Three Paradoxes in Utility Theory

Background: The expected utility hypothesis is the hypothesis in economics that the utility of an agent facing uncertainty is calculated by considering utility in each possible state and constructing a weighted average. The weights are the agent's estimate of the probability of each state. The expected utility is thus an expectation in terms of probability theory. To determine utility according to this method, the decision maker must rank their preferences according to the outcomes of various decision options. According to the theory, if someone prefers A to B and B to C, then weights for the weighted average must exist such that she is indifferent between receiving B outright and gambling - with the specified weights - between A and C.

Daniel Bernoulli (1738) gave the earliest known written statement of this hypothesis as a way to resolve the St. Petersburg Paradox. In the expected utility theorem, v. Neumann and Morgenstern proved that any "normal" preference relation over a finite set of states can be written as an expected utility. (Therefore, it is also called von-Neumann Morgenstern utility.) Von Neumann and Morgenstern published this in their Theory of Games and Economic Behavior in 1944. It is important because it was developed shortly after the Hicks-Allen "ordinal revolution" of the 1930's, and it revived the idea of cardinal utility in economic theory. Economics has not resolved whether (and in what cases) utility is cardinal or ordinal.

A related concept is the certainty equivalent of a gamble. The more risk-averse a person is, the more he will be prepared to pay to eliminate risk, for example accepting $1 instead of a 50% chance of $3, even though the expected value of the latter is more. People may be risk-averse or risk-loving depending on the amounts involved and on whether the gamble relates to becoming better off or worse off; this is a possible explanation for why the same person may buy both an insurance policy and a lottery ticket. However, expected utility as a descriptive model of decisions under risk has in recent years been replaced by more sophisticated variants that take irrational deviations from the expected utility model into account; compare Prospect theory and the general article on Behavioral finance.

Framing: Framing means the manner in which a rational choice problem has been presented. Amos Tversky and Daniel Kahneman have shown that framing can affect the outcome (ie. the choices one makes) of choice problems, to the extent that several of the classic axioms of rational choice do not hold. Tversky and Kahneman (1981) demonstrated systematic reversals of preference when the same problem is presented in different ways, for example in the 'Asian disease' problem. Participants were asked to "imagine that the U.S. is preparing for the outbreak of an unusual Asian disease, which is expected to kill 600 people. Two alternative programs to combat the disease have been proposed. Assume the exact scientific estimate of the consequences of the programs are as follows." The first group of participants were presented with a choice between two programs:

- Program A: "200 people will be saved"

- Program B: "there is a one-third probability that 600 people will be saved, and a two-thirds probability that no people will be saved"

72 percent of participants preferred program A (the remainder, 28 percent, opting for program B). The second group of participants were presented with the choice between:

- Program C: "400 people will die"
- Program D: "there is a one-third probability that nobody will die, and a two-third probability that 600 people will die"

In this decision frame, 78 percent preferred program D, with the remaining 22 percent opting for program C. However, programs A and C, and programs B and D, are effectively identical in accordance with von-Neumann's expected utility hypothesis, in which the value of the outcome of an event is multiplied by the probability of its occurrence. A change in the decision frame between the two groups of participants produced a preference reversal, with the first group preferring program A/C and the second group preferring B/D. Ed Zelinsky has shown that framing effects can explain some observed behaviors of legislators.

Framing biases affecting investing, lending, borrowing decisions make one of the themes of behavioral finance. Preference reversals and other associated phenomena are of wider relevance within behavioural economics, as they contradict the predictions of rational choice, the basis of traditional economics.

Preference Reversals over Uncertain Outcomes: Starting with studies such as Lichtenstein & Slovic (1971), it was discovered that subjects sometimes exhibit signs of preference reversals with regards to their certainty equivalents of different lotteries. Specifically, when eliciting certainty equivalents, subjects tend to value "p bets" (lotteries with a high chance of winning a low prize) lower than "$ bets" (lotteries with a small chance of winning a large prize). When subjects are asked which lotteries they prefer in direct comparison, however, they frequently prefer the "p bets" over "$ bets." Many studies have examined this "preference reversal," from both an experimental (e.g., Plott & Grether, 1979) and theoretical (e.g., Holt, 1986) standpoint, indicating that this behavior can be brought into accordance with neoclassical economic theory under certain assumptions.

