Haberdashers' School

Occasional Papers Series

Occasional Paper Number Ninety-One

Von Neumann and Morgenstern's Cardinal Measure of Utility and the Shape of the Utility of Wealth Function

Dr Ian St John

Department of Economics, Haberdashers' School

istjohn@habselstree.org.uk

November 2025

Haberdashers' School Occasional Paper Number Ninety-One
November 2025

Von Neumann and Morgenstern's Cardinal Measure of Utility and the Shape of the Utility of Wealth Function

Ian St John

Abstract

Economists have tended to assume that it is impossible for consumers to attach a specific value to the degree of utility or satisfaction that derive from a good. Hence, demand theory in the first-half of the twentieth century involved attempts to model consumer behaviour by making the most minimal possible assumptions regarding the consumers' ability to quantify the utility they derived from consuming a good – first with Indifference Curve analysis and then with Revealed Preference theory, in which consumers were assumed to only be able to say they preferred one bundle of goods to another. However, in the course of constructing their analysis of the prospective pay-offs from choices in the context of Game Theory, Von Neumann and Morgenstern developed a technique by which cardinal values can be attached to different combinations of goods based on the degrees of probability of differing expectations being realised. We can gauge the degree to which one bundle of goods will yield less utility than another by the extent to which the likelihood of that inferior bundle needs to be raised to make the individual indifferent between the two. In this way we can generate an index of cardinal utility for expected consumption bundles. Using this insight it is possible to construct a utility function for money, which in turn helps us to explain risk-averse and risk-loving forms of behaviour, and the willingness of some people to take out insurance and the willingness of others to gamble.

Economists usually state that a cardinal measure of utility, in which consumers are able to attach a particular number-value to the utility of good A as opposed to good B, is impossible. When the concept of utility was first applied to the problem of prices in the 1870s it was hoped that a definite measure of utility – the 'util' – would be possible. In which case, the theory of decision making could be resolved into positive statements about actual psychological data regarding the level of pleasure or satisfaction people derived from consuming one good or another or doing one thing or another. Yet this ambition was soon abandoned: while, it was agreed, people *could* say that they preferred good A to good B, or were indifferent between good A and good B, they could not say *by how much* they preferred A to B. Equally, it couldn't be said *how much* utility one person derived from A compared to another: inter-personal comparisons of utility were impossible. Hence, in the development of consumer behaviour and price theory, the concept of a measurable marginal utility from the consumption of a good was discarded; instead, demand theory was constructed on the less stringent assumption that a consumer could place a series of goods in a hierarchy of preferences, saying that they preferred A to B and B to C and hence A to C. It is this assumption that underlies the indifference curve theory of demand.

However, in the process of developing their theory of games in the 1940s, which involves formulating strategies in the light of possible outcomes, Von Neumann and Morgenstern constructed a model of *expected* future utility, and on the basis of this model they developed a *cardinal* measure of the utility of *expected* outcomes.[1] This measure is similar to that used to estimate temperature. By referring to a thermometer we can say, not only that today is hotter than yesterday (an ordinal measure), we can say *by how much* it is hotter – say, 5 degrees or 10 degrees. We say today is 25 degrees centigrade and yesterday was 20 – so today is 5 degrees or 25 per cent hotter than yesterday. These are precise statements of orders of magnitude, but they are *not* an actual measure of heat. They refer only to points on a Celsius scale, which is only one of many possible scales for measuring temperature. If, for example, we used a Fahrenheit scale, then yesterday would be 68 degrees Fahrenheit (compared to 20 degrees Celsius) and today 77 degrees, in which case today is 9 degrees or 13 per cent hotter than yesterday. Each scale gives specific results, but the numbers are arbitrary, depending on the scale selected. It is just such a measurable yet arbitrary scale that Von Neumann and Morgenstern constructed for possible future consumption utilities.

[1] J. Von Neumann and O. Morgenstern, *Theory of Games and Economic Behavior* (Princeton University Press, Princeton, 1943)

Von Neumann and Morgenstern

An Expected Utility Function

Typically, we might say that a consumer can purchase two goods, C1 and C2, and that their total utility is a function of their consumption of these two goods:

$U = U(C_1, C_2)$ (1)

A person will maximise their utility by choosing a particular combination of C1 and C2. In reality, things are not so simple. Suppose a person is deciding where to go on holiday, where C1 is to spend a week in Bournemouth and C2 is to go skiing in New Zealand. The utility of each depends, obviously, on the subjective preferences of each person: how much they like sunbathing, how much they enjoy skiing etc. But the utility of each also depends upon unpredictable risk factors. If I go to Bournemouth, what is the likelihood that it will rain all week or there will be a train strike? If I go skiing, what are the chances that it won't snow or that my skis will get lost by the airline? All choices to buy anything involve *risk* and to choose rationally we need attach probabilities to such risks. Should I buy a new car or a second-hand one? A new car is more expensive but the risk it will have a fault is low (though not zero); a second-hand car is cheaper, yet the risk it will have a defect is higher. In deciding which choice will maximise our utility we need a way to weigh up the probable outcomes of our decision.

To take an example from Hal Varian. Suppose I have £4 and am confronted by two choices: to purchase a cup of coffee or buy a lottery ticket offering a £1 million jackpot.

