

Game Theory: Preferences and Expected Utility

Branislav L. Slantchev

Department of Political Science, University of California – San Diego

April 19, 2005

Contents.

1	Preferences	2
2	Utility Representation	4
3	Choice Under Uncertainty	5
3.1	Lotteries	5
3.2	Preferences Over Lotteries	7
3.3	vNM Expected Utility Functions	9
3.4	The Expected Utility Theorem	11
4	Risk Aversion	16

1 Preferences

We want to examine the behavior of an individual, called a *player*, who must choose from among a set of outcomes. Begin by formalizing the set of outcomes from which this choice is to be made.

Let X be the (finite) set of outcomes with common elements x, y, z . The elements of this set are mutually exclusive (choice of one implies rejection of the others). For example, X can represent the set of candidates in an election and the player needs to choose for whom to vote. Or it can represent a set of diplomatic and military actions—bombing, land invasion, sanctions—among which a player must choose one for implementation.

The standard way to model the player is with his *preference relation*, sometimes called a *binary relation*. The relation on X represents the relative merits of any two outcomes for the player with respect to some criterion. For example, in mathematics the familiar weak inequality relation, ' $\geq$ ', defined on the set of integers, is interpreted as “integer x is at least as big as integer y ” whenever we write $x \geq y$. Similarly, a relation “is more liberal than,” denoted by ' P ', can be defined on the set of candidates, and interpreted as “candidate x is more liberal than candidate y ” whenever we write xPy .

More generally, we shall use the following notation to denote strict and weak preferences. We shall write $x \succ y$ whenever we mean that x is strictly preferred to y and $x \succeq y$ whenever we mean that x is weakly preferred to y . We shall also write $x \sim y$ whenever we mean that the player is indifferent between x and y . Notice the following logical implications:

$$x \succ y \Leftrightarrow x \succeq y \wedge \neg(y \succeq x)$$

$$x \sim y \Leftrightarrow x \succeq y \wedge y \succeq x$$

$$x \succeq y \Leftrightarrow \neg(y \succ x).$$

Suppose we present the player with two alternatives and ask him to rank them according to some criterion. There are four possible answers we can get:

1. x is better than y and y is not better than x
2. y is better than x and x is not better than y
3. x is not better than y and y is not better than x
4. x is better than y and y is better than x

Although logically possible, the fourth possibility will prove quite inconvenient, so we immediately exclude it with a basic assumption.

ASSUMPTION 1. Preferences are **asymmetric**: There is no pair x and y from X such that $x \succ y$ and $y \succ x$.

We shall also require the player to be able to make judgments about every option that is of interest to us. In particular, he should be able to compare a third option, z , to the original two options. This assumption is quite strong for it implies that the player cannot refuse to rank an alternative.

ASSUMPTION 2. Preferences are **negatively transitive**: If $x \succ y$, then for any third element z , either $x \succ z$, or $z \succ y$, or both.

There are several properties that we now define and all of which are implied by the two assumptions above. The three properties are:

1. *Irreflexivity*: For no x is $x \succ x$.
2. *Transitivity*: If $x \succ y$ and $y \succ z$, then $x \succ z$.
3. *Acyclicity*: If, for a given finite integer n , $x_1 \succ x_2, x_2 \succ x_3, \dots, x_{n-1} \succ x_n$, then $x_n \neq x_1$.

Let's prove that transitivity is implied by negative transitivity and asymmetry. Suppose that the two properties are satisfied, and assume (a) $x \succ y$, and (b) $y \succ z$. Prove $x \succ z$.

- | | |
|---|---|
| 1. $\neg[z \succ y]$ | from (b), asymmetry |
| 2. $[x \succ z] \vee [z \succ y] \vee [[x \succ z] \wedge [z \succ y]]$ | from (a), negative transitivity |
| 3. $[x \succ z] \vee [[x \succ z] \wedge [z \succ y]]$ | from (1) and (2), disjunctive syllogism |
| 4. $[[x \succ z] \vee [x \succ z]] \wedge [[x \succ z] \vee [z \succ y]]$ | from (3), distribution |
| 5. $[x \succ z] \vee [x \succ z]$ | from (4), simplification |
| 6. $x \succ z$ | from (5), tautology |

By the rule of conditional proof, we're done. Hence, asymmetry and negative transitivity imply transitivity. The following proposition states some other implications of these two assumptions for the weak and indifference preference relations.

PROPOSITION 1. *If ' $\succ$ ' is asymmetric and negatively transitive, then*

1. $\succeq$ is **complete**: For all $x, y \in X, x \neq y$, either $x \succeq y$ or $y \succeq x$ or both;
2. $\succeq$ is **transitive**: If $x \succeq y$ and $y \succeq z$, then $x \succeq z$;
3. $\sim$ is **reflexive**: For all $x \in X, x \sim x$
4. $\sim$ is **symmetric**: For all $x, y \in X, x \sim y$ implies $y \sim x$;
5. $\sim$ is **transitive**: If $x \sim y$ and $y \sim z$, then $x \sim z$;
6. If $w \sim x, x \succ y$, and $y \sim z$, then $w \succ y$ and $x \succ z$.

An important concept we hear a lot about is that of rationality. Here is its precise definition in terms of preference relations:

DEFINITION 1. The preference relation $\succeq$ is **rational** if it is complete and transitive.

This means that given *any* two alternatives, the individual can determine whether he likes one at least as much as the other (completeness) and no sequence of pairwise choices will result in a cycle (transitivity). As the above proposition shows, these two properties are implied by the more basic asymmetry and negative transitivity of the strict preference relation. This is the precise definition of rationality we shall use in this course.

If you are interested in social choice theory, the next step is to examine conditions that preference relations must meet in order for the set X to have maximal elements. However, for our purposes, we skip right to utility representations instead.