St. Petersburg paradox

The St. Petersburg paradox relates to probability theory and decision theory. It is based on a particular (theoretical) lottery game (sometimes called *St. Petersburg Lottery*) that leads to a random variable with infinite expected value, i.e. infinite expected payoff, but would nevertheless be considered to be worth only a very small amount of money. The St. Petersburg paradox is a classical situation where a naïve decision theory (which takes only the expected value into account) would recommend a course of action that no (real) rational person would be willing to take. The paradox can be resolved when the decision model is refined via the notion of marginal utility or by taking into account the finite resources of the participants. Some economists claim that the paradox is resolved by noting that one simply cannot buy that which is not sold (and sellers would not produce a lottery whose expected loss to them were unacceptable).

The paradox is named from Daniel Bernoulli's presentation of the problem and his solution, published in 1738 in the *Commentaries of the Imperial Academy of Science of Saint Petersburg* (Bernoulli 1738). However, the problem was invented by Daniel's cousin Nicolas Bernoulli who first stated it in a letter to Pierre Raymond de Montmort of 9 September 1713.

The Paradox

In a game of chance, you pay a fixed fee to enter, and then a fair coin will be tossed repeatedly until a "tail" first appears, ending the game. The "pot" starts at 1 dollar and is doubled every time a "head" appears. You win whatever is in the pot after the game ends. Thus you win 1 dollar if a tail appears on the first toss, 2 dollars if on the second, 4 dollars if on the third, 8 dollars if on the fourth, etc. In short, you win 2^{k-1} dollars if the coin is tossed k times until the first tail appears.

What would be a fair price to pay for entering the game? To answer this we need to consider what would be the average payout: With probability 1/2, you win 1 dollar; with probability 1/4 you win 2 dollars; with probability 1/8 you win 4 dollars etc. The expected value is thus

$$E = \frac{1}{2}\cdot 1 + \frac{1}{4}\cdot 2 + \frac{1}{8}\cdot 4 + \frac{1}{16}\cdot 8 + \cdots$$
$$= \frac{1}{2} + \frac{1}{2} + \frac{1}{2} + \frac{1}{2} + \cdots$$
$$= \sum_{k=1}^{\infty} \frac{1}{2} = \infty$$

This sum diverges to infinity; on average you can expect to win an infinite amount of money when playing this game. So according to traditional notions, and assuming that the casino has infinite resources, no matter how much you pay to enter you can expect to come out ahead in the long run, the idea being that on the very rare occasions when a large payoff comes along, it will far more than repay however much money you have paid to play. Yet in published descriptions of the paradox, e.g. (Martin, 2004), many people expressed disbelief in the result. Martin quotes Ian Hacking as saying "few of us would pay even $25 to enter such a game" and says most commentators would agree.

Solutions of the paradox: There are different approaches for solving the “paradox”.

Expected utility theory : The classical resolution of the paradox involved the explicit introduction of a utility function, an expected utility hypothesis, and the presumption of diminishing marginal utility of money. In Daniel Bernoulli's own words:

> The determination of the value of an item must not be based on the price, but rather on the utility it yields…. There is no doubt that a gain of one thousand ducats is more significant to the pauper than to a rich man though both gain the same amount.

Using a utility function, e.g., as suggested by Bernoulli himself, the logarithmic function $u(x) = \ln(x)$ (known as "log utility"), the expected utility of the lottery (for simplicity assuming an initial wealth of zero) becomes finite:

$$EU = \sum_{k=1}^{\infty} p_k \cdot u(2^{k-1}) = \sum_{k=1}^{\infty} \frac{\ln(2^{k-1})}{2^k} = \ln 2 = u(2) < \infty$$

(This particular utility function suggests that the lottery is as useful as 2 dollars.)

Before Daniel Bernoulli published, in 1978, another Swiss mathematician, Gabriel Cramer, found already parts of this idea (also motivated by the St. Petersburg Paradox) in stating that

> the mathematicians estimate money in proportion to its quantity, and men of good sense in proportion to the usage that they may make of it.

He demonstrated in a letter to Nicolas Bernoulli that a square root function describing the diminishing marginal benefit of gains can resolve the problem. However, unlike Daniel Bernoulli, he did not consider the total wealth of a person, but only the gain by the lottery.

The solution by Cramer and Bernoulli, however, is not yet completely satisfying, since the lottery can easily be changed in a way that the paradox reappears: To this aim, we just need to change the game so that it gives the (even larger) payoff e^{2^k}. Again, the game should be worth an infinite amount. More generally, one can find a lottery that allows for a variant of the St. Petersburg paradox for every unbounded utility function, as was first pointed out by (Menger, 1934).

