurant, and this alters his views about the comparative merits of the salmon and steak. Other common violations of SEU_2 are clearly irrational. For example, Redelmeier and Shafir [1995] show that physicians are *less* likely to prescribe ibuprofen to patients in pain when they have the option of prescribing the inferior drug piroxicam than when piroxicam is unavailable. While this sort of behavior does not discredit SEU_2 as a normative principle, it does show that it is inaccurate as a *description* of human behavior.

Ordering. The third principle rules out "preference cycles" in which $A \gtrsim B$, $B \gtrsim C$, but $A > C$, and it requires that the preference ranking be *complete* in the sense that exactly one of $A > B$, $A \approx B$ or $B > A$ always hold.

SEU_3 *Ordering*. Preference rankings *completely order* the set of acts.

Though some have disputed anti-cyclicality, and Fishburn [1991] has even developed an acyclic decision theory, the prohibition against cycles remains among the most widely accepted principles of rational preference. On views that equate preferences and choices, preference cycles are irrational because they leave the agent open to exploitation as a 'money pump': she will freely trade *C* for *B* and *B* for *A*, and then pay a fee to exchange *C* for *A*, thereby getting nothing for something. Even if choice is not equated with preference, cycles are still problematic. Many seemingly rational cycles treat preferences as partial, rather than "all-things-considered," evaluations. For instance, one might prefer an expensive shirt to a moderately-priced one on the basis of style, and prefer the moderately priced shirt to a cheap shirt on the basis of durability, but prefer the cheap shirt to the expensive one on the basis of price. Here what seems to be a rational preference cycle is really a failure to integrate considerations of style, durability, and price into an "all-things-considered" value judgment.

Failures of evaluative discrimination can also seem to generate rational preference cycles. Suppose a vinophile, who cares only about how his wine tastes, cannot taste any difference between wine *A* and wine *B*, or between wine *B* and wine *C*, but can taste that *C* is better than *A*. It is tempting to think that the vinophile should be indifferent between *A* and *B* and between *B* and *C*, but should prefer *C* to *A*. A clearer understanding of the situation shows that this is incorrect. A person should only be indifferent between prospects when he lacks any reason, on balance, for preferring one to the other. Our vinophile, however, has reason to favor *B* over *A* since *B* is indistinguishable in taste from a wine superior to *A*. He also has reason to favor *C* over *B* since *B* is indistinguishable from a wine inferior to *C*. Properly speaking, then, the vinophile is not indifferent between *A* and *B* or between *B* and *C*: his preference run $B \succsim A$, $C \succsim B$, and $C > A$, but neither $A \approx B$ nor $B \approx C$ is true.

One might worry that our vinophile's reasons seem insufficient to justify strict preferences. It would, for example, be silly for him to pay anything to trade a bottle of *A* for a bottle of *B* (unless he could convert the latter into a bottle of *C* for a small enough fee). While this is a legitimate concern, it tells against completeness rather than anti-cyclicality. When an agent cannot precisely discriminate the qualities of prospects on which his evaluations depend, or when these qualities are themselves vague or indeterminate, his preference ranking will be incomplete: for certain options, all three of $A > B$, $A \approx B$ and $B > A$ will fail. Sometimes both $A >_{\sim} B$ and $B >_{\sim} A$ will fail as well, in which case the agent has no views about the comparative merits of A and B. Alternatively, as in our vinophile example, the agent might determinately *weakly* prefer B to A even though he neither strictly prefers B to A nor is indifferent between them. So, while $A \approx B$ and $B > A$ each entail $B >_{\sim} A$, the latter is consistent with the falsity of both $A \approx B$ and $B > A$. In addition to indeterminacy or vagueness in values, incompleteness in preferences can arise via an imprecision in credences. In both sorts of cases it can be perfectly rational to have an incomplete preference ranking.