2 Utility Representation

Consider a set of alternatives X . A utility function $u(x)$ assigns a numerical value to $x \in X$, such that the rank ordering of these alternatives is preserved. More formally,

DEFINITION 2. A function $u : X \rightarrow \mathbb{R}$ is a **utility function representing preference relation** $\succeq$ if the following holds for all $x, y \in X$:

$$x \succeq y \Leftrightarrow u(x) \geq u(y).$$

There are many utility functions that can represent the same preference relation. Note that preferences are *ordinal*, that is, they specify the ranking of the alternatives, but not how far apart they are from each other (intensity). *Cardinal* properties are those that are not preserved under strictly increasing transformations. For example, because the function assigns numerical values to various alternatives, the magnitude of any differences in the utility between two alternatives is cardinal.

We want to use the utility function because instead of examining conditions under which preference relations produce maximal elements for a set of alternatives, it is easier to specify the numerical representation and then apply standard optimization techniques to find the maximum. Thus, the “best” options from the set X are precisely the options that have the maximum utility.

At this point, it is worth emphasizing that *players do not have utility functions*. Rather, they have preferences, which we can represent (for analytical purposes) with utility functions.

We now want to know when a given set of preferences admits a numerical representation. Not surprisingly, the result is closely linked to rationality.

PROPOSITION 2. A preference relation $\succeq$ can be represented by a utility function only if it is rational.

Proof. To prove this proposition, we must show that rationality is a necessary condition for representability. In other words, if there exists a utility function that represents preferences $\succeq$, then $\succeq$ must be complete and transitive. So, assume that $u(\cdot)$ represents $\succeq$.

Completeness. Because $u(\cdot)$ is a real-valued function defined on X , for any $x, y \in X$, either $u(x) \geq u(y)$ or $u(y) \geq u(x)$. Because $u(\cdot)$ is a utility function representing the relation $\succeq$, this implies that either $x \succeq y$ or $y \succeq x$. Hence, $\succeq$ must be complete.

Transitivity. Suppose $x \succeq y$ and $y \succeq z$. Because $u(\cdot)$ represents $\succeq$, we have $u(x) \geq u(y)$ and $u(y) \geq u(z)$. Therefore, $u(x) \geq u(z)$, which in turn implies $x \succeq z$. $\square$

One may wonder if any rational preference ordering $\succeq$ can be represented by some utility function. In general, the answer is no. However, if X is finite, then we can always represent a rational preference ordering with a utility function. The following proposition summarizes these results for the more basic primitives.

PROPOSITION 3. If the set X on which $>$ is defined is finite, then $>$ admits a numerical representation if, and only if, it is asymmetric and negatively transitive.

Sketch of Proof. We have to show *necessity*: $>$ admits a numerical representation only if it is asymmetric and negatively transitive; and *sufficiency*: if $>$ is asymmetric and negatively transitive and X is finite, then $>$ admits a numerical representation.

We proved necessity for $\succeq$, so this step is analogous. Note that it does not require restricting X to be finite. To prove sufficiency, consider the preferences over any two alternatives, and generalize from there (which you can do because X is finite). $\square$

Thus, we know that given a rational preference ordering over a finite set of alternatives, we can always find a utility function that can represent this ordering. Alternatively, if a utility function represents a preference ordering, then that ordering must be rational.

3 Choice Under Uncertainty

Until now, we have been thinking about preferences over alternatives. That is, choices that result in certain outcomes. However, most interesting applications deal with occasions when the player may be uncertain about the consequences of choices at the time the decision is made. For example, when you choose to buy a car, you are not sure about its quality. When you chose to start a fight, you may be uncertain about whether you will win or lose.

3.1 Lotteries

Imagine that a decision-maker faces a choice among a number of risky alternatives. Each alternative may result in a number of possible **outcomes**, but which of these outcomes will occur is uncertain at the time the choice is made.

The von Neumann-Morgenstern (vNM) Expected Utility Theory models uncertain prospects as probability distributions over outcomes. These probabilities are given as part of the description of the outcomes—they are *objective*. Thus, X is now a set of outcomes but there is also a larger set of *probability distributions* over these outcomes denoted by $\mathcal{P}$. We shall mostly be dealing with *simple* probability distributions (these involve only a finite number of possible outcomes).

DEFINITION 3. A **simple probability distribution** p on X is specified by:

1. a finite subset of X , called the **support** of p and denoted by $\text{supp}(p)$; and
2. for each $x \in \text{supp}(p)$, a number $p(x) > 0$, with $\sum_{x \in \text{supp}(p)} p(x) = 1$.

The set of simple probability distributions on X will be denoted by $\mathcal{P}$.

We shall call p (the probability distributions) also *lotteries*, and *gambles* interchangeably. In perhaps simpler words, X is the set of outcomes and p is a set of probabilities associated with each possible outcome. All of these probabilities must be nonnegative and they all must sum to 1. $\mathcal{P}$ then is the set of all such lotteries.

For example, suppose I am participating in a game where I could either roll a die or flip a coin. If I roll the die and the number that comes up is less than 3, I get \$120, otherwise, I get nothing. If I flip the coin and it comes up heads, I get \$100, and if it comes up tails, I get nothing. In our lingo, the set of outcomes is $X = \{0, 100, 120\}$.

The roll of the die is one objective probability distribution (lottery), which assigns the outcome \$0 a probability $2/3$, and the outcome \$120 a probability $1/3$. The support of this lottery consists of these two outcomes only. That is, the outcome \$100 is not in the support, which is another way of saying that this lottery assigns it a probability of zero. Let's denote this lottery by p . We now have $p(0) = 2/3$, $p(120) = 1/3$, and $p(100) = 0$.

The coin flip is another lottery. Let's denote it by q . The support of this lottery consists of the outcomes \$0 and \$100. Because the coin is fair, we have $q(0) = q(100) = 1/2$, and $q(120) = 0$.