There are basically two ways of solving this generalized paradox, which is sometimes called the *Super St. Petersburg paradox*:

- We can take into account that a casino would only offer lotteries with a finite expected value. Under this restriction, it has been proved that the St. Petersburg paradox disappears as long as the utility function is concave, which translates into the assumption that people are (at least for high stakes) risk averse [Compare (Arrow, 1974)].

- It is possible to assume an upper bound to the utility function. This does not mean that the utility function needs to be constant at some point, an example would be $u(x) = 1 - e^{-x}$.

Recently, expected utility theory has been extended to arrive at more behavioral decision models. In some of these new theories, as in Cumulative Prospect Theory, the St. Petersburg paradox again appears in certain cases, even when the utility function is concave, but not if it is bounded (Rieger and Wang, 2006).

Probability weighting : Nicolas Bernoulli himself proposed an alternative idea for solving the paradox. He conjectured that people will neglect unlikely events[4]. Since in the St. Petersburg lottery only unlikely events yield the high prizes that lead to an infinite expected value, this could resolve the paradox. The idea of probability weighting resurfaced much later in the work on prospect theory by Daniel Kahneman and Amos Tversky. However, their experiments indicated that, very much to the contrary, people tend to overweight small probability events. Therefore the proposed solution by Nicolas Bernoulli is nowadays not considered to be satisfactory.

One can't buy what isn't sold: Some economists attempt to resolve the paradox by arguing that, even if an entity had infinite resources, the game would never be offered. If the lottery represents an infinite expected gain to the player, then it also represents an infinite expected loss to the host. No one could be observed paying to play the game because it would never be offered. As Paul Samuelson describes the argument:

> Paul will never be willing to give as much as Peter will demand for such a contract; and hence the indicated activity will take place at the equilibrium level of zero intensity. (Samuelson,1960)

Finite St. Petersburg lotteries

The classical St. Petersburg lottery assumes that the casino has infinite resources. This assumption is often criticized as unrealistic, particularly in connection with the paradox, which involves the reactions of ordinary people to the lottery. Of course, the resources of an actual casino (or any other potential backer of the lottery) are finite. More importantly, the expected value of the lottery only grows logarithmically with the resources of the casino. As a result, the expected value of the lottery, even when played against a casino with the largest resources realistically conceivable, is quite modest. This can be seen from a consideration of the finite variant of the St. Petersburg lottery:

If the total resources of the casino are *W* dollars, then the maximum payoff and therefore the maximum number of rounds is "capped", and the expected value of the lottery becomes

$$E = \sum_{k=1}^{L} p_k 2^{k-1} + 2^{L-1} \sum_{k=L+1}^{\infty} p_k$$

$$= \sum_{k=1}^{L} \frac{1}{2} + 2^{L-1} \left(1 - \left(1 - \frac{1}{2^L}\right)\right) = \frac{L+1}{2},$$

where $L = 1 + \text{floor}(\log_2(W))$. L is the maximum number of times the casino can play before it can no longer cover the next bet. The function $\log_2(W)$ is the base-2 logarithm of W, which can be computed as $\log(W)/\log(2)$ in any other base. The floor function gives the greatest integer less than or equal to its argument. The logarithm function becomes infinite as its argument becomes infinite, but does so very, very slowly. This logarithmic growth is the inverse behavior of exponential growth.

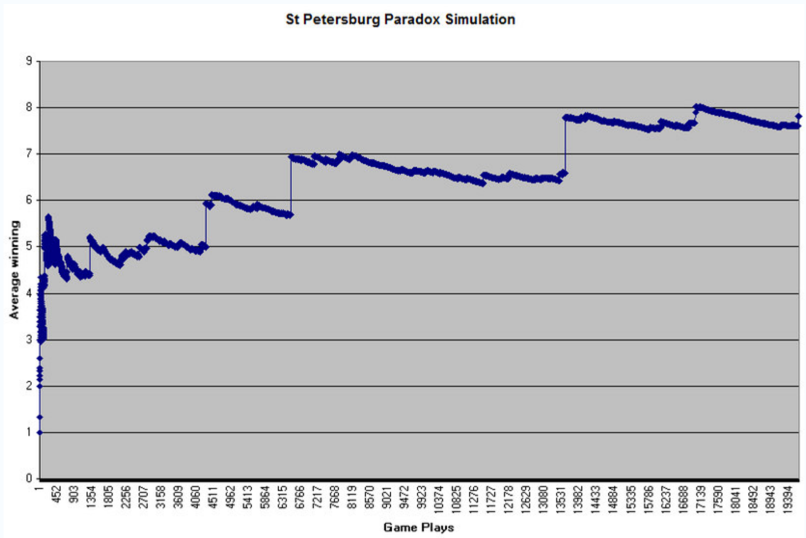

A typical graph of average winnings over one course of a St. Petersburg Paradox lottery shows how occasional large payoffs lead to an overall very slow rise in average winnings. After 20,000 gameplays in this simulation the average winning per lottery was just under 8 dollars. The graph encapsulates the paradox of the lottery: The overall upward slope in the average winnings graph shows that average winnings tend upward to infinity, but the slowness of the rise in average winnings (a rise that becomes yet slower as gameplay progresses) indicates that a tremendously huge number of lottery plays will be required to reach average winnings of even modest size.