One response to these considerations, which is advocated in Levi [1980], Jeffrey [1983] and Kaplan [1980], is to construe SEU_3's completeness clause as a requirement of *coherent extendibility*. Instead of asking an agent to completely order acts, one demands merely that there be at least one complete preference ranking (usually there will be many) that satisfies all other requirements of rationality, and that agrees with the agent's preferences whenever she has definite preferences. One then represents vague or indeterminate preferences by giving up the idea that the agent's attitudes can be modeled by a single probability/utility pair (given a unit and zero for utility). Rather, there will be a *representing set* $\boldsymbol{R}$ of ($\boldsymbol{P}$, $\boldsymbol{u}$) pairs that agree with the agent's preferences in the sense that, for any options A and B, each of $A >, >_{\sim}, \approx B$ hold if and only if, respectively, $\boldsymbol{Exp}_{\boldsymbol{P},\boldsymbol{u}}(A) >, \geq, = \boldsymbol{Exp}_{\boldsymbol{P},\boldsymbol{u}}(B)$ holds for *every* ($\boldsymbol{P}$, $\boldsymbol{u}$) pair in $\boldsymbol{R}$. Act A is unambiguously choiceworthy only if maximizes expected utility relative to every ($\boldsymbol{P}$, $\boldsymbol{u}$) pair in $\boldsymbol{R}$. It is *admissible* when it maximizes expected utility relative to some such pair. There is no generally accepted procedure for handling situations where no admissible act is unambiguously choiceworthy. Some theorists would say that the agent's beliefs and desires are too indefinite to justify any choice as rational. Others, most notably Isaac Levi [1980], maintain that principles of decision making that outrun expected utility maximization come into play in this situation. For example, Levi allows agents to decide among admissible options using 'maximin', i.e., by selecting the act whose worst consequence is at least as good as the worst consequence of any alternative.

Comparative Probability. The next principle of *SEU* forges a link between rational preference and rational belief. A *wager* on event E is an act of the form $[c]_E \cup [d]_{\sim E}$ where $[c] > [d]$. Such a wager produces the desirable consequence c in every state consistent with E and the undesirable consequence d in every state consistent with $\sim E$. Intuitively, a person should prefer such a wager more strongly the more likely she takes E to be. More precisely, given any events E and F, $[c]_E \cup [d]_{\sim E}$ should be preferred to $[c]_F \cup [d]_{\sim F}$ exactly if E as more probable than F. The following axiom is meant to ensure that this is so.

> SEU_4 *Comparative Probability*. Assuming $[c] > [d]$, if the agent prefers $[c]_E \cup [d]_{\sim E}$ to $[c]_F \cup [d]_{\sim F}$, she must also prefer $[c^*]_E \cup [d^*]_{\sim E}$ to $[c^*]_F \cup [d^*]_{\sim F}$ for any consequences such that $[c^*] > [d^*]$.

SEU_4 can seem implausible when the values of consequences vary with the world's state. Suppose, for example, that c and d are monetary fortunes that one might have in ten years, say c = $500,000 and d = $400,000. Let E and F be hypotheses about the cumulative rate of inflation over the decade: E puts the figure at sixty percent, while F puts it at ten percent. Even if one regards E as the more probable hypothesis, one might still prefer to wager on *F* since one's fortune will be worth more in the event that *F* is true.

There are two standard responses to this problem. Savage maintained that decision problems of this sort, in which the values of consequences depend on states, are ill-formed. He argued any such problem could be transformed into a well-formed decision by a suitable subdivision of consequences. In the example above, *c* would be split into c_1 = '$500,000 after cumulative inflation of sixty percent', and c_2 = '$500,000 after cumulative inflation of ten percent'. Alternatively, one might opt for a *state-dependent utility* theory, which replaces SEU_4 by a weaker condition and allows the values of consequences of vary with states. See, Karni [1993] and Schervish, et. al., [1990] for details.

Independence and the Sure-thing Principle. The most controversial tenet of *SEU* is the *Independence* axiom.

SEU_5 *Independence*. Preference among acts that have exactly the same consequences when *E* is false should depend exclusively on what happens when *E* is true. If $A_E \cup C_{\sim E}$ is preferred to $B_E \cup C_{\sim E}$ for some act *C*, then $A_E \cup D_{\sim E}$ is preferred to $B_E \cup D_{\sim E}$ for all acts *D*.

To illustrate, consider the following act *types*, where *c*, *d*, *c** and *d** are known consequences, and *x* ranges over possible consequences.