In a simple lottery, the outcomes that result are certain. A straightforward generalization is to allow outcomes that are simple lotteries themselves. Suppose now we have two simple probability distributions, p and q , and some number $\alpha \in [0, 1]$. These can form a new probability distribution, r , called a *compound lottery*, written as $r = \alpha p + (1 - \alpha)q$. This requires two steps:

1. $\text{supp}(r) = \text{supp}(p) \cup \text{supp}(q)$
2. for all $x \in \text{supp}(r)$, $r(x) = \alpha p(x) + (1 - \alpha)q(x)$, where $p(x) = 0$ if $x \notin \text{supp}(p)$ and $q(x) = 0$ if $x \notin \text{supp}(q)$

That is, if an outcome is in the support of either one of the simple lotteries, it is also in the support of the compound lottery. The probability associated with an outcome in the compound lottery is a linear combination of the probabilities for this outcome from the simple lotteries.

DEFINITION 4. Given K simple lotteries p_i , and probabilities $\alpha_i \geq 0$ with $\sum_i \alpha_i = 1$, the **compound lottery** $(p_1, \dots, p_K; \alpha_1, \dots, \alpha_K)$ is the risky alternative that yields the simple lottery p_i with probability α_i for all $i = 1, \dots, K$.

Returning to our example, we have two simple lotteries, p (the die) and q (the coin), so $K = 2$. So we need two probabilities, α_1 and α_2 , such that $\alpha_1 + \alpha_2 = 1$. To put it simply, we would have $\alpha_1 = \alpha$, and $\alpha_2 = 1 - \alpha$. The probability α is simply the probability of choosing the die lottery. Its complement is the probability of choosing the coin lottery. Let's suppose that α is determined by the roll of two dice such that α is the probability of their sum equaling either 5 or 6. That is, $\alpha = 1/4$.¹ Thus, $(p, q; 1/4, 3/4)$ is the compound lottery where the simple lottery p occurs with probability $1/4$, and the simple lottery q occurs with probability $1 - 1/4 = 3/4$.

For any compound lottery, we can calculate a corresponding **reduced lottery**, which is a simple lottery that generates the same probability distribution over the outcomes. In other words, we can reduce any compound lottery to a simple lottery.

DEFINITION 5. Let $(p_1, \dots, p_K; \alpha_1, \dots, \alpha_K)$ denote some compound lottery consisting of K simple lotteries. $\hat{p}$ is the reduced lottery that generates the same probability distribution over outcomes, and it is defined as follows. For each $x \in X$,

$$\hat{p}(x) = \sum_{i=1}^K \alpha_i p_i(x).$$

That is, to get a probability of some outcome x in the reduced lottery, you multiply the probability that each lottery p_i arises, α_i , by the probability that p_i assigns to the outcome x , $p_i(x)$, and then adding over all i .

¹You should verify that you can calculate this probability. There are 4 ways to get a sum of 5 and 5 ways to get a sum of 6. The probability of getting either one or the other is $4/36 + 5/36 = 1/4$.

Returning to our example, let's calculate the reduced lottery associated with our compound lottery induced by the roll of the two dice. We have three outcomes, and therefore:

$$\begin{aligned}\hat{p}(0) &= \alpha p(0) + (1 - \alpha)q(0) = (1/4)(2/3) + (3/4)(1/2) = 13/24 \\ \hat{p}(100) &= \alpha p(100) + (1 - \alpha)q(100) = (1/4)(0) + (3/4)(1/2) = 9/24 \\ \hat{p}(120) &= \alpha p(120) + (1 - \alpha)q(120) = (1/4)(1/3) + (3/4)(0) = 2/24\end{aligned}$$

Clearly, $\sum_{x \in X} \hat{p}(x) = 1$, as required. Note further that $\text{supp}(\hat{p}) = \text{supp}(p) \cup \text{supp}(q)$. Thus, the simple lottery that assigns probability $13/24$ to getting nothing, $3/8$ to getting \$100, and $1/12$ to getting \$120, generates the same probability distribution over the outcomes as the compound lottery.

Using this principle, we can define more complicated lotteries that consist of compound lotteries themselves. It should be obvious how to extend the method to these cases.

3.2 Preferences Over Lotteries

We now have a way of modeling risky alternatives. The next step is to define the preferences over them. We shall assume that for any risky alternative, only the reduced lottery over outcomes is of relevance to decision-makers. This is known as the *consequentialist premise*. It does not matter whether probabilities arise from simple, compound, or complex compound lotteries. The only thing that should matter for the decision-maker is the probability distribution over the outcomes, not how it arises.

In our example, the consequentialist hypothesis requires that I view the compound lottery $(p, q; 1/4, 3/4)$ and the reduced lottery $\hat{p} = (13/24, 9/24, 2/24)$ as equivalent.

More generally, it requires that I view *any* two lotteries that generate the same reduced lottery as equivalent. Let's go back to our example with three outcomes. Let the simple lottery $p = (p(0), p(100), p(120))$ denote the probabilities over the three outcomes. So, $p = (2/3, 0, 1/3)$ is our original roll of a single die simple lottery. Figure 1 illustrates a case where two different compound lotteries generate the same reduced lottery. The consequentialist hypothesis requires that I regard both of them as equivalent.




Figure 1: Two Compound Lotteries with the Same Reduced Lottery.

Before we continue with the definition of preference relations over lotteries, we note the special case of a **degenerate lottery**. This is a simple lottery that assigns probability 1 to some outcome, and 0 to all others. We denote it by p^x , where $x \in X$ is the outcome to which the lottery assigns probability 1. For example, going back to the three-outcome case, $p^0 = (1, 0, 0)$, $p^{100} = (0, 1, 0)$, and $p^{120} = (0, 0, 1)$. The degenerate lotteries will prove quite useful for replacing outcomes with equivalent lotteries. Why? Because for any outcome x we

can get the appropriate degenerate lottery p^x . By the consequentialist premise, a decision maker would not care whether he is receiving outcome x directly or as a result of a lottery that yields it with certainty.