$U(C_1)$ = utility of a cup of coffee

$U(C_2)$ = utility if I win the prize of £1 million

What should a rational consumer do? The utility of winning the lottery far exceeds the utility of a coffee. Yet the odds of my £4 translating into a decent cup of coffee are significantly higher than are my odds of winning the lottery![1] Hence our utility depends on a combination of the subjective utility of a prospective product and the probability that the utility of the outcome actually occurs. We therefore say:

Expected Utility $U_E = U_E(C_1, C_2, \pi_1, \pi_2)$ (2)

where C_1 and C_2 are the utilities of coffee and winning the lottery, and π_1 and π_2 are the probabilities of a good cup of coffee and winning the lottery. In this case, we might say that the chance of getting a good cup of coffee is 0.95 and the chance of winning the lottery 0.0001. Thus, in deciding whether to spend by £4 on buying a cup of coffee or a lottery ticket, I must take into account, not just the benefit to me of the two products, but the *likelihood* of their happening. A lottery win of £1 million promises huge utility, but the odds on this happening are one in 10,000. Many of us would prefer to take our 95 per cent chance on a cup of coffee.

Let us consider the lottery case more closely. Suppose that there are only two possible outcomes from buying a lottery ticket: either your number comes up and you win the prize (A), or your number does not and you get nothing (B). We call these two contrasting outcomes *states of affairs*, and which occurs is said to be *contingent* on the random luck of the draw. Given that there are only two outcomes – a win and a non-win – their probabilities sum to one. If the chance of winning the lottery is 0.10, then the chance of not winning is 0.9, and the sum of both is one or 100 per cent. Because the two outcomes are mutually exclusive, we can write the *expected utility* of purchasing a lottery ticket as follows:

$E[U(L_1)] = \pi_1 U(A) + (1 - \pi_1)U(B)$ (3)

This equation states that the *expected* utility of purchasing a lottery ticket is the utility of a winning ticket (A) multiplied by the probability of having a winning ticket (π_1) plus the utility of not having the winning ticket (B – which we can assume is zero) multiplied by the probability of not winning the lottery ($1 - \pi_1$). According to Von Neumann and Morgenstern, a rational consumer will wish to maximise their *expected* utility. Thus, if there were another lottery where the outcomes were a win (C) and a loss (D), with probabilities of π_2 and $(1 - \pi_2)$, then the expected utility from this second lottery would be:

$E[U(L_2)] = \pi_2 U(C) + (1 - \pi_2)U(D)$ (4)

In this case, if:

$E[U(L_2)] > E[U(L_1)]$

then the consumer would prefer (assuming the cost of the ticket were the same) to play the second lottery to the first since the expected utility from purchasing a ticket is

[1] Still not 100 per cent of course. I once purchased a cup of coffee from *Subway* and set off down the road only to find that the milk was off and the entire cup had to be chucked away.

higher. This must be because either the winning prize is greater and/or the chances of winning are higher.

Let us assume that the prizes and prices of the ticket remain constant in the two lotteries, but that we steadily raise the probability of drawing a winning prize in the first lottery. As the probability (π_1) of winning the jackpot on the first lottery rises then the expected utility from entering the first lottery will rise, and there will eventually become a point at which the expected utilities from the two lotteries are the same; i.e.

$$E[U(L_2)] = E[U(L_1)] \qquad (5)$$

When this is true the consumer will be *indifferent* between the two lotteries since the expected utility they will derive from each is the *same*. This insight, namely that by varying the probabilities of certain outcomes we can bring about a position in which the consumer is *indifferent* between them, is what allowed Von Neumann and Morgenstern to construct a cardinal measure of utility. For the essential point is this: the *more* lottery one is perceived by the chooser to be inferior to lottery two, the *higher* must be the probability of winning the prize in lottery one if the person making the choice is to be indifferent between them. Probability is an inverse proxy for utility.

Constructing a Cardinal Utility Index from Expected Utilities

Suppose a person has the possibility of consuming a range of bundles of two goods, tea and coffee.[1] Assume there are seven such bundles, each consisting of varying quantities of tea and coffee, and the person is able to arrange such bundles in ascending order of preference, calling the least preferred bundle, bundle A, and the most preferred, bundle G. We can arrange the bundle options in a table.

[1] This example is based on R.G.D. Allen, *Mathematical Economics* (Macmillan, London, 1956), p. 671.

Bundle	Rank order of Preference	Pounds per week		Index of Cardinal Utility
		Tea	Coffee	
A	1	2	0	1.5
B	2	2	1	2
C	3	2	2	2.5
D	3	3	0	2.5
E	5	2	3	3
F	6	4	2	5
G	7	4	3	5.6

Table 1. Rank Orderings of Different Combinations of Tea and Coffee

In this table, bundle B is preferred to bundle A, bundle C to bundle B, bundle D to bundle C, and so on.

Assume, now, that the consumer has an actual bundle B. This bundle is *certain* and they can *definitely* have 2 pounds of tea and 1 pound of coffee. The consumer is then offered the chance to gamble this bundle B for the chance to have either bundle A (less preferred to B) or bundle C (preferred to B). The odds of getting bundle A are π and the odds of getting C are (1 – π). For example, think of 100 balls in a jar, where 20 per cent of the balls are A and 80 per cent are C, and the consumer is offered the chance to trade in their bundle B for the chance to do this lucky dip. The expected utility of the gamble is:

$$E(U) = \pi U_A + (1 - \pi) U_C \qquad (6)$$

Here U_A and U_C are the utilities the consumer would derive from bundles A and C respectively.