	S_1	S_2	S_3
A_x	c	d	x
B_x	c^*	d^*	x

SEU_5 says that an agent's preference between A_x and B_x should not depend on x's value. More generally, it requires agents to have well-defined *conditional* preferences: A is preferred to B in the event of E just in case $A_E \cup C_{\sim E} > B_E \cup C_{\sim E}$ for some (hence any) C.

SEU_5 has the following intuitive consequence:

> *Sure-Thing Principle*: Let $E_1, E_2, \ldots, E_n$ be mutually exclusive, collectively exhaustive events. If A is weakly preferred to B conditional on each E_i, then A is weakly preferred to B simpliciter. Moreover, if A is strictly preferred to B conditional on some event that is not judged certainly false, then A is strictly preferred to B.

Independence and the sure-thing principle have been quite controversial. Some apparent failures of SEU_5 arise in ill-formed decision problem whose states are not independent of acts. For example, imagine a man who has to drive home from a party where alcohol is being served. He likes to drink, but worries about getting home safely. Suppose he frames his decision like this:

	Car Accident	No Accident
Drink	-100	1
Teetotal	-101	0

Since the consequences of drinking are better that those of refraining both in the event of an accident and otherwise, it looks as if the sure-thing principle advocates drinking, which is clearly bad advice given that drinking increases the probability of an accident. Problems of this sort led Richard Jeffrey [1965] to develop an *evidential* version of decision theory in which Independence is only valid for decisions in which acts provide no evidence about the occurrence of any state. Reflections on 'Newcomb problems', in which acts and states are causally independent but evidentially correlated, led *causal* decision theorists like Robert Stalnaker [1981], Allan Gibbard and William Harper [1978], and Brian Skyrms [1980] to insist that the two principles be restricted to decisions in which the choice of an act has no causal influence over states.

The most famous objections to SEU_5 are the "paradoxes" of Maurice Allais [1953] and Daniel Ellsberg [1961], which seem to show that *SEU* rules out certain rational attitudes toward risk and uncertainty. An act involves *risk* when the agent knows the objective probabilities with which its consequences will obtain. It involves *uncertainty* when the agent's information allows a range of possible risk profiles for consequences. SEU_5 entails that, insofar as decision-making is concerned, all legitimate considerations of risk and uncertainty are fully captured in expected utilities. The Allais and Ellsberg paradoxes suggest, to the contrary, that risk and uncertainty are *non-separable* quantities: one cannot express them as weighted averages of their values conditional on disjoint events. If this is correct, then an agent need not have any fixed preference between the act types A_x and B_x because x's value might provide information about the relative *risk* or uncertainty of the two options, and this information might justifiably influence the agent's preferences.

The Allais paradox envisions an agent who chooses between *A* and *A** and then between *B* and *B** (with the know probabilities listed).

	0.10	**0.01**	**0.89**

A	$1,000,000	$1,000,000	$1,000,000
*A**	$5,000,000	$0	$1,000,000
B	$1,000,000	$1,000,000	$0
*B**	$5,000,000	$0	$0

Empirical studies show that people systematically violate Independence when presented with such choices. They 'play it safe' and select *A* over *A** in the first choice, but favor the 'riskier' option *B** over *B** in the second. The standard rationale for these choices assumes (a) that there is more risk involved in choosing *A** over *A* than there is in choosing *B** over *B*, and (b) that it is rational minimize this risk even when doing so violates Independence.

Ellsberg's Paradox shows something similar with respect to judgments of uncertainty. Suppose a ball will be drawn at random from an urn that holds 30 red balls, and 60 white or blue balls in unknown proportion. One chooses between *C* and *C** and then between *D* and *D**.

	Red	**White**	**Blue**
C	$100	$0	$0
*C**	$0	$100	$0
D	$100	$0	$100
*D**	$0	$100	$100

Here most people prefer *C* to *C** and *D** to *D.* Interestingly, when gains are replaced by losses, people still violate Independence, but both choices are reversed. People thus seem to prefer risk to uncertainty when they have something to gain, but prefer uncertainty to risk when they have something to lose. Those who regard Ellsberg's paradox as a counterexample to *SEU* maintain that such non-separable preferences for risk over uncertainty or uncertainty over risk are entirely rational.