We now proceed just like we did in the case of preference relations. We take the set of alternatives, denoted (as you should recall) by $\mathcal{P}$, to be *the set of all simple lotteries over the set of outcomes X* . Next, assume that the decision maker has a preference relation $\succ$ defined on $\mathcal{P}$. (Recall that we can construct both weak and indifference relations from the strict one.) As before, we assume that this relation is rational. It is important to remember that we cannot *derive* the preferences over lotteries from preferences over outcomes, we have to *assume* them as part of the description of the model.

ASSUMPTION 3 (Rationality Axiom). The strict relation $\succ$ on $\mathcal{P}$ is asymmetric and negatively transitive.

Recall that these two assumptions imply that $\succ$ is complete and transitive, and therefore rational. I should note that this requirement of rationality is a bit more demanding than the original one, where the alternatives were outcomes instead of lotteries.

The next assumption is one of continuity. It tells us that very small changes in probabilities do not affect the ordering between two lotteries.

ASSUMPTION 4 (Continuity Axiom). Let $p, q, r \in \mathcal{P}$ be such that $p \succ q \succ r$. Then there exist $\alpha, \beta \in (0, 1)$, such that $\alpha p + (1 - \alpha)r \succ q \succ \beta p + (1 - \beta)r$.

Intuitively, since p is preferred to q , then no matter how bad r is, we can find some mixture of p and r weighted by α that is close enough to p , so that this mixture is better than q . Similarly, we can find another mixture between the two, this time weighted by β that is close enough to r , so that q is preferred to this new mixture. In other words, the preference relation is *continuous*. This assumption is sometimes called the Archimedean axiom.²

The next assumption, sometimes called the *substitution axiom*, is important and controversial. It tells us that if we mix each of two lotteries with a third one, then the preference ordering of the resulting mixtures does not depend on which particular third lottery we used. That is, it is independent of the third lottery.

ASSUMPTION 5 (Independence Axiom). The preference ordering $\succ$ on $\mathcal{P}$ satisfies the *independence axiom* if for all $p, q, r \in \mathcal{P}$ and any $\alpha \in (0, 1)$, the following holds:

$$p \succ q \Leftrightarrow \alpha p + (1 - \alpha)r \succ \alpha q + (1 - \alpha)r.$$

In the compound lottery the player is getting r with the same probability, $(1 - \alpha)$, and thus the “same part” should not affect his preferences. The entire difference should be in how the player evaluates the lotteries p and q against each other. Since $\alpha > 0$, there is a chance that the difference $p \succ q$ will matter, and so the player must prefer the first compound lottery to

²This rules out the so-called **lexicographic** preferences, where one of the outcomes has the highest priority in determining the preference ordering. The name comes from the way a dictionary is organized: the first letter of each word has highest priority, and only then the rest follow. For example, suppose that outcomes are defined by the pair (m, c) , where m is the make of car I get, and c is its color. Define $x \succeq y$ if either “ $[m_x \succ m_y]$ ” or “ $[m_x = m_y] \wedge [c_x \succeq c_y]$ ”. That is, in the preference ordering on the car-color set, the make of car has highest priority. If I prefer one make to another, then the color is irrelevant. If I am indifferent between two makes, only then does color come into play. It is worth noting that even though the lexicographic preference ordering is rational, no utility function exists that can represent it.

the second one. This is sometimes called the substitution axiom because one can substitute any third lottery while preserving the ordering of the original two.

Suppose that $p \succ q$ and $\alpha = 5/6$, and let r denote some other lottery. The compound lottery $5/6p + 1/6r$ can be thought of as arising from a die roll, in which we get p if the number is anything but 6, and r if the number is 6. Similarly, $5/6q + 1/6r$ would be the roll where anything but six gives q , and six gives r . Conditional on a number less than six, we have $5/6p + 1/6r \succ 5/6q + 1/6r$ because $p \succ q$. Conditional on six occurring, the two lotteries are equivalent because they both result in r . The substitution axiom requires that $5/6p + 1/6r \succ 5/6q + 1/6r$, a pretty sensible conclusion.

That is, the preference between two lotteries p and q should determine which of the two the decision maker prefers to have as a part of a compound lottery regardless of the other possible outcome of the compound lottery, say r . This other outcome r is irrelevant for the choice because r does not occur together with either p or q , but only instead of them.

3.3 vNM Expected Utility Functions

The independence axiom is crucial for the representability of preferences over lotteries by a utility function that has an *expected utility form*. Let's see what we mean by that.

DEFINITION 6. Let X denote a finite set of outcomes, and let $N = |X|$ be number of outcomes. The utility function $U : \mathcal{P} \rightarrow \mathbb{R}$ has an **expected utility form** if there is an assignment of numbers $(u_1, \dots, u_N)$ to the N outcomes such that for every simple lottery $p \in \mathcal{P}$, we have

$$U(p) = \sum_{i=1}^N p(x_i) u_i.$$

A utility function with the expected utility form is called a **von Neumann-Morgenstern (vNM) expected utility function**.

That is, suppose that given a set of outcomes, we assign some numerical value (utility) to each outcome. The *expected utility* of every simple lottery over these outcomes is given by the utility of each outcome multiplied by the probability the lottery assigns to it, summed over all possible outcomes. The resulting utility function (which takes a simple lottery as an argument and returns its expected utility) is the vNM expected utility function.³ We want to see if there exist such functions that can represent preferences over lotteries.

For the remainder of this section, you may want to refresh your memory about the rules of the summation notation. Further, you may want to recall what we mean by a **linear function**. A linear function $f(\cdot)$ satisfies two properties:

1. Additivity: $f(x + y) = f(x) + f(y)$
2. Homogeneity: $f(\alpha x) = \alpha f(x)$.

Note that $U(p) = \sum_{i=1}^N p(x_i) u_i$ is a general form for a function that is *linear in the probabilities* $p = (p(x_1), \dots, p(x_N))$. Thus, the vNM expected utility function is linear in the probabilities. In fact, we can prove the following.

³Although this is called the vNM expected utility representation, it is much older, going back to Daniel Bernoulli in the 18th century.