The following table shows the expected value of the game with various potential backers and their bankroll:

Backer	*Bankroll*	*Expected value of lottery*
Friendly game	$64	$3.50
Millionaire	$1,050,000	$10.50
Billionaire	$1,075,000,000	$15.50

Bill Gates	$51,000,000,000 (2005)	$18.00
U.S. GDP	$11.7 trillion (2004)	$22.00
World GDP	$40.9 trillion (2004)	$23.00
Googolnaire	10^{100}	$166.50

Notes: The slightly higher bankrolls for "millionaire" and "billionaire" allow a final round of play at those levels; otherwise for each, the maximum payout would be half as much and the expected value would be $0.50 less. A "googolnaire" is a hypothetical person worth a googol dollars (10^{100}).

An average person might not find the lottery worth even the modest amounts in the above table, arguably showing that the naive decision model of the expected return causes the same problems as for the infinite lottery, however the possible discrepancy between theory and reality is far less dramatic.

The assumption of infinite resources can produce other apparent paradoxes in economics. A reference may also be made to martingale (roulette system) and gambler's ruin.

Iterated St. Petersburg lottery

Players may assign a higher value to the game when the lottery is repeatedly played. This can be seen by simulating a typical series of lotteries and accumulating the returns, compare the illustration

Allais paradox

The Allais paradox, more neutrally described as the *Allais problem*, is a choice problem designed by Maurice Allais to show an inconsistency of actual observed choices with the predictions of expected utility theory. The problem arises when comparing participants' choices in two different experiments, each of which consists of a choice between two gambles, A and B. The payoffs for each gamble in each experiment are presented in the table below.

It has been found that presented with the choice between 1A and 1B, most people choose 1A. Presented with the choice between 2A and 2B, most people choose 2B. This is inconsistent with expected utility. The point is that both gambles give the same outcome 89% of the time (the top row; $1 million for Gamble 1, and zero for Gamble 2), so, in expected utility, these equal outcomes should have no effect on the desirability of the gamble. If the 89% 'common consequence' is disregarded, both gambles offer the same choice; a 10% chance of getting $5 million and 1% chance of getting nothing as against an 11% chance of getting $1 million. (It may help to re-write the payoffs. 1A offers an 89% chance of winning 1 million and a 11% chance of winning 1 million, where the 89% chance is irrelevant. 2B offers an 89% chance of winning nothing, a 1% chance of winning nothing, and a 10% chance of winning 5 million, with the 89% chance of nothing disregarded. Hence, choice 1A and 2A should now clearly be seen as the same choice, and 1B and 2B as the same choice).

Experiment 1				**Experiment 2**			
Gamble 1A		**Gamble 1B**		**Gamble 2A**		**Gamble 2B**	
Winnings	**Chance**	**Winnings**	**Chance**	**Winnings**	**Chance**	**Winnings**	**Chance**
$1 million	100%	$1 million	89%	Nothing	89%	Nothing	90%
		Nothing	1%	$1 million	11%		
		$5 million	10%			$5 million	10%

Allais presented his paradox as a counterexample to the independence axiom (also known as the "sure thing principle" of expected utility theory. Independence means that if an agent is indifferent between simple lotteries L_1 and L_2, the agent is also indifferent between L_1 mixed with an arbitrary simple lottery L_3 with probability p and L_2 mixed with L_3 with the same probability p. Violating this principle is known as the "common consequence" problem (or "common consequence" effect). The idea of the common consequence problem is that as the prize offered by L_3 increases, L_1 and L_2 become consolation prizes, and the agent will modify preferences between the two lotteries so as to minimize risk and disappointment in case they do not win the higher prize offered by L_3.

Difficulties such as this gave rise to a number of alternatives to, and generalizations of, the theory, notably including prospect theory, developed by Daniel Kahneman and Amos Tversky, weighted utility (Chew) and rank-dependent expected utility by John Quiggin. The point of these models was to allow a wider range of behavior than was consistent with expected utility theory.