The consumer, who is assumed to wish to maximise their expected utility, has a choice: do they keep the certain bundle B or give it up for the uncertain outcomes A or C? This depends on the relative utility (desirability) of A and C compared to B and the odds that either will occur. If we take the relative of utilities of A, B, and C to be constant, then the decision of whether or not to gamble will depend upon the probability of bundle A coming up (π) and, hence, the probability of C occurring (1-π). If the consumer is told that the distribution of balls in the glass has gone from 20 A and 80 C to 40 A and 60 C, then their chances of experiencing a *fall* in utility from the gamble (giving up B and receiving A) have increased and they will be *less* likely to take the gamble. By contrast, if the proportion of C balls rises from 80 to 90 per cent, then their chances of improving their utility through the gamble have increased and they will be more likely to trade in the certainty of B for the chance of C.

Assume, now, that there is a set of odds attached to A and C which are such that the consumer is *indifferent* between keeping B for sure or gambling on A or C. Assume, in other words, that there is a probability of securing bundle C in the gamble which is sufficiently high to ensure that the consumer doesn't mind if they keep B or take a gamble on A or C. As an example, imagine that the consumer is indifferent between a 65 per cent chance of C and the certainty of B. If the chance of a C bundle were 66 per cent they would definitely take the gamble, and if the chance of a C bundle were 64 per cent, then they would not. Only the probability rate of 65 per cent leaves the consumer undecided between B and the chances of A or C. Let us call this indifference probability ($1-\pi^*$). In this case we can say:

$$U_B = \pi^* U(A) + (1-\pi^*)U(C) \qquad (7)$$

Equation (7) means that, when the consumer is indifferent between certain B and possibilities of A and C, then the utility of B is *equal to the expected utility of the gamble*. Since we are holding the utilities of tea and coffee bundles A, B, and C constant, it is by varying the probabilities of A and C that the two sides of the equation are brought into equality.

To see how, using this concept of indifference between expected utilities, we can construct a cardinal measure of utility, consider again **Table 1**. To construct a scale of utilities for these various bundles we arbitrarily attach numbers to two of the bundles: we designate the utility of bundle B as 2 and of bundle E as 3 (shown in the fifth column). These numbers do not mean anything, just as the numbers in a temperature scale do not mean anything in themselves. They are simply a device enabling us to construct a scale permitting the comparison of bundles of tea and coffee (temperatures).

Suppose our person *has* bundle C of tea and coffee, with 2lbs of each. They are certain of C. However, they are offered the chance to gamble C for the *chance* of B or E. B is inferior to C (they get one less pound of coffee), but E is superior (they would get the same amount of tea as C but more coffee). We then ask the person what odds they would require of gaining E through a gamble to cause them to be indifferent between keeping C or gambling C on the chance of B or E. Suppose they say they would be indifferent between C and B or E when the probability of getting E is 50 per cent (50-50). This means there is a 50 per cent of losing a pound of coffee and a 50 per cent chance of gaining a pound of coffee. Given this, we can now assign a value to the utility of C:

$$U_C = \frac{1}{2}\,U_B + \frac{1}{2}\,U_E$$

Since have already given values of 2 to U_B and 3 to U_E, then:

$$U_C = \frac{1}{2}\,(2) + \frac{1}{2}\,(3) = \frac{1}{2}(2 + 3) = \frac{1}{2}(5) = 2.5$$

Hence, the utility of bundle C is 2.5.

Now imagine the person *has* bundle B, but is offered the chance of gambling B for the possibility of A or E. The question again is: what odds would the person need to be offered to gain an improved bundle E in order for them to be indifferent between keeping B or exchanging this for the chances of a deterioration (A) or an improvement (E) in their state of affairs? Compared to the previous gamble, in this case the person is starting from a worse position since B is a less preferred bundle of tea and coffee than C. The person has less to lose than in the first case. Further, the move from B to E represents a larger improvement in their situation than the move from C to E. Hence, we might assume that they will accept the gamble of B in return for the chance of E at a lower rate of probability than their previous gamble. For example, we might say that they will be indifferent between the certainty of B and gambling for A/E when there is a one-third chance of E and hence a two-thirds chance of A. In which case:

$$U_B = \frac{2}{3} U_A + \frac{1}{3} U_E$$

We already know the utilities of B and E, so we can write:

$$2 = \frac{2}{3} U_A + \frac{1}{3} (3) = \frac{2}{3} U_A + 1$$

$$1 = \frac{2}{3} U_A$$

$$U_A = 1.5$$

Thus, the utility of bundle A is 1.5, 25 per cent less desirable than bundle B and half as good as bundle E.