Some proponents of *SEU*, see Broome [1990], respond by arguing that the consequences in the Allais and Ellsberg paradoxes are *underdescribed*. For example, the standard pattern of preferences in Allais can be rationalized by noting that, when the 0.01 event occurs, agents who choose *A** over *A* may feel regret (because they passed up a sure thing), while those who choose *B** over *B** will feel no regret (because they probably would have ended up with nothing anyhow). For such agents, the decision matrix really looks like this:

	0.10	**0.01**	**0.89**
A	$1,000,000	$1,000,000	$1,000,000
*A**	$5,000,000	$0 with regret	$1,000,000
B	$1,000,000	$1,000,000	$0
*B**	$5,000,000	$0 without regret	$0

Likewise, if an agent feels uneasy when gains ride on uncertain prospects (or losses ride on risky prospects), then the correct description of the Ellsberg problem is this:

	Red	**White**	**Blue**
C	$100	$0	$0
*C**	$0 with uneasiness	$100 with uneasiness	$0 with uneasiness
D	$100 with uneasiness	$0 with uneasiness	$100 with uneasiness
*D**	$0	$100	$100

If these matrices accurately describe the decisions, then neither the Allais or Ellsberg paradoxes provide a genuine counterexample to SEU_5.

These sorts of 'rationalizing' response are weakened by their dependence on substantive assumptions about the psychology of risk, uncertainty and regret that are not universally accepted, see Loomes and Sudgen [1982] and Weber [1998]. An alternative is to argue that the usual preferences in the Allais and Ellsberg paradoxes are simply irrational. In Allais, for example, agents assume that the disparity in risk between *A* and *A** exceeds the disparity in risk between *B* and *B**. This may be a mistake. One way to determine differences in risk is to consider the costs of insuring against the incremental risk one incurs by trading one option for another. Someone who switches from *A** to *A* in Allais can offset this risk by purchasing an insurance policy that pays out $1,000,000 contingent on the 0.01 event. Notice, however, that the risk incurred by switching from *B** to *B* can be offset by the very same policy. Since a single policy eliminates both risks there is reason to think that the actual change in risk is the same in each case. Similar things can be said about the Ellsberg choosers, who implicitly assume that they decrease their uncertainty more by switching from *C** to *C* than they do by switching from *D** to *D*. So, if one measures disparities in risk or uncertainty by the costs of insuring against it, then *SEU* is safe from the Allais and Ellsberg examples.

Opponents of *SEU* will, of course, deny that risks should be measured by the costs of insuring against them. Ultimately, the issue will be resolved by the development of a convincing measure of risk. While there is a well-known theory of risk *aversion* within *SEU* there is no universally accepted method for quantifying risk itself. The best work in this area, which builds on Rothschild and Stiglitz [1970], suggests that risk is indeed separable.

Alternatives to SEU. While subjective expected utility theory remains firmly ensconced as the ‘standard model’ of rational decision making for individuals, a number of alternatives have been developed. One kind of approach seeks to relax Independence while preserving most other aspects of *SEU*. Especially noteworthy here is the ‘generalized expected utility analysis’ of Machina [1982], and the ‘weighted utility model’ of Chew and MacCrimmon [1979]. Alternatively, one can reject maximizing conceptions of rationality altogether, and see decision making as matter of ‘satisficing’ relative to fixed constraints. For example, Gigerenzer, et al. [1999] seek to replace the single all-purpose prescription to maximize expected utility by an ecological model of rationality in which decision makers employ a set of simple, highly localized decision heuristics. These heuristics efficiently generate choices that produce desirable consequences in the contexts where they tend to be employed, but they can go badly awry when used in out of context. For discussion of further ‘non-standard’ decision theories see Sugden [2004].

Interesting though these alternatives are, none has seriously challenged the normative status of *SEU*. Though highly idealized, and far from adequate as a description of human behavior, *SEU* remains out best overall account of rational decision-making.

JAMES M. JOYCE

REFERENCES

Allais, Maurice "Le Comportement de l'Homme Rationnel devant le Risque: Critique des postulats et axiomes de l'École Americaine," 1953, *Econometrica*.