Also relevant here is the framing theory by Daniel Kahneman and Amos Tversky. Identical items will result in different choices if presented to agents differently (i.e. a surgery with a 70% survival rate vs. a 30% chance of death) However, the main point Allais wishes to make, is that the independence axiom of expected utility theory may not be a necessary axiom. The independence axiom states that two identical outcomes within a gamble should be treated as irrelevant to the analysis of the gamble as a whole. However, this overlooks the notion of complementarities, the fact your choice in one part of a gamble may depend on the possible outcome in the other part of the gamble. In the above choice, 1B, there is a 1% chance of getting nothing. However, this 1% chance of getting nothing also carries with it a great sense of disappointment if you were to pick that gamble and lose, knowing you could have won with 100% certainty, if you had chosen 1A. This feeling of disappointment however, is contingent on the outcome in the other portion of the gamble (i.e. the feeling of certainty). Hence, Allais argues that it is not possible to evaluate portions of gambles or choices independently of the other choices presented, as the independence axiom requires, and thus is a poor judge of our rational action(1B cannot be valued independently of 1A as the independence or sure thing principle requires of us). We don't act irrationally when choosing 1A and 2B, rather expected utility theory is not robust enough to capture such "bounded rationality" choices that in this case arise because of complementarities.

Mathematical Proof of Inconsistency

Using the values above and a utility function of u(W), where W is wealth, we can demonstrate exactly how the paradox manifests.

Because the typical individual prefers 1A to 1B and 2B to 2A, we can write conclude that the expected utilities of the preferred is greater than the expected utilities of the second choices, or,

$$1.00U(1m) > 0.89U(1m) + 0.01U(0) + 0.1U(5m)$$

$$0.89U(0) + 0.11U(1m) < 0.9U(0) + 0.1U(5m)$$

We can rewrite the latter equation as,

$$0.11U(1m) < 0.01U(0) + 0.1U(5m)$$

$$1U(1m) - 0.89U(1m) < 0.01U(0) + 0.1U(5m)$$

$$1U(1m) < 0.01U(0) + 0.1U(5m) + 0.89U(1m)$$

Which contradicts the first bet which shows the player prefers the sure thing over the gamble.

Ellsberg paradox

The Ellsberg paradox is a paradox in decision theory and experimental economics in which people's choices violate the expected utility hypothesis. It is generally taken to be evidence for ambiguity aversion. The paradox was popularized by Daniel Ellsberg, although a version of it was noted considerably earlier by John Maynard Keynes (Keynes 1921, pp.75-76, p.315, ft.2).

The paradox: Suppose you have an urn containing 30 red balls and 60 other balls that are either black or yellow. You don't know how many black or yellow balls there are, but that the total number of black balls plus the total number of yellow balls equals 60. The balls are well mixed so that each individual ball is as likely to be drawn as any other. You are now given a choice between two gambles:

Gamble A	Gamble B
You receive $100 if you draw a red ball	You receive $100 if you draw a black ball

Also you are given the choice between these two gambles (about a different draw from the same urn):

Gamble C	Gamble D
You receive $100 if you draw a red or yellow ball	You receive $100 if you draw a black or yellow ball

Since the prizes are exactly the same, it follows that you will *prefer* Gamble A to Gamble B *if, and only if*, you believe that drawing a red ball is more likely than drawing a black ball (according to expected utility theory). Also, there would be no clear preference between the choices if you thought that a red ball was as likely as a black ball. Similarly it follows that you will *prefer* Gamble C to Gamble D *if, and only if*, you believe that drawing a red or yellow ball is more likely than drawing a black or yellow ball. If drawing a red ball is more likely than drawing a black ball, then drawing a red or yellow ball is also more likely than drawing a black or yellow ball. So, supposing you *prefer* Gamble A to Gamble B, it follows that you will also *prefer* Gamble C to Gamble D. And, supposing instead that you *prefer* Gamble D to Gamble C, it follows that you will also *prefer* Gamble B to Gamble A.

When surveyed, however, most people *strictly prefer* Gamble A to Gamble B and Gamble D to Gamble C. Therefore, some assumptions of the expected utility theory are violated.