PROPOSITION 4. A utility function $U : \mathcal{P} \rightarrow \mathbb{R}$ has an expected utility form if, and only if, it is linear. That is, if, and only if, it satisfies the following property:

$$U\left(\sum_{i=1}^K \alpha_i p_i\right) = \sum_{i=1}^K \alpha_i U(p_i).$$

Proof. (Sufficiency.) Suppose that $U(\cdot)$ is linear. Let p^i denote the degenerate lottery that yields outcome $i \in X$ with probability one. We can write any lottery p as a convex combination of the degenerate lotteries $(p^1, \dots, p^N)$ as follows: $p = \sum_{i=1}^N p(x_i) p^i$.⁴ We now have:

$$U(p) = U\left(\sum_{i=1}^N p(x_i) p^i\right) = \sum_{i=1}^N p(x_i) U(p^i) = \sum_{i=1}^N p(x_i) u_i.$$

The first equality follows from our definition of p in terms of the degenerate lotteries. The second equality follows from our assumption that $U(\cdot)$ is linear. The last equality follows from the fact that if p^i is a degenerate lottery, then $U(p^i) = u_i$ (this is from the definition of $U(\cdot)$ itself). Thus, we have just shown that if $U(\cdot)$ is linear, then it has the expected utility form.

(Necessity.) Suppose that $U(\cdot)$ has the expected utility form. Consider any compound lottery $p = (p_1, \dots, p_K; \alpha_1, \dots, \alpha_K)$. Its reduced lottery is $\hat{p} = \sum_{k=1}^K \alpha_k p_k$ (recall this from our previous calculations), where $p_k = (p_k(x_1), \dots, p_k(x_N))$. Therefore, we have the following:

$$U\left(\sum_{k=1}^K \alpha_k p_k\right) = \sum_{i=1}^N \left(\sum_{k=1}^K \alpha_k p_k(x_i)\right) u_i = \sum_{k=1}^K \alpha_k \left(\sum_{i=1}^N p_k(x_i) u_i\right) = \sum_{k=1}^K \alpha_k U(p_k).$$

The first equality follows from the assumption that $U(\cdot)$ has the expected utility form. The second equality follows from simple algebra. The third equality follows again from the assumption that $U(\cdot)$ has the expected utility form. Thus, we have just shown that if $U(\cdot)$ has the expected utility form, then it is linear. $\square$

We must emphasize that the expected utility property is a *cardinal* property of utility functions defined on the space of lotteries. This means that it is preserved only by increasing *linear* transformations. We can prove (but won't although Proposition 4 makes the proof straightforward) the following result.

PROPOSITION 5. If $U : \mathcal{P} \rightarrow \mathbb{R}$ is a vNM expected utility function, then $\hat{U} : \mathcal{P} \rightarrow \mathbb{R}$ is another vNM expected utility function if, and only if, there are scalars $a > 0$ and b such that:

$$\hat{U}(p) = aU(p) + b$$

for every $p \in \mathcal{P}$.

⁴This involves vector addition. To see what this statement means, suppose there are $N = 3$ outcomes and let $p = (1/3, 1/3, 1/3)$ be some simple lottery. The three degenerate lotteries are $p^1 = (1, 0, 0)$, $p^2 = (0, 1, 0)$, and $p^3 = (0, 0, 1)$. Defining p in terms of these degenerate lotteries is easy:

$$\sum_{i=1}^N p(x_i) p^i = 1/3 p^1 + 1/3 p^2 + 1/3 p^3 = (1/3, 0, 0) + (0, 1/3, 0) + (0, 0, 1/3) = (1/3, 1/3, 1/3) = p.$$

This result tells us (among other things) that every vNM utility function can be transformed into another vNM utility function through a linear operation. This will be especially useful when we want to rescale utilities to make calculations easier. The point is that there are many vNM utility functions that can represent preferences over lotteries, but the representation is preserved only up to linear transformations.

An important consequence of Proposition 5 is that for vNM expected utility functions, *differences in utilities have meaning* because they imply the preferences over lotteries. For example, suppose there are four outcomes with utilities u_1, u_2, u_3 , and u_4 . We would write the statement “the difference in utility between outcomes x_1 and x_2 is greater than the difference in utility between outcomes x_3 and x_4 ” as $u_1 - u_2 > u_3 - u_4$, which is equivalent to:

$$\frac{1}{2}u_1 - \frac{1}{2}u_2 > \frac{1}{2}u_3 - \frac{1}{2}u_4 = \frac{1}{2}u_1 + \frac{1}{2}u_4 > \frac{1}{2}u_2 + \frac{1}{2}u_3.$$

This now implies that the lottery $p = (\frac{1}{2}, 0, 0, \frac{1}{2})$ is preferred to the lottery $q = (0, \frac{1}{2}, \frac{1}{2}, 0)$. This ranking is preserved by all linear transformations of the vNM utility function.

It is worth emphasizing that it is incorrect to say that a decision-maker prefers an outcome x_1 over outcome x_2 because the utility of x_1 is higher than the utility of x_2 , or $u_1 > u_2$. Rather, *because the decision-maker prefers x_1 to x_2* , the utilities that represent these outcomes are such that $u_1 > u_2$. This is frequent confusion (even in print), which you have to avoid scrupulously. Remember that utility functions are abstract entities that use fictitious numerical values to represent preferences over risky alternatives. The preference orderings are fundamental, and the utility functions are their (non-unique) representations only.

3.4 The Expected Utility Theorem

With all these preliminaries, we are now in position to tackle the central result in the theory of choice under uncertainty. The expected utility theorem first proved by von Neumann and Morgenstern in 1944 is the cornerstone of game theory. It states that if the decision-maker’s preferences over lotteries satisfy assumptions 3, 4, and 5, then these preferences are representable with the expected utility form. That is, if the axioms of rationality, continuity, and independence are accepted, then we can assign real numbers to the outcomes such that the relationship between the expected utilities of any two lotteries preserves their rank ordering. In other words, we can move from using preference orderings to using numbers (utilities), which opens up the entire mathematical arsenal for analysis.

