


Utility

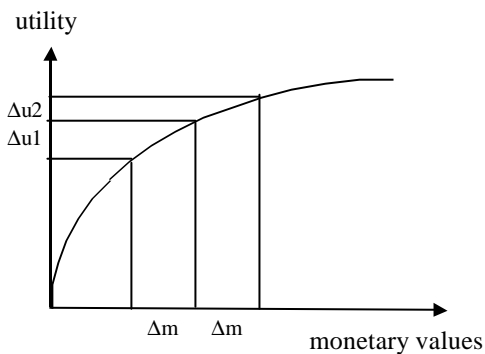
Tests and observations indicate that most people are unwilling to pay as much as the expected monetary value for an uncertain investment. They are normally *averse to risk*, and frequently prefer a sure return of lower monetary value to an uncertain return of higher expected monetary value.

Risk aversion can best be perceived by putting it in a personal context. Suppose that you were presented with an opportunity to invest money in a 50:50 proposition to earn £100 overnight. The probabilities are equal that you will earn £100 or nothing, so that the expected value of the investment is £50. Consequently any investment (payment for the opportunity to play the game) below £50 yields an expected profit. What is the most that you would be willing to invest in order to obtain this opportunity? In fact, average bids are around £25.




As Bernoulli suggested, “*The pertinent variable to be averaged ... is not the actual monetary worth of the outcomes, but rather the intrinsic worth of their monetary values*”. The intrinsic value or satisfaction derived from gaining different amounts of money is measured by an individual’s *utility function*. Decision makers generally select alternatives according to expected utility. It is precisely because people’s utility functions are non-linear that they exhibit *risk preference*, or, more commonly, *risk aversion*.

The utility of an outcome depends upon the decision maker himself. Although most people are willing to gamble small amounts such as 10p to get a 50:50 chance at a 20p prize, they are usually not prepared to make an expected value bid when the stake is higher, even though the expected value of the return is correspondingly higher. But, in certain circumstances, some people have utility functions with increasing returns along part of their range. They are prepared to act as if a large prize had a utility more than R times as large as a prize 1/R times the size. To them the larger amount, the bonanza, opens up a whole new way of life. Those who “invest” on the National Lottery are an example of people who exhibit this risk preference.



However, for practical decision-making about large scale problems it is useful to act, as do private and public organizations alike, as if the utility function has decreasing returns. *Bernoulli’s Law of Diminishing Marginal Utility* describes his observation of such risk aversion: “... any increase in wealth, no matter how insignificant, will always result in an increase in utility which is inversely proportionate to the quantity of goods already possessed”.

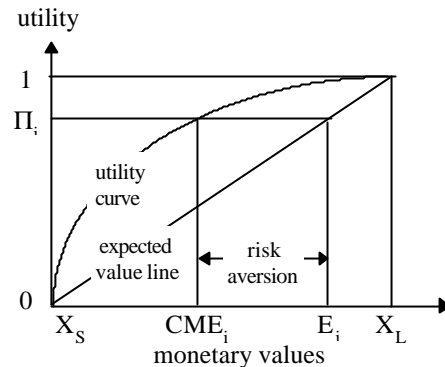


A *standard alternative* or *P_i lottery* has two outcomes, a larger one, X_L , with probability Π_i , and a smaller one, X_S , with probability $1 - \Pi_i$. In the above game we had a 0.50 lottery for £100 and £0.

The *certainty monetary equivalent CME* is the maximum amount a person would be willing to pay for a Π_i lottery, or, equivalently, the minimum amount a person owning the opportunity would be willing to sell it for. The certainty monetary equivalent is only rarely identical to the expected monetary value of an opportunity. Generally, the CME is less than the EMV, since people are risk averse. In the above game the certainty monetary equivalent was about £25, even though the expected monetary value was £50.

Constructing a Decision-Maker's Utility Function: There are a variety of closely related techniques for obtaining the utility of outcomes of choice in a probabilistic situation. The methods generate a unique number for each possible outcome which can be used to order all choices according to their desirability to the decision maker. This ultimate ranking number is derived from a combination of the measurable rewards, the probability of the occurrence of each outcome, and the decision-maker's regard for risk.