Mathematical demonstration

Mathematically, your estimated probabilities of each color ball can be represented as: *R*, *Y*, and *B*. If you *strictly prefer* Gamble A to Gamble B, by utility theory, it is presumed this preference is reflected by the expected utilities of the two gambles: specifically, it must be the case that

$$R \cdot U(\$100) + (1 - R) \cdot U(\$0) > B \cdot U(\$100) + (1 - B) \cdot U(\$0)$$

where $U(\cdot)$ is your utility function. If $U(\$100) > U(\$0)$ (you strictly prefer $100 to nothing), this simplifies to:

$$R > B$$

If you also strictly prefer Gamble D to Gamble C, the following inequality is similarly obtained:

$$B \cdot U(\$100) + Y \cdot U(\$100) + R \cdot U(\$0) > R \cdot U(\$100) + Y \cdot U(\$100) + B \cdot U(\$0)$$

This simplifies to:

$$B > R$$

This contradiction indicates that your preferences are inconsistent with expected-utility theory.

Generality of the paradox

Note that the result holds regardless of your utility function. Indeed, the amount of the payoff is likewise irrelevant. Whichever gamble you choose, the prize for winning it is the same, and the cost of losing it is the same (no cost), so ultimately, there are only two outcomes: you receive a specific amount of money, or you receive nothing. Therefore it is sufficient to assume that you prefer receiving some money to receiving nothing (and in fact, this assumption is not necessary -- in the mathematical treatment above, it was assumed U($100) > U($0), but a contradiction can still be obtained for U($100) < U($0) and for U($100) = U($0).

In addition, the result holds regardless of your risk aversion. All the gambles involve risk. By choosing Gamble D, you have a 1 in 3 chance of receiving nothing, and by choosing Gamble A, you have a 2 in 3 chance of receiving nothing. If Gamble A was less risky than Gamble B, it would follow that Gamble C was less risky than Gamble D (and vice versa), so, risk is not averted in this way.

However, because the exact chances of winning are known for Gambles A and D, and not known for Gambles B and C, this can be taken as evidence for some sort of ambiguity aversion which cannot be accounted for in expected utility theory. It has been demonstrated that this phenomenon occurs only when the choice set permits comparison of the ambiguous proposition with a less vague proposition (but not when ambiguous propositions are evaluated in isolation; See Fox and Tversky, 1995).

Possible Explanations

There have been various attempts to provide decision-theoretic explanations of Ellsberg's observation. Since the probabilistic information available to the decision-maker is incomplete, these attempts sometimes focus on quantifying the non-probabilistic ambiguity which the decision-maker faces. That is, these alternative approaches sometimes suppose that the agent formulates a subjective (though not necessarily Bayesian) probability for possible outcomes.

One such attempt is based on info-gap decision theory. The agent is told precise probabilities of some outcomes, though the practical meaning of the probability numbers is not entirely clear. For instance, in the gambles discussed above, the probability of a red ball is 30/90, which is a precise number. Nonetheless, the agent may not distinguish, intuitively, between this and, say, 30/91. No probability information whatsoever is provided regarding other outcomes, so the agent has very unclear subjective impressions of these probabilities.

In light of the ambiguity in the probabilities of the outcomes, the agent is unable to evaluate a precise expected utility. Consequently, a choice based on *maximizing* the expected utility is also impossible. The info-gap approach supposes that the agent implicitly formulates info-gap models for the subjectively uncertain probabilities. The agent then tries to satisfice the expected utility and to maximize the robustness against uncertainty in the imprecise probabilities. This robust-satisficing approach can be developed explicitly to show that the choices of decision-makers should display precisely the preference reversal which Ellsberg observed (Ben-Haim, 2006, section 11.1).

Info-gap decision theory

Info-gap decision theory is a non-probabilistic decision theory seeking to optimize robustness to failure, or opportunity of windfall. This differs from classical decision theory, which typically maximizes the expected utility.

In many fields, including engineering, economics, management, biological conservation, medicine, homeland security, and more, analysts use models and data to evaluate and formulate decisions. An **info-gap** is the disparity between what *is known* and what *needs to be known* in order to make a reliable and responsible decision. Info-gaps are Knightian uncertainties: a lack of knowledge, an incompleteness of understanding. Info-gaps are non-probabilistic and cannot be insured against or modelled probabilistically. A common info-gap, though not the only kind, is uncertainty in the shape of a probability distribution. Another common info-gap is uncertainty in the functional form of a property of the system, such as friction force in engineering, or the Phillips curve in economics.

Info-gap commonly focuses on making decisions in such a way that unacceptably poor outcomes are avoided. Its focus on the worst possible outcome shares many features with minimax decision theory.

Info-gap models: Info-gaps are quantified by info-gap models of uncertainty. An info-gap model is an unbounded family of nested sets all sharing a common structure. A frequently encountered example is a family of nested ellipsoids all having the same shape. The structure of the sets in an info-gap model derives from the information about the uncertainty. In general terms, the structure of an info-gap model of uncertainty is chosen to define the smallest or strictest family of sets whose elements are consistent with the prior information.