Broome, John. *Weighing Goods*. Oxford: Blackwell, 1991.

Ellsberg Daniel. "Risk, ambiguity, and the Savage axioms," *Quarterly Journal of Economics* **75** (1961): 643-669.

Fishburn, Peter C. 1991. "Nontransitive Preferences in Decision Theory," *Journal of Risk and Uncertainty* **4** (1991): 113-34.

Gibbard, Allan and William Harper. "Counterfactuals and Two Kinds of Expected Utility." In *Foundations and Applications of Decision Theory*, edited by C. Hooker, J. Leach, and E. McClennen. Dordrecht: Reidel, 1978.

Gigerenzer, G, Todd, P., and the ABC Research Group. *Simple Heuristics that Make us Smart*. New York: Oxford University Press, 1999.

Hammond, Peter. "Consequentialist Foundations for Expected Utility," Theory and Decision, **25** (1988): 25-78.

Jeffrey, Richard. *The Logic of Decision*, New York, McGraw Hill, 1965. 2nd revised edition, Chicago: University of Chicago Press, 1983.

Joyce, James M. *The Foundations of Causal Decision Theory*. Cambridge, UK: Cambridge University Press, 1999.

Kaplan, Mark. "Decision Theory as Philosophy," *Philosophy of Science* **50** (1983): 549-77.

Karni, E. "Subjective Expected Utility Theory with State-Dependent Preferences," *Journal of Economic Theory*, **60** (1993): 428–438.

Levi, Isaac. *The Enterprise of Knowledge*. Cambridge, MA: MIT Press, 1980.

Loomes, Graham, and Robert Sugden. "Regret theory: An Alternative Theory of Rational Choice Under Uncertainty," *Economic Journal*, **92** (1982): 805-824.

Luce, R. Duncan and Howard Raiffa. *Games and Decisions*. John Wiley & Sons, New York, 1957

Machina, Mark. "'Expected Utility' Analysis Without the Independence Axiom," *Econometrica* **50** (1982): 277-323.

McClennen, Edward. *Rationality and Dynamic Choice: Foundational Explorations*. Cambridge, UK: Cambridge University Press, 1990.

Ramsey, Frank P. "Truth and Probability," in *The Foundations of Mathematics and Other Logical Essays*, edited by Richard Braithwaite. London: Kegan Paul, 1931.

Redelmeier, D., & Shafir, E. Medical Decision Making in Situations that Offer Multiple Alternatives. *Journal of the American Medical Association* **273, #4** (1995): 302-305.

Rothschild, M. and J. E. Stiglitz ."Increasing Risk: I. A Definition," *Journal of Economic Theory* **2** (1970): 225-43.

Savage, Leonard J. *The Foundations of Statistics*. New York: Wiley, 1954. 2nd revised edition, New York: Dover, 1972.

Schervish, M., T. Seidenfeld, and J. Kadane. State-Dependent Utilities. *Journal of the American Statistical Association* **85** (1990): 840-47.

Sen, Amartya. "Rational Fools: A Critique of the Behavioral Foundations of Economic Theory," *Philosophy and Public Affairs* **6** (1977): 317–344.

Shafir, E., & Tversky, A. "Decision making," in Edward Smith and Daniel Osherson, editors, *An Invitation to Cognitive Science, Second Edition (Volume 3: Thinking).* Cambridge, MA: MIT Press, 1995: 77-100.

Skyrms, Brian. *Causal Necessity*. New Haven: Yale University Press, 1980.

Stalnaker, Robert. “Letter to David Lewis, May 21, 1972,” in Robert Stalnaker, William Harper and Glen Pearce, editors, *IFs: Conditionals, Belief, Decision, Chance, and Time*. University of Western Ontario Series in Philosophy of Science, **15** (1981) Dordrecht & Boston: Reidel, 1981.

Sugden, Robert. “Alternatives to Expected Utility: Formal Theories,” in Peter Hammond, Salvador Barberà and Christian Seidl, editors, *Handbook of Utility Theory, Vol. 2* Kluwer Academic 2004.

Weber, Michael. “The Resilience of the Allais Paradox,” *Ethics* **109** (1998): 94-118.