To construct a utility function for probabilistic situations, one first selects the reference rewards X_L and X_S that bound the rewards associated with all other alternatives of interest. For a lottery with prizes X_L and X_S and $\Pi_i = 1$, the CME is X_L . The monetary sum X_L is assigned the utility value 1, equal to the probability of the lottery. For similar reasons, the monetary sum X_S is assigned the utility value 0. The two points $(X_L, 1)$ and $(X_S, 0)$ are taken to be the end points of the decision maker's utility curve.



To obtain intermediate points on the utility curve, the decision-maker is confronted with a series of Π_i lotteries and asked to define his CMEs for each such lottery. For each Π_i lottery, the utility of CME_i is set to the value of probability Π_i . In other words, each point (CME_i, Π_i) is on the utility curve.

The expected monetary value of any Π_i lottery is:
$$E_i = \Pi_i * X_L + (1 - \Pi_i) * X_S.$$
 Each point (E_i, Π_i) lies on the straight line between $(X_L, 1)$ and $(X_S, 0)$. The extent of risk aversion is measured by the difference between CME_i and E_i for any Π_i lottery or probabilistic situation.

In practice, it may be hard for a decision-maker to distinguish his feelings about slightly different probabilistic situations (involving, say, a probability of 0.75 as against a probability of 0.8), and so it may be more generally useful to define lotteries in terms of fixed probabilities which can be visualized readily, and to vary the prizes instead. Here is an example of an operational method that uses only conceptually simple 50:50 gambles, in a situation where X_L is £80,000 and X_S is -£30,000.

The analyst initially sets $utility(£80,000) = 1$, and $utility(-£30,000) = 0$.

ANALYST: *What certain outcome would you consider equivalent to a 50:50 gamble on outcomes of £80,000 and -£30,000?*

DECISION-MAKER: £10,000

Since the expected utility of the above gamble is $0.5 * 1 + 0.5 * 0 = 0.5$, and the decision-maker is indifferent between this gamble and £10,000 for certain, we deduce that $utility(£10,000) = 0.5$.

ANALYST: *What certain outcome would you consider equivalent to a 50:50 gamble on outcomes of £80,000 and £10,000?*

DECISION-MAKER: £25,000

Since the expected utility of the above gamble is $0.5 * 1 + 0.5 * 0.5 = 0.75$, and the decision-maker is indifferent between this gamble and £25,000 for certain, we deduce that $utility(£25,000) = 0.75$.

ANALYST: *What certain outcome would you consider equivalent to a 50:50 gamble on outcomes of £10,000 and -£30,000?*

etc.

In this way, the interval between the best and worst outcomes is bisected repeatedly. Soon the questioner will have enough information to plot the decision-maker's utility function. A freehand curve can then be drawn to illustrate the decision-maker's utility curve.


Expected Utility Value Criterion: Not only do many people exhibit decision making behaviour generally in line with the selection of alternatives according to expected utility, but it can be shown that, under a very reasonable set of assumptions, rational decision making should be based on the principle of expected utility. In other words, one should use the same method as described earlier under the Expected Monetary Value criterion, but use utility values instead of monetary payoffs.

Assumptions of the utility function: The concept of the utility function and rational choice rests on fundamental axioms or assumptions. These utility axioms are intuitively appealing insofar as they set forth the kind of conditions which common sense indicates should surround a rational decision process. These six axioms are sufficient to define a utility function in a probabilistic context. Unless one is prepared to define rational decision-making in terms of a contrary set of assumptions, one is forced to accept the principle that decisions should be based on the principle of maximizing expected utility.


Axiom 1. A decision-maker can compare and make consistent choices between alternatives. Faced with a pair of alternatives A and B, he will always prefer one to the other or be indifferent between them.

Axiom 2. A decision-maker's choices are transitive. If he prefers A to B and B to C, then he will prefer A to C. Likewise, if he is indifferent between A and B and indifferent between B and C, then he will be indifferent between A and C.


Axiom 3. If two choices are indifferent to the decision-maker, they can be substituted for each other. For example, if the decision-maker is indifferent between:




and




then he will also be indifferent between:




and




Axiom 4. A decision-maker will be indifferent between a compound lottery and a simple lottery which offers the same rewards with the same probabilities. For example, the decision-maker will be indifferent between:




and




Axiom 5. If two lotteries A and B lead to the same outcomes, C and D, the decision-maker will choose the lottery in which the preferred outcome C has the greater probability. For example, the decision-maker will prefer:




to




Axiom 6. If the decision-maker prefers A to B and B to C, then there is a Π_i lottery with prizes A and C which is indifferent to B:



and



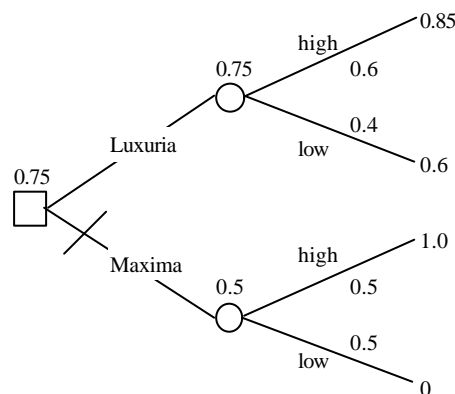
Example: A business woman who is organizing an exhibition in a provincial town has to choose between two venues: the Luxuria Hotel and the Maxima Centre. To simplify her problem, she decides to estimate her potential profit at these locations on the basis of two scenarios: high attendance and low attendance at the exhibition. If she chooses the Luxuria Hotel, she reckons that she has a 60% chance of achieving a high attendance and hence a profit of £30,000 (after taking into account the costs of advertising, hiring the venue, etc.). There is, however, a 40% chance that attendance will be low, in which case her profit will be just £11,000. If she chooses the Maxima Centre, she reckons she has a 50% chance of high attendance, leading to a profit of £60,000, and a 50% chance of low attendance leading to a loss of £10,000.



The business woman's expected profit is £22,400 if she chooses the Luxuria Hotel and £25,000 if she chooses the Maxima Centre. By the EMV criterion she should choose the Maxima Centre, but this is the riskier option, offering high rewards if things go well but losses if things go badly.

After having been asked a series of questions about her indifference between various hypothetical lotteries, we have managed to identify the following points on the business woman's utility curve:


Monetary sum	Utility
£60,000	1.0
£30,000	0.85
£11,000	0.60
-£10,000	0




These results are now applied to the decision tree by replacing the monetary values with their utilities. By treating these utilities in the same way as the monetary values we are able to identify the course of action which leads to the highest expected utility. Choosing the Luxuria Hotel gives an expected utility of 0.75. Choosing the Maxima Centre gives an expected utility of 0.5. Thus the business woman should choose the Luxuria Hotel as the venue for her exhibition. Clearly, the Maxima Centre is too risky.

Various types of Utility Functions: If we plot the business woman's utility function, we obtain the concave curve of someone who is *risk averse*.

Suppose that the business woman had assets of £30,000 and were offered a gamble that would give her a 50% chance of doubling her money and a 50% chance of losing it all. Her current assets have a utility of 0.85. If she gambles she has a 50% chance of increasing her assets so that their utility would increase to 1.0, and a 50% chance of ending with assets with a utility of about 0.35. Hence the expected utility of the gamble is about 0.675, which is less than the utility of her current assets. The increase in utility that will occur if she wins is far less than the loss in utility she will suffer if she loses.



Other typical utility curves are shown below:



The left utility curve indicates a *risk-seeking* attitude. A person with a utility function like this would have accepted the above gamble. The linear utility function in the centre demonstrates a *risk-neutral* attitude. If a person’s utility function looks like this then the EMV criterion will represent their preferences. The utility curve on the right indicates both a risk-seeking attitude and risk aversion, depending on the current level of assets.


Utility Functions for Non-Monetary Attributes: Utility functions can be defined for attributes other than money.

Example: A drug company is hoping to develop a new product. If it succeeds with its existing research methods it estimates that there is a 0.4 probability that the drug will take 6 years to develop and a 0.6 probability that the development will take 4 years. However, recently a “short-cut” method has been proposed which might lead to significant reductions in the development time, and the company, which has limited resources available for research, has to decide whether to take a risk and switch completely to the proposed new method. The head of research estimates that, if the new approach is adopted, there is a 0.2 probability that development will take a year, a 0.4 probability that it will take 2 years and a 0.4 probability that the approach will not work and, because of the time wasted, it will take 8 years to develop the product.

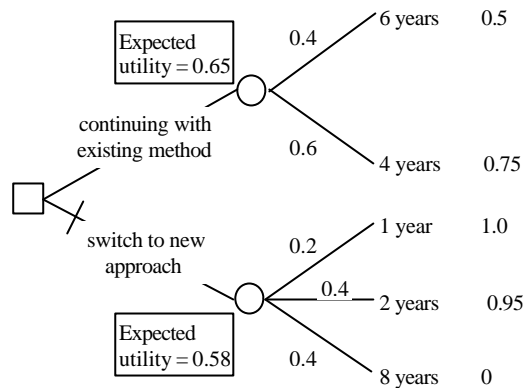
Adopting the new approach is risky, and so we need to derive utilities for the utility times. The worst development time is 8 years, so $u(8 \text{ years}) = 0$, and the best time is 1 year, so $u(1 \text{ year}) = 1.0$. After being asked a series of questions, the head of research is able to say that he is indifferent between a development time of 2 years and engaging in a lottery which will give him a 0.95 probability of a 1-year development and a 0.05 probability of an 8-year development time. Thus:

$$\begin{aligned} u(2 \text{ years}) &= 0.95 \cdot u(1 \text{ year}) + 0.05 \cdot u(8 \text{ years}) \\ &= 0.95 \cdot 1.0 + 0.05 \cdot 0 = 0.95 \end{aligned}$$

By a similar process we find that $u(4 \text{ years}) = 0.75$ and $u(6 \text{ years}) = 0.5$. The utility function has a concave shape indicating risk aversion.



The utilities are shown in this decision tree, where it can be seen that continuing with the existing method gives the highest expected utility.



It is also possible to derive utility functions for attributes which are not easily measured in numerical terms. For example, consider the choice of design for a chemical plant. Design A may have a small probability of failure which may lead to pollution of the local environment. An alternative, design B, may also carry a small probability of failure which would not lead to pollution but would cause damage to some expensive equipment. If a decision-maker ranks the possible outcomes from best to worst as: (i) no failure, (ii) equipment damage, and (iii) pollution, then $u(\text{no failure}) = 1$ and $u(\text{pollution}) = 0$. The value of $u(\text{equipment damage})$ could then be determined by posing questions such as:

Which would you prefer:

- (1) A design which was certain at some stage to fail, causing equipment damage; or
- (2) A design which had a 90% chance of not failing and a 10% chance of failing and causing pollution?

Once a point of indifference was established, $u(\text{equipment damage})$ could be derived.

A similar application in the electronics industry involves designs of electronic circuits for cardiac pacemakers. The designs carry a risk of particular malfunctions and the utilities relate to outcomes such as “pacemaker not functioning at all”, “pacemaker working too fast”, “pacemaker working too slowly” and “pacemaker functioning OK”.

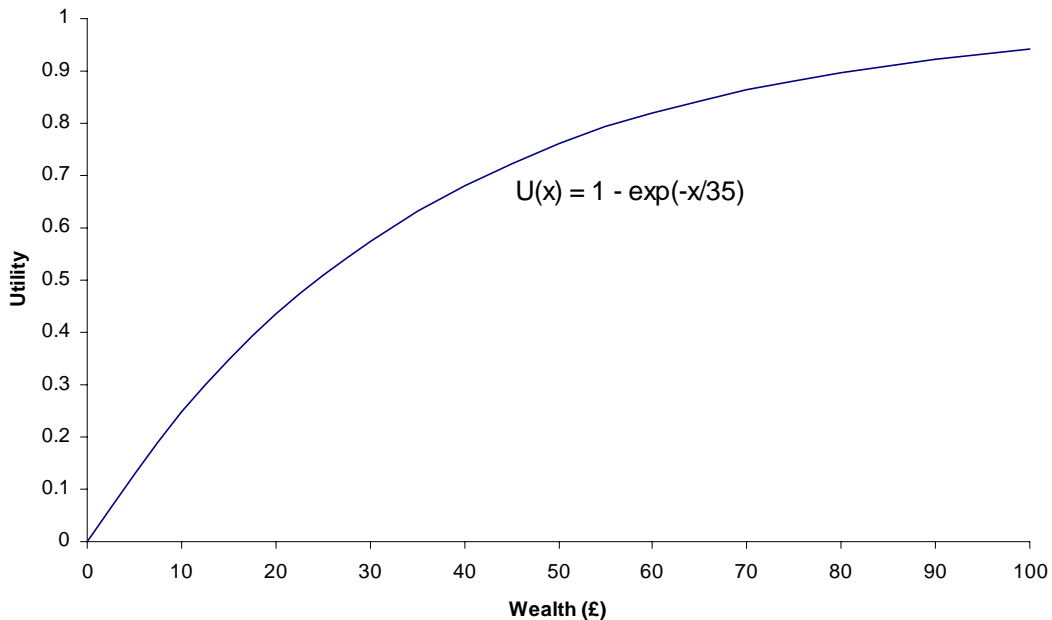
Risk Tolerance

The process of assessing a utility function subjectively involves asking the decision maker quite a few questions about certainty equivalents to various lotteries, in order to plot sufficient points on the curve to fit a smooth line. If we are willing to assume that the utility function follows a well-defined mathematical curve, then all we need to do is to estimate the parameter(s) involved, which generally involves less work.

For example, consider the exponential utility function:

$$U(x) = 1 - e^{-x/R}$$

The function is concave and can be used to represent risk-averse behaviour. $U(x)$ tends to 1 as x becomes large, $U(0) = 0$ and $U(x)$ is negative for negative x (being in debt).



The parameter R (equal to £35 in the above diagram) determines how risk averse a person is, and is called the **risk tolerance**. A large value of R implies a flatter curve and high risk tolerance. A low value implies a highly curved function and low risk tolerance.

R can be approximated relatively easily by asking the decision maker to consider the following choice:



For which sum of money $£C$ is the decision maker indifferent between the 50-50 gamble with prizes ($£C$ and $-£C/2$) and simply breaking even? This amount $£C$ is approximately equal to the value of the decision maker's risk tolerance $£R$.

Once you know R (approximately) you can plug the value into the exponential function and can then calculate utility values for different amounts of money. Once you know the expected utility of a lottery you can find the certainty monetary equivalent by inverting the utility function and using natural logarithms.

Within any given problem the utility function can be re-scaled or shifted by multiplying by a constant or moving the origin, e.g.,

$$U(x) = S*(1 - e^{-(x-T)/R})$$

The exponential utility function implies an attitude towards risk called **constant risk aversion**. No matter how much wealth a decision maker had a particular gamble would be viewed in the same way - in other words the risk premium (EMV minus CME) would be the same.

The trouble is that most people's risk tolerance is not constant, but increases (and risk aversion decreases) with the amount of money they have. Despite this reservation, the exponential utility function can be a useful "quick-and-dirty" method of modelling someone's utility function. If a decision maker is facing a choice between two alternatives and the right decision, based on expected utilities is quite clear cut, then this utility function may be sufficient. If the choice is not so clear, then a more careful assessment of the utility function may be called for.

Increasing risk tolerance (and so decreasing risk aversion) implies that the more money a person has, the less nervous they are about a particular bet and so the smaller its risk premium. The logarithmic utility curve:

$$U(x) = \log_e(x)$$

corresponds to risk preference that exhibits increasing risk tolerance. This function can also be re-scaled and shifted.

Examples of (re-scaled) logarithmic and exponential utility curves are sketched below. Notice how similar they are. It does not take a large change in the utility curve's shape to alter the nature of the individual's risk attitude.