A common example of an info-gap model is the **fractional error model.** The best estimate of an uncertain function $u(x)$ is $\tilde{u}(x)$, but the fractional error of this estimate is unknown. The following unbounded family of nested sets of functions is a fractional-error info-gap model:

$$\mathcal{U}(\alpha, \tilde{u}) = \{u(x) : \ |u(x) - \tilde{u}(x)| \leq \alpha \tilde{u}(x), \text{ for all } x\}, \quad \alpha \geq 0$$

At any **horizon of uncertainty** α, the set $\mathcal{U}(\alpha, \tilde{u})$ contains all functions $u(x)$ whose fractional deviation from $\tilde{u}(x)$ is no greater than α. However, the horizon of uncertainty is unknown, so the info-gap model is an unbounded family of sets, and there is no worst case or greatest deviation.

There are many other types of info-gap models of uncertainty. All info-gap models obey two basic axioms:

- **Nesting.** The info-gap model $\mathcal{U}(\alpha, \tilde{u})$ is nested if $\alpha < \alpha'$ implies that:

$$\mathcal{U}(\alpha, \tilde{u}) \subset \mathcal{U}(\alpha', \tilde{u})$$

- **Contraction.** The info-gap model $\mathcal{U}(0, \tilde{u})$ is a singleton set containing its center point:

$$\mathcal{U}(0, \tilde{u}) = \{\tilde{u}\}$$

The nesting axiom imposes the property of "clustering" which is characteristic of info-gap uncertainty. Furthermore, the nesting axiom implies that the uncertainty sets $\mathcal{U}(\alpha, u)$ become more inclusive as α grows, thus endowing α with its meaning as an horizon of uncertainty. The contraction axiom implies that, at horizon of uncertainty zero, the estimate $\tilde{u}$ is correct.

Robustness and opportuneness: Uncertainty may be either pernicious or propitious. That is, uncertain variations may be either adverse or favorable. Adversity entails the possibility of failure, while favorability is the opportunity for sweeping success. Info-gap decision theory is based on quantifying these two aspects of uncertainty, and choosing an action which addresses one or the other or both of them simultaneously. The pernicious and propitious aspects of uncertainty are quantified by two "immunity functions": the robustness function expresses the immunity to failure, while the opportuneness function expresses the immunity to windfall gain.

The **robustness function** expresses the greatest level of uncertainty at which failure cannot occur; the **opportuneness function** is the least level of uncertainty which entails the possibility of sweeping success. The robustness and opportuneness functions address, respectively, the pernicious and propitious facets of uncertainty.

Let *q* be a decision vector of parameters such as design variables, time of initiation, model parameters or operational options. We can verbally express the robustness and opportuneness functions as the maximum or minimum of a set of values of the uncertainty parameter α of an info-gap model:

$$\hat{\alpha}(q) = \max\{\alpha : \text{ minimal requirements are always satisfied}\} \quad \text{(robustness)} \quad (1)$$

$$\hat{\beta}(q) = \min\{\alpha : \text{ sweeping success is possible}\} \quad \text{(opportuneness)} \quad (2)$$

We can "read" eq. (1) as follows. The robustness $\hat{\alpha}(q)$ of decision vector *q* is the greatest value of the horizon of uncertainty α for which specified minimal requirements are always satisfied. $\hat{\alpha}(q)$ expresses robustness — the degree of resistance to uncertainty and immunity against failure — so a large value of $\hat{\alpha}(q)$ is desirable. Eq. (2) states that the opportuneness $\hat{\beta}(q)$ is the least level of uncertainty α which must be tolerated in order to enable the possibility of sweeping success as a result of decisions *q*. $\hat{\beta}(q)$ is the immunity against windfall reward, so a small value of $\hat{\beta}(q)$ is desirable. A small value of $\hat{\beta}(q)$ reflects the opportune situation that great reward is possible even in the presence of little ambient uncertainty. The immunity functions $\hat{\alpha}(q)$ and $\hat{\beta}(q)$ are complementary and are defined in an anti-symmetric sense. Thus "bigger is better" for $\hat{\alpha}(q)$ while "big is bad" for $\hat{\beta}(q)$. The immunity functions — robustness and opportuneness — are the basic decision functions in info-gap decision theory.

The robustness function involves a maximization, but not of the performance or outcome of the decision. The greatest tolerable uncertainty is found at which decision *q* **satisfices** the performance at a critical survival-level. One may select an action *q* according to its robustness $\hat{\alpha}(q)$, whereby the robustness function underlies a satisficing decision algorithm which maximizes the immunity to pernicious uncertainty.