Let’s make sure we understand what exactly it is that we are after here. We know that when we are dealing with preferences over certain outcomes, we can (under specific conditions) represent the outcomes with numbers that preserve the preference ordering over the outcomes. However, in many cases we will be dealing with risky choices that involve uncertain outcomes. We saw how to represent this situation with lotteries. Once we have the lotteries, we *could* calculate the expected utility of playing any lottery *if* the outcomes were numbers. (This we do with a function that has the expected utility form.) However, in most cases outcomes will not be numbers. So, we want to know whether it is possible to represent outcomes with numbers such that the preference ordering over the lotteries is preserved. The following theorem tells us that it is possible. Therefore, we can assign numbers to outcomes and then calculate the expected utility of a lottery in the way we know how. The ordering of the expected utilities preserves the preference ordering of the lotteries. The decision-maker chooses the lottery that yields highest expected utility because this is his most preferred lottery, *an extremely powerful result*. Without further ado, on to the theorem.

THEOREM 1 (Expected Utility Theorem). *A preference relation $\succ$ on the set $\mathcal{P}$ of simple lotteries on X satisfies assumptions 3, 4, and 5 if, and only if, we can assign a number u_x to each outcome $x \in X$ such that for any two lotteries $p, q \in \mathcal{P}$, the following holds:*

$$p \succ q \Leftrightarrow \sum_{x \in \text{supp}(p)} u_x p(x) > \sum_{x \in \text{supp}(q)} u_x q(x).$$

In words, the equation stated in the theorem reads “a lottery p is preferred to lottery q if, and only if, the expected utility of p is greater than the expected utility of q .” Obviously, to calculate the expected utility of a lottery, we must know the utilities attached to the actual outcomes.

The theorem makes two claims. First, it states that if the preference relation $\succ$ is represented by a vNM utility function, then it is rational and satisfies the continuity and independence axioms. This is the necessity (only if) part of the claim, which we won’t prove here.

Second, and more importantly, the theorem states that if the preference relation is rational and satisfies the continuity and independence axioms, then it is representable by a vNM utility function. This is the sufficiency (if) part of the claim, to whose proof we now turn.

The proof of this claim proceeds in several steps. We assume that the preference relation is rational, continuous, and satisfies the independence axiom. The idea is to find a utility function that can represent the preferences over lotteries and meets the condition in the theorem; that is, it has the expected utility form.

We show that for each lottery p , there exists a unique number $\alpha_p \in [0, 1]$ such that the decision-maker is indifferent between p and a simple lottery where the best outcome occurs with probability α_p and the worst outcome occurs with probability $1 - \alpha_p$. In simple lotteries over the best and worst outcomes, the decision-maker always prefers those that yield the best outcome with higher probability. Since all lotteries can be reduced to such simple lotteries, we can compare them in the reduced form (by the consequentialist hypothesis). Thus, each lottery $p \in \mathcal{P}$ has such unique number α_p associated with it. We then show that the utility function that assigns this number is linear, which in turn tells us that it has the expected utility form, which is another way of saying that it allows us to assign numbers to the outcomes themselves. This is what we were after in the first place. And now, on to the proof itself.

Proof. For simplicity, assume that there are at least two outcomes in X , which are best and worst for the player.⁵ Let b denote the best outcome and w denote the worst outcome in X . Now let p^b denote the degenerate lottery that yields b with probability 1, and p^w be the degenerate lottery that yields w with probability 1. Then, for any lottery $p \in \mathcal{P}$, we have $p^b \succeq p$ and $p \succeq p^w$. If $p^b \sim p^w$, the conclusion in the theorem follows trivially because this implies that all lotteries in $\mathcal{P}$ are indifferent. Therefore, from now on, we assume that $p^b \succ p^w$. We first show that a non-degenerate mixture of two lotteries will hold a preference position strictly between these two lotteries.

STEP 1. If $p \succ q$ and $\alpha \in (0, 1)$, then $p \succ \alpha p + (1 - \alpha)q \succ q$.

Since $p \succ q$, the independence axiom implies: $p = \alpha p + (1 - \alpha)p \succ \alpha p + (1 - \alpha)q \succ \alpha q + (1 - \alpha)q = q$. Next, we show that in lotteries involving only the best and the worst outcomes, the player always strictly prefers a higher probability of winning the better outcome.

⁵With our assumption of finite number of outcomes, this can be established as a consequence of the independence axiom.

STEP 2. For any $\alpha, \beta \in (0, 1)$, $\alpha p^b + (1 - \alpha)p^w \succ \beta p^b + (1 - \beta)p^w$ if, and only if, $\alpha > \beta$.

(Sufficiency.) Suppose that $\alpha > \beta$. Note that $\alpha p^b + (1 - \alpha)p^w = \gamma p^b + (1 - \gamma)[\beta p^b + (1 - \beta)p^w]$, where $\gamma = \frac{\alpha - \beta}{1 - \beta} \in (0, 1]$, where the fact that γ is strictly positive follows from $\alpha > \beta$. By Step 1, it follows that $p^b \succ \beta p^b + (1 - \beta)p^w$. Another application of Step 1 implies that $\gamma p^b + (1 - \gamma)[\beta p^b + (1 - \beta)p^w] \succ \beta p^b + (1 - \beta)p^w$, where we make use of the fact that $\gamma > 0$. Thus, we conclude that $\alpha p^b + (1 - \alpha)p^w \succ \beta p^b + (1 - \beta)p^w$.

(Necessity.) Suppose that $\alpha \leq \beta$. If $\alpha = \beta$, then $\alpha p^b + (1 - \alpha)p^w \sim \beta p^b + (1 - \beta)p^w$, so the claim does not hold in this case. Suppose now that $\alpha < \beta$. By the argument in the previous paragraph (reversing α and β), we obtain $\beta p^b + (1 - \beta)p^w \succ \alpha p^b + (1 - \alpha)p^w$, so the claim does not hold in this case either. Therefore, $\alpha > \beta$ is a necessary condition too.

Next, we show that we can calibrate the player preference for any lottery in terms of a lottery involving only the best and worst outcomes:

STEP 3. For any $p \in \mathcal{P}$, there is a unique $\alpha_p \in [0, 1]$ such that $p \sim \alpha_p p^b + (1 - \alpha_p)p^w$.

Existence follows from continuity and the fact that p^b and p^w are the best and worst lotteries, respectively. Uniqueness follows from Step 2. Suppose that there exist two numbers α_1 and α_2 such that $p \sim \alpha_1 p^b + (1 - \alpha_1)p^w$ and $p \sim \alpha_2 p^b + (1 - \alpha_2)p^w$. Suppose without loss of generality that $\alpha_1 > \alpha_2$. By Step 2, $\alpha_1 p^b + (1 - \alpha_1)p^w \succ \alpha_2 p^b + (1 - \alpha_2)p^w$. But p cannot be indifferent to two mixtures, one of which is strictly preferred to the other, a contradiction. Therefore $\alpha_1 = \alpha_2$, i.e. the number is unique.

STEP 4. The function $U : \mathcal{P} \rightarrow \mathbb{R}$ that assigns $U(p) = \alpha_p$, where α_p is defined in Step 3, for all $p \in \mathcal{P}$ represents the preference relation $\succ$.

Step 3 implies that for any two lotteries $p, q \in \mathcal{P}$, we have $p \succ q$ if, and only if, $\alpha_p p^b + (1 - \alpha_p)p^w \succ \alpha_q p^b + (1 - \alpha_q)p^w$. Thus, by Step 2, $p \succ q$ if, and only if, $\alpha_p > \alpha_q$.

Finally, we want to show that this function $U(p) = \alpha_p$ has the expected utility form. From Proposition 4, we know that it would be sufficient to demonstrate that it is linear, for then the vNM form follows.

STEP 5. The utility function $U(p) = \alpha_p$ is linear. That is, for any $p, q \in \mathcal{P}$ and $\alpha \in [0, 1]$, we have:

$$U(\alpha p + (1 - \alpha)q) = \alpha U(p) + (1 - \alpha)U(q).$$

By definition, we have

$$p \sim U(p)p^b + (1 - U(p))p^w,$$

and

$$q \sim U(q)p^b + (1 - U(q))p^w.$$

Applying the independence axiom gives us:⁶

$$\begin{aligned} \alpha p + (1 - \alpha)q &\sim \alpha[U(p)p^b + (1 - U(p))p^w] + (1 - \alpha)q \\ &\sim \alpha[U(p)p^b + (1 - U(p))p^w] + (1 - \alpha)[U(q)p^b + (1 - U(q))p^w] \\ &\equiv [\alpha U(p) + (1 - \alpha)U(q)]p^b + [1 - \alpha U(p) - (1 - \alpha)U(q)]p^w. \end{aligned}$$

⁶The independence axiom for $\succ$ implies the same axiom for the indifference relation. That is, if $p \sim q$, r is any third lottery, and $\alpha \in [0, 1]$, then $\alpha p + (1 - \alpha)r \sim \alpha q + (1 - \alpha)r$. This follows from the independence axiom and the definition of the indifference preference ordering.

That is, the compound lottery in the second line has the same reduced lottery as the compound lottery in the third line above. Thus, we conclude that:

$$\alpha p + (1 - \alpha)q \sim [\alpha U(p) + (1 - \alpha)U(q)]p^b + [1 - \alpha U(p) - (1 - \alpha)U(q)]p^w.$$

We are now done, but let's make the final step a bit clearer. Let $\gamma = \alpha U(p) + (1 - \alpha)U(q)$. This is the probability we assigned to p^b such that $\alpha p + (1 - \alpha)q \sim \gamma p^b + (1 - \gamma)p^w$. The definition of $U(\cdot)$ from Step 4 means that:

$$U(\alpha p + (1 - \alpha)q) = \gamma = \alpha U(p) + (1 - \alpha)U(q),$$

which is what we needed to show. That is, $U(\cdot)$ is linear.

Together, steps 1 to 5 establish the existence of a linear utility function $U(\cdot)$ representing $\succ$. By Proposition 4, this function has the expected utility form. Thus, for each outcome $x \in X$, there exist a number u_x such that the claim of the proposition holds. This completes the proof. $\square$

Let's look at an example. Suppose that we have four possible outcomes, $X = \{x_1, x_2, x_3, x_4\}$, and the decision-maker's preference ordering over these outcomes is:

$$x_1 \succ x_2 \sim x_3 \succ x_4.$$

Suppose now that we also have two lotteries, $p = (1/2, 1/3, 1/6)$ and $q = (1/3, 1/3, 1/3)$ such that $\text{supp}(p) = \{x_1, x_2, x_3\}$, and $\text{supp}(q) = \{x_2, x_3, x_4\}$. Finally, suppose that $p \succ q$.

We want to assign some real number u_x to each outcome $x \in X$ the expected utility of p is strictly greater than the expected utility of q . Recall that to calculate the expected utility of a lottery, you simply multiply the utility of an outcome times the probability of it occurring assigned by the lottery, and sum over all outcomes in the lottery's support. That is, the expected utility of p is $\sum_{x \in \text{supp}(p)} p(x)u_x$. The Expected Utility Theorem tells us that we can do this, and also tells us how to do it.

First, let $p^b = (1, 0, 0, 0)$ be the degenerate simple lottery that yields x_1 (the most preferred outcome) with certainty. Also, let $p^w = (0, 0, 0, 1)$ be the degenerate simple lottery that yields x_4 (the worst outcome) with certainty. We know that there exist two numbers $\alpha_p \neq \alpha_q \in [0, 1]$ such that $p \sim \alpha_p p^b + (1 - \alpha_p)p^w$, and $q \sim \alpha_q p^b + (1 - \alpha_q)p^w$. Furthermore, we know that $\alpha_p > \alpha_q$ because $p \succ q$. So, the utility function $U(\cdot)$ would assign $U(p) = \alpha_p$ and $U(q) = \alpha_q$, and so $U(p) > U(q)$.

Let's assign $u_1 = 1 > u_2 = u_3 = 1/2 > u_4 = 0$. Since U has the expected utility form (proven by the theorem), we have:

$$\begin{aligned} U(p) &= 3/6 u_1 + 2/6 u_2 + 1/6 u_3 = 3/4 = \alpha_p, \\ U(q) &= 1/3 u_2 + 1/3 u_3 + 1/3 u_4 = 1/3 = \alpha_q. \end{aligned}$$

As required, $\alpha_p > \alpha_q$. Is the decision maker indifferent between p and the lottery $\hat{p} = \alpha_p p^b + (1 - \alpha_p)p^w$? If he is, then the expected utilities of these lotteries must be the same. Let's check:

$$\begin{aligned} U(\hat{p}) &= \alpha_p U(p^b) + (1 - \alpha_p)U(p^w) = 3/4(1) + 1/4(0) = 3/4 = U(p) \\ U(\hat{q}) &= \alpha_q U(p^b) + (1 - \alpha_q)U(p^w) = 1/3(1) + 2/3(0) = 1/3 = U(q). \end{aligned}$$

This example might be a bit misleading because we assigned utilities to the outcomes such that each utility was a number between 0 and 1. This ensures that $U(p)$ will also always produce a number in that range. (Can you see why?) In other words, we guaranteed that $U(p)$ would produce something that is a valid probability. This, in turn, allowed us to assign α_p and α_q directly.

The point of our representation, however, is that any numbers that preserve the ordering of the outcomes would do. To this end, let's assign $u_1 = 3 > u_2 = u_3 = 1 > u_4 = -11$. Calculating the expected utilities of the two lotteries now yields:

$$\begin{aligned} U(p) &= \frac{3}{6}u_1 + \frac{2}{6}u_2 + \frac{1}{6}u_3 = 2, \\ U(q) &= \frac{1}{3}u_2 + \frac{1}{3}u_3 + \frac{1}{3}u_4 = -3. \end{aligned}$$

These numbers sure don't look like valid probabilities, even though $U(p) > U(q)$, as required. How do we get α_p and α_q to check our results? This is where the linearity of $U(\cdot)$ comes into play because it allows us to *rescale* the utilities such that they always lie between 0 and 1 while preserving their rank ordering.⁷

Our original utilities lie in the interval $[-11, 3]$. We want to find a function $f(\cdot)$ that would take any number $a \in [-11, 3]$ and return another number $\hat{a} \in [0, 1]$ such that for any two numbers $a, b \in [-11, 3]$ with $a > b$, $f(a) > f(b)$. Note that $f(-11) = 0$ and $f(3) = 1$. Here's the function that we want:

$$f(u) = \frac{u - (-11)}{3 - (-11)} = \frac{u + 11}{14}.$$

Make sure you understand that it works (and why). We now have new utilities, as follows: $\hat{u}_1 = f(u_1) = f(3) = 1$, $\hat{u}_2 = \hat{u}_3 = \frac{6}{7}$, and $\hat{u}_4 = 0$. Note that the original ordering is preserved. We now calculate the expected utilities of the two lotteries:

$$\begin{aligned} \hat{U}(p) &= \frac{3}{6}\hat{u}_1 + \frac{2}{6}\hat{u}_2 + \frac{1}{6}\hat{u}_3 = \frac{13}{14} = \alpha_p, \\ \hat{U}(q) &= \frac{1}{3}\hat{u}_2 + \frac{1}{3}\hat{u}_3 + \frac{1}{3}\hat{u}_4 = \frac{4}{7} = \alpha_q. \end{aligned}$$

These numbers do look like valid probabilities and $\alpha_p > \alpha_q$, as required. Let's do the back check. If the decision maker is indifferent between q and $\hat{q} = \alpha_q p^b + (1 - \alpha_q) p^w$, then the expected utilities of these two lotteries must be equal:

$$\begin{aligned} \hat{U}(\hat{p}) &= \alpha_p \hat{U}(p^b) + (1 - \alpha_p) \hat{U}(p^w) = \frac{13}{14}(1) + \frac{1}{14}(0) = \frac{13}{14} = \hat{U}(p) \\ \hat{U}(\hat{q}) &= \alpha_q \hat{U}(p^b) + (1 - \alpha_q) \hat{U}(p^w) = \frac{4}{7}(1) + \frac{3}{7}(0) = \frac{4}{7} = \hat{U}(q). \end{aligned}$$

And so they are. More generally, we can rescale any set of utilities that reflect the preference ordering over outcomes such that the transformed utilities are constrained to lie in the interval between 0 and 1. This guarantees that the expected utility function will also always return numbers in that interval, and so we can treat these numbers as valid probabilities. But when the expected utility function does return such numbers, they are precisely the numbers we used in the proof of the theorem. That is, each such number, unique for every lottery, is the probability with which the equivalent reduced lottery picks the best outcome while picking the worst outcome with complementary probability.

⁷In fact, we can rescale them to lie between any two arbitrary numbers. The problem set asks you to do precisely that.

The main result in this section is that the preference ordering $\succ$ has an expected utility representation. Each possible outcome has a corresponding utility, and the value of a probability distribution is measured by the expected utility that it provides. This utility function is unique up to a positive affine transformation.

Thus, we began by taking a choice space with a special structure—the set of probability distributions. We then used this structure to pose some axioms about the player's preferences. Using these axioms, we showed that a numerical representation for player preferences can be created that exploits this structure. In other words, given primitive preferences over lotteries, we can represent them with a vNM expected utility function.

We can thus dispense with the preferences and instead use utility functions from now on. However, you should always remember the relationship between them: **It is *not* the case that the decision-maker prefers one lottery over another because it yields a higher expected utility. Rather, the lottery yields a higher expected utility *because it is preferred to the other*. Decision-makers do not have utilities, they have preferences. Utilities are only representations of these basic preferences. In essence, decision-