Next, assume the person has bundle B (worth 2) and is offered the gamble of A (1.5) or F. Since a move from B to F would be a significant improvement (doubling their amount of tea and coffee) they might be indifferent to a keeping B and having a $\frac{1}{7}$ chance of F. In this case the expected utilities are as follows:

$$U_B = \frac{6}{7} U_A + \frac{1}{7} U_F$$

$$2 = \frac{6}{7} (1.5) + \frac{1}{7} U_F = 1\frac{2}{7} + \frac{1}{7} U_F$$

$$\frac{5}{7} = \frac{1}{7} U_F$$

$$U_F = 5$$

So bundle F has a utility rating of 5. Lastly, imagine our person is certain of bundle B but offered the chance of A or G. G is still better than F, so if they would be indifferent to B and a $\frac{1}{7}$ chance of F, then they might be indifferent between B and a $\frac{1}{8}$ chance of G. In which case the indifference condition is:

$U_B = \frac{7}{8} U_A + \frac{1}{8} U_G$

$2 = \frac{7}{8} (1.5) + \frac{1}{8} U_G$

$2 = (1.3) + \frac{1}{8} U_G$

$0.7 = \frac{1}{8} U_G$

$U_G = 5.6$

So the utility index of the most preferred bundle is 5.6.

Provided, therefore, we can arrange goods in a hierarchy of preference, and can attach random numbers to the utilities of two of them, then, by using probabilities to arrive at combinations of goods between which the consumer is indifferent, we can construct a cardinal index series of utilities.

Measuring the Utility of Money

Thus far we have offered an individual the chance to gamble certain amounts of two goods for the chance to gain more of those goods. We can proceed in the same way if we analyse the decisions of an individual who has a certain sum of *money* and is offered the chance to gamble that sum for the uncertain results of a lottery. In this way we can derive a cardinal index measure for the utility of sums of money.[1]

Consider a person with £1 in money wealth. They have this money and hence have a certain level of utility U(£1). Suppose they are offered the chance of keeping this £1 or using it to buy a lottery ticket which carries a prize of £100. As before, we ask the individual to choose that probability of winning the prize at which they would be *indifferent* between keeping the £1 or gambling it on the chance of £100. Call this probability π_1. Then:

$$U(£1) = \pi_1 U(£100) + (1-\pi_1)U(0) \qquad (8)$$

Here the odds of *not* winning the prize are $(1-\pi_1)$, in which case the individual will receive nothing. For example, we might say the person will be indifferent between keeping their pound and buying a lottery ticket when there is a one per cent chance of winning and a ninety-nine per cent chance of losing their money.

Now assume that the individual is offered the choice of a *certain* £2 or a gamble between £100 and £0. Since they now have slightly more to lose (£2 instead of £1) the odds on a win will need to improve slightly if the person is to be indifferent between

[1] For this section I have used S. Estrin, D. Laidler, and M. Dietrich, *Microeconomics* (Prentice Hall, Harlow, Fifth Edition, 2008), pp. 118-124.

keeping the £2 and entering the lottery. The condition for the person to be indifferent between the utility £2 and the expected utility of the gamble is then:

$U(£2) = \pi_2 U(£100) + (1-\pi_2)U(0)$ (9)

where $\pi_2 > \pi_1$. So, for example, the person my want a two per cent chance of winning the £100 to enter the bet, compared to one per cent before. If we offered the choice between U(£3) and the chance of U(£100) or nothing, then they would be indifferent when:

$U(£3) = \pi_3 U(£100) + (1-\pi_3)U(0)$ (10)

Again, $\pi_3 > \pi_2 > \pi_1$.

Proceeding in this fashion, we arrive at a series of equations of indifference:

$U(£1) = \pi_1 U(£100) + (1-\pi_1)U(0)$

$U(£2) = \pi_2 U(£100) + (1-\pi_2)U(0)$

$U(£3) = \pi_3 U(£100) + (1-\pi_3)U(0)$

.

.

$U(£99) = \pi_{99} U(£100) + (1-\pi_{99})U(0)$

$U(£100) = \pi_{100} U(£100) + (1-\pi_{100})U(0)$

The probability of a successful gamble needs to rise as the individual is asked to pay up larger and larger sums of *certain* money and hence utility. A person with £99 would only gamble on the chance of £100 if the odds were something like 99.5 per cent. Once a person has £100 for certain then they would only be indifferent between this and a *chance* of £100 if the odds on winning the £100 were 100 per cent. Hence, $\pi_{100} = 1$.

Since the outcome of not winning the lottery is always U(0), whatever the odds, we can ignore these outcomes in terms of expected utility. Rearranging the above series we have:

$$\frac{U(£1)}{U(£100)} = \pi_1$$

$$\frac{U(£2)}{U(£100)} = \pi_2$$

$$\frac{U(£3)}{U(£100)} = \pi_3$$

.

.

$$\frac{U(£99)}{U(£100)} = \pi_{99}$$

$$\frac{U(£100)}{U(£100)} = 1$$

As can be seen, the probability numbers on the right (which, remember, are the odds a person would require to be indifferent between keeping the certain numerator in each case and gambling for the uncertain denominator) form a cardinal index measure of the utility of each sum of money. As the probability required to induce the person to be indifferent between a given some of money for certain and the chance of a greater sum increases, so we can infer that the utility of the initial certain some of money is increasing also. A person attaches greater utility to £20 than to £1, so to be willing to gamble that £20 they require a higher chance of success than they would to gamble £1. As we said, a person might be willing to gamble £1 for a 0.01 chance of £100, but to gamble £20 for the chance of £100 they might require a 0.15 chance: and this increase from 0.01 to 0.15 is a measure of the increased utility of £20 compared to £1. Again: these numbers do not tell us the *actual* utility of a sum of money, but they enable us to *compare* the utilities of different sums of money.

The Shape of the Money Wealth Utility Function

What happens to the marginal utility of money as people have more of it? It is conventionally believed that, while the marginal utility of money is positive, it decreases as the stock of money a person has increases.[1] In terms of our cardinal index of the utility of money, this means that as a person's certain stock of money increases from £1 towards £100, the greater must be the increase in the probability of a 'win' to ensure they are willing to make the gamble. The more money a person has the less attractive, relatively, becomes the winning prize compared to the loss in money

[1] This point was first clearly expounded in the context of the expected utility of money income by the eighteenth-century mathematician, Daniel Bernoulli, who, in his 'St Petersburg Paradox', argued that people were prepared to pay only limited finite amounts of money to play games with infinite expected payoff value. The reason was, he said, that people valued, not the money but the utility ('moral value') the money yielded, and this declined as people had more of it. C.f. W. Nicholson, *Microeconomic Theory: Basic Principles and Extensions* (The Dryden Press, Illinois, Second Edition, 1978), pp. 150-51.

they will suffer if they lose – so the probability of a win must rise to compensate for this effect. Such a utility function with respect to money wealth is represented in **Figure 1**.

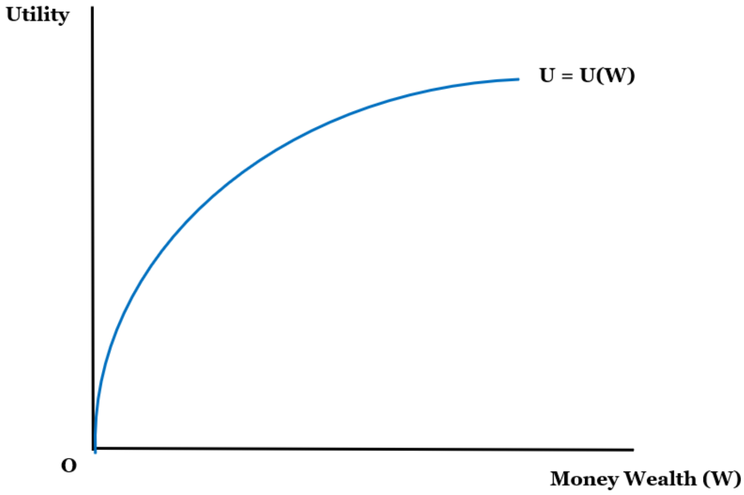

Figure 1. Utility of Money Wealth Function

In this utility function the marginal utility of money is positive (dU/dW > 0), but the marginal utility of money wealth declines as the stock of money wealth increases (d^2U/d^2 < 0). We say there is a diminishing marginal utility of money. A typical such function would be:

U = $W^{\frac{1}{2}}$ (11)

where the first derivative of utility with respect to money wealth (marginal utility) is:

$\frac{dU}{dW}$ = 0.5$W^{-0.5}$ (12)

which is positive, and the second derivative measuring the change in marginal utility:

$\frac{d^2U}{dW^2}$ = -0.25$W^{-1.5}$ (13)

is negative, showing that marginal utility decreases as wealth increases.

A utility function such as (9) gives important insights into how a person will act when confronted with choices between certain and expected utilities of the kind we have been considering thus far. **Figure 2** illustrates.

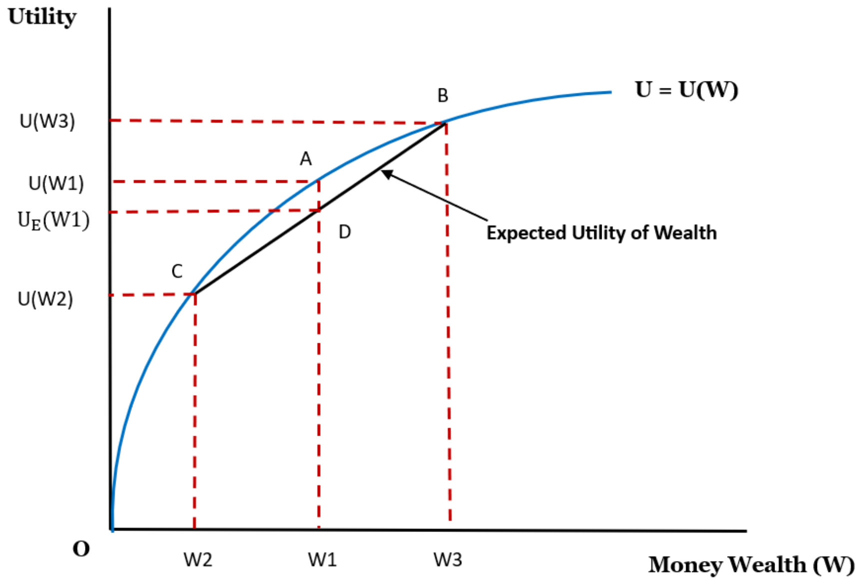

Figure 2. Utility of Wealth and the Expected Utility of Wealth

Suppose a person initially occupies point A, with a level of wealth W1 and a *certain* utility of U(W1). This person is now offered a 'fair' 50-50 gamble (by tossing a coin for example): they have a 50 per cent chance of gaining W3 – W1 wealth or of losing W1 – W2 (where these amounts are equal in absolute terms).[1] Thus, we might say W1 = £100 and the person is offered a 50-50 chance of £150 or £50. Do we expect them to take this gamble? In purely wealth terms we would expect the person to be *indifferent* about taking the gamble. This is because the expected wealth associated with the gamble is:

$$W_E = \frac{1}{2}W2 + \frac{1}{2}W3 \qquad (14)$$

$$W_E = \frac{1}{2}(W2 + W3)$$

$$W_E = W1$$

Thus, if they gamble, the equal odds of winning and losing the same amount mean that the average value of their pay-out will be the same as the initial stock of wealth – so the person will be no worse or better off from gambling, and hence indifferent. However, what *expected utility* has to do with is not expected wealth as such, but expected utility *from* that wealth. The expected utility from this gamble is calculated as follows:

$$U_E = \frac{1}{2}[U(W2)] + \frac{1}{2}[U(W3)] \qquad (15)$$

[1] A 'fair' gamble is one where the long run expected payoff is zero. Thus, in this case, the person has a 50 per cent chance of winning or losing an equal amount, so if this game were played numerous times wins would equal losses and the net outcome is zero. The person would be no better or worse off at the end. A coin that was unbalanced would yield more of one outcome than another and the game would not be 'fair' and the person would emerge with a net gain or a net loss.

How expected utility U_E stands in relation to U(W1) depends on the individual's utility function with respect to money. Since expected wealth is equal to actual wealth we might expect U_E = U(W1). But this is not what is usually observed or assumed. As we have seen, the marginal utility of wealth is believed to decline as total wealth increases. This means that, in utility terms, the *decline* in utility from the loss of a certain given amount of wealth (W1 – W2) will exceed the *gain* in utility for an equal increase in wealth (W3 – W1). From the perspective of the decision taker, the fall in utility when the gamble is lost [U(W1) – U(W2)] is greater than the gain in utility should the gamble be won [U(W3) – U(W1)]. In this case, the person would *not* take the gamble even though it was a fair one and we say the person is *risk averse*, where risk is defined as the variability of uncertain outcomes. A risk-averse person prefers a certain sum of money to an equal but uncertain *expected* sum of money.

In the diagram, the expected wealth of the person with a 50-50 chance of W3 or W2 is W1, which is the same, of course, as their assumed initial holding of money wealth. This initial (certain) holding of wealth is associated with a total utility of U(W1), as shown by the U = U(W) function. However, when the level of wealth W1 is *not* certain but is rather the *expected* pay-off when the person gambles at odds of 50-50 between the chances of W3 and W2, then this level of expected wealth yields a *lower* level of utility than when W1 was certain. Because the total utility function has a concave shape, the utility gained from a lottery win is less than the utility lost from a no-win and hence the *expected utility of wealth* W1 (being an average of the two) is less than the *actual* utility from money wealth W1. The straight line between B and C is the expected utility of wealth line, with a mid-point of D corresponding to the expected average level of wealth W1. The expected utility corresponding to the 50-50 gamble between W2 and W3 is U_E(W1), which is below point A, which is the utility corresponding to the actual possession of W1. Thus, a person with wealth W1 who is offered the chance of keeping W1 or gambling for W2 or W3 with an expected average pay-out of W1 would *not* be indifferent between these two results but would *prefer* to keep the certain W1 since the *expected* utility of the gamble would be less. Put simply, while the expected *wealth* from the gamble is the same, the expected *utility* is not. The vertical distance between A and D is a measure of the person's *risk aversion*. The more concave the total utility function, and hence the greater the distance between A and D, the more a person appreciates the *certain* possession of W1 over its *expected* possession through the results of a gamble and hence the *more risk-averse* they are said to be. Note that the smaller the gamble (the shorter the distance between B and C) the less risk averse a person will be: if I have £100 and am offered 50-50 odds on a £5 bet I would be more likely to take the bet than if I were offered 50-50 odds on a £50 bet.

A Numerical Example

Assume a person's utility of wealth function is:

$U = \sqrt{W} = W^{\frac{1}{2}}$ (16)

Suppose a person has wealth of £2,000. They have the chance of a 50-50 bet of winning £1,000 or losing £1,000, ending up with either £3,000 or £1,000. The expected wealth of this bet is:

$W_E = 0.5(3{,}000) + 0.5(1{,}000)$ (17)

$W_E = 1{,}500 + 500 = £2{,}000$

So, their *expected* wealth from the bet is the same as their *current* wealth. If the person was risk-neutral they would be indifferent about taking the bet. But what determines their willingness to take the bet is the *utility* of the various outcomes. Given the utility function, when the person's wealth is a certain £2,000 their utility is:

$U = \sqrt{2000} = 44.7$

Now, compare this to the expected utility of the two possible gamble outcomes:

$U(1000) = \sqrt{1000} = 31.6$

$U(3000) = \sqrt{3000} = 54.8$

Hence, the expected utility from the fair gamble is:

$U_E = 0.5(54.8) + 0.5(31.6)$ (18)

$U_E = 27.4 + 15.8 = 43.2$

Thus, while the expected wealth from the gamble is the same as the certain wealth, the expected utility is *lower*: 43.2 compared to 44.7. This is why a utility maximising risk-averse individual will not accept a fair-bet where the expected payoff is the same as their current wealth. **Figure 3** illustrates this case.

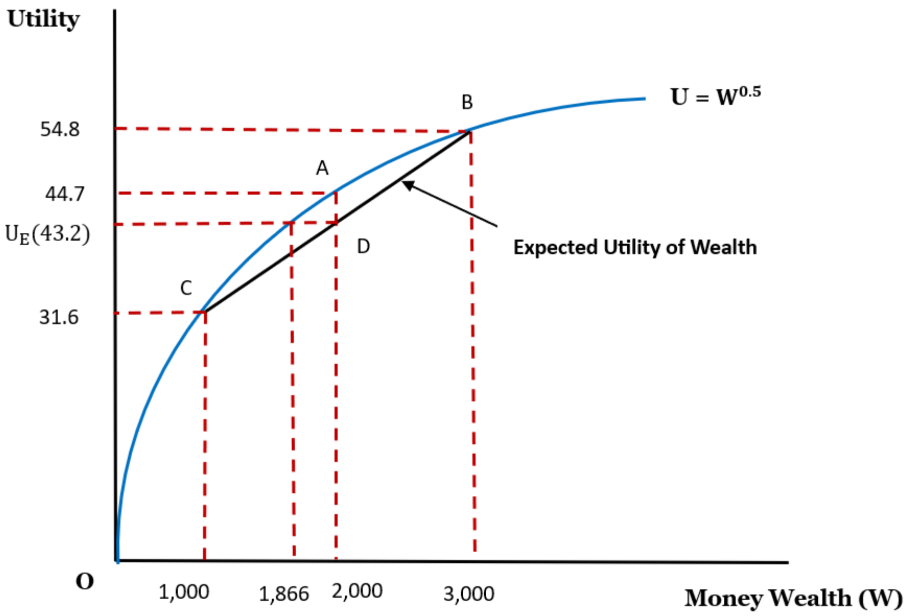

Figure 3. Expected Wealth and Utilities for a £1000 gamble

It can be seen in **Figure 3** that the expected utility from the gamble (43.2) is equal to the utility the person would enjoy from a certain wealth of £1,866. In other words, a risk-averse individual would rather have £1,866 for certain than an expected gamble pay-off of £2,000, where that gamble payoff involved a 50-50 chance of ending up with £3,000 or £1,000. We can use the difference (2,000 – 1,866 = 134) as a measure of how risk averse a person is: how much certain wealth they would forgo to avoid a fair gamble that would leave them, in expected wealth terms, no worse off than their initial starting point. It is this payment they would sacrifice for a certain over an uncertain wealth holding that explains the prevalence of insurance: our individual would pay an insurance premium up to £134 to render their wealth certain as opposed to uncertain.

Thus far we have assumed that the consumer is risk averse – which does seem to be the typical case. But it is not universal and we can obviously identity two alternative possibilities: that the individual is *risk neutral* and that the individual is *risk loving*.

Risk Neutral

An individual whose wealth-utility function is linear is said to be *risk neutral.* In this case, if we take any given certain level of wealth and compare it with the expected wealth from a 50-50 gamble, then the certain utility and the expected utility from the gamble will be the same and the person will be indifferent between a certain income and a random income with fair 50-50 odds. Diagrammatically, the utility function and the expected utility line are identical and a certain income carries no utility premium. **Figure 4** exhibits this situation.

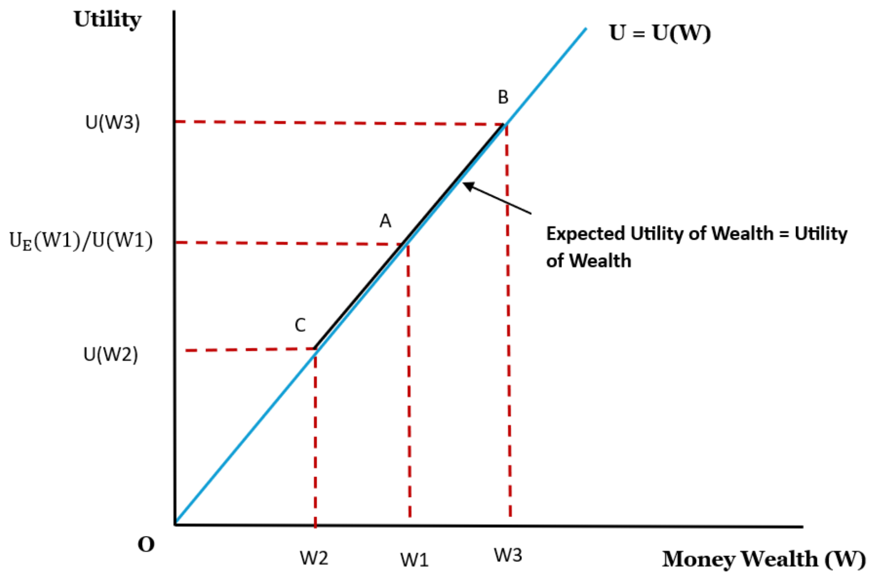

Figure 4. Actual and Expected Utility Functions for a Risk-Neutral Person

We see that a level of certain money wealth W1 corresponds to a total utility of U(W1). A 50-50 gamble of W3-W1 against W1-W2 yields an expected level of wealth W1. Since, in this case, the utility function is linear and not concave the additional utility from a successful gamble is equal to the loss in utility from an unsuccessful gamble, so the expected utility from the gamble is U_E(W1), which is the same as the certain utility at W1. This individual is indifferent between a certain level of wealth and a fair 50-50 gamble.

Risk Loving

Anyone who has waited in a grocery store queue while the person ahead buys lottery tickets or has walked past (or even into) a betting shop will know that not all people are risk averse or even risk neutral. Many people behave as if they positively relish risk and would trade in a certain amount of money wealth for the chance to win more – at the risk of losing the same amount. These are risk-seeking or *risk-loving* people: the kind of person who, if you found them with £100, and offered to toss a coin, so if heads came up you would give them £50 and if tails came up, they would give you £50, would accept the offer. Such people are the opposite of risk averse: they prefer a random distribution of wealth to a certain distribution. The utility function of a risk-loving person is shown below.

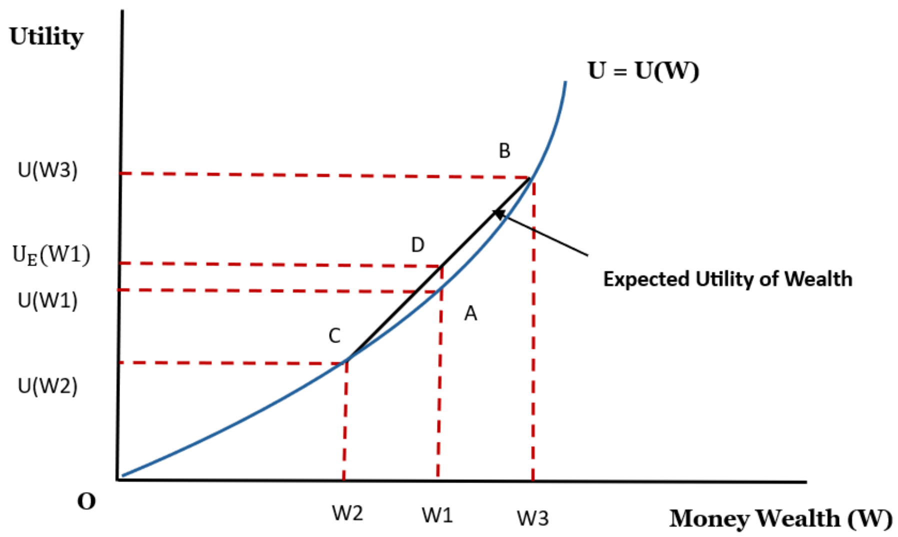

Figure 5. Utility Function of a Risk-Loving Individual

As previously, we suppose a person initially occupies point A, with a level of wealth W1 and a *certain* utility of U(W1). This person is now offered a 50 per cent chance of gaining W3–W1 wealth or of losing W1–W2 (where these amounts are equal in absolute terms). The expected wealth associated with the gamble is:

$W_E = \frac{1}{2}$ W2 + $\frac{1}{2}$ W3 (14)

$W_E = \frac{1}{2}$ (W2 + W3)

W_E = W1

The expected utility from this gamble is:

$U_E = \frac{1}{2}$ [U(W2)] + $\frac{1}{2}$ [U(W3)] (15)

In this case, because the total utility function is *convex* (rather than concave, as before), the expected utility from the gamble, U_E, exceeds the utility associated with a certain level of wealth W1. For the risk-loving individual, the prospective gain of utility resulting from a win (W3–W1) *exceeds* the expected loss in utility from an equal loss of (W1–W2). For a risk-taking individual, the fall in utility when the gamble is lost [U(W1) – U(W2)] is *less than* the gain in utility should the gamble be won [U(W3) – U(W1)]. Hence, a risk-loving individual *will* accept the offer of a fair 50-50 gamble.

In **Figure 5** we see that, at a level of wealth W1, when that level of wealth is *certain* the individual's utility is at point A, corresponding to U(W1), but when that level of wealth is the *expected* result of a gamble that could yield W3 or W2, then their utility is represented by point D, associated with a higher expected utility of U_E(W1). The degree to which expected utility from the gamble exceeds the certain utility from W1 (i.e. D – A) is a measure of how risk-loving this individual is. Put simply, they attach a higher utility to prospect of a win than they do to the equal prospect of a loss –

which is why such people are always drawn to gamble despite the risk of losing their money.

A Numerical Example

Assume a person's utility of money wealth function is:

$U = W^2$ (19)

Suppose their wealth was £200, and they were offered a 50-50 chance of winning or losing £50, leaving them with £150 or £250. The expected wealth from this gamble is:

$W_E = 0.5(250) + 0.5(150) = 125 + 75 = 200$ (20)

Thus, the expected payoff from the gamble equals the individual's current wealth. Their *certain* utility is:

$U(200) = (200)^2 = 40{,}000$

Their *expected* utility is:

$U(150) = (150)^2 = 22{,}500$

$U(250) = (250)^2 = 62{,}500$

Thus, for this person, the expected utility from the gamble is:

$U_E = 0.5(22{,}500) + 0.5(62{,}500) = 11{,}250 + 31{,}250 = 42{,}500$

Here we see that, for this person, the prospective utility from entering into the gamble *exceeds* the utility they expect from keeping their initial holding of wealth. This person would rather gamble than not. For a risk-loving expected-utility maximiser it is rational to gamble current wealth for a 50-50 chance of increased or reduced wealth. And the larger the gamble the larger the more expected utility exceeds actual utility: a risk-loving person would rather gamble larger sums than smaller.