The opportuneness function in eq. (2) involves a minimization, however not, as might be expected, of the damage which can accrue from unknown adverse events. The least horizon of uncertainty is sought at which decision *q* enables (but does not necessarily guarantee) large windfall gain. Unlike the robustness function, the opportuneness function does not satisfice, it "windfalls". When $\hat{\beta}(q)$ is used to choose an action *q*, one is "windfalling" by optimizing the opportunity from propitious uncertainty in an attempt to enable highly ambitious goals or rewards.

Given a scalar reward function *R*(*q*,*u*), depending on the decision vector *q* and the info-gap-uncertain function *u*, the minimal requirement in eq. (1) is that the reward *R*(*q*,*u*) be no less than a critical value r_c. Likewise, the sweeping success in eq. (2) is attainment of a "wildest dream" level of reward r_w which is much greater than r_c. Usually neither of these threshold values, r_c and r_w, is chosen irrevocably before performing the decision analysis. Rather, these parameters enable the decision maker to explore a range of options. In any case the windfall reward r_w is greater, usually much greater, than the critical reward r_c:

$$r_w > r_c$$

The robustness and opportuneness functions of eqs. (1) and (2) can now be expressed more explicitly:

$$\hat{\alpha}(q, r_c) = \max\left\{\alpha : \left(\min_{u \in \mathcal{U}(\alpha, \tilde{u})} R(q,u)\right) \geq r_c\right\} \quad (3)$$

$$\hat{\beta}(q, r_w) = \min\left\{\alpha : \left(\max_{u \in \mathcal{U}(\alpha, \tilde{u})} R(q,u)\right) \geq r_w\right\} \quad (4)$$

$\hat{\alpha}(q, r_c)$ is the greatest level of uncertainty consistent with guaranteed reward no less than the critical reward r_c, while $\hat{\beta}(q, r_w)$ is the least level of uncertainty which must be accepted in order to facilitate (but not guarantee) windfall as great as r_w. The complementary or anti-symmetric structure of the immunity functions is evident from eqs. (3) and (4).

These definitions can be modified to handle multi-criterion reward functions.

The robustness function generates **robust-satisficing preferences** on the options. A decision maker will usually prefer a decision option q over an alternative q' if the robustness of q is greater than the robustness of q' at the same value of critical reward r_c. That is:

$$q \succ_r q' \quad \text{if} \quad \hat{\alpha}(q, r_c) > \hat{\alpha}(q', r_c) \quad (5)$$

Let $\mathcal{Q}$ be the set of all available or feasible decision vectors *q*. A robust-satisficing decision is one which maximizes the robustness on the set $\mathcal{Q}$ of available *q*-vectors and satisfices the performance at the critical level r_c:

$$\hat{q}_c(r_c) = \arg \max_{q \in \mathcal{Q}} \hat{\alpha}(q, r_c)$$

Usually, though not invariably, the robust-satisficing action $\hat{q}_c(r_c)$ depends on the critical reward r_c.

The opportuneness function generates **opportune-windfalling preferences** on the options. A decision maker who chooses to focus on windfall opportunity will prefer a decision *q* over an alternative q' if *q* is more opportune than q' at the same level of reward r_w. Formally:

$$q >_o q' \quad \text{if} \quad \hat{\beta}(q, r_w) < \hat{\beta}(q', r_w) \tag{6}$$

The opportune-windfalling decision, $\hat{q}_w(r_w)$, *minimizes* the opportuneness function on the set of available decisions:

$$\hat{q}_w(r_w) = \arg \min_{q \in \mathcal{Q}} \hat{\beta}(q, r_w)$$

The two preference rankings, eqs. (5) and (6), as well as the corresponding the optimal decisions $\hat{q}_c(r_c)$ and $\hat{q}_w(r_w)$, may be different.

The robustness and opportuneness functions have many properties which are important for decision analysis. Robustness and opportuneness both trade-off against aspiration for outcome: robustness and opportuneness deteriorate as the decision maker's aspirations increase. Robustness is zero for model-best anticipated outcomes. Robustness curves of alternative decisions may cross, implying reversal of preference depending on aspiration. Robustness may be either sympathetic or antagonistic to opportuneness: a change in decision which enhances robustness may either enhance or diminish opportuneness. Various theorems have also been proven which show how the probability of success is enhanced by enhancing the info-gap robustness, without of course knowing the underlying probability distribution.

(Compiled from Wikipedia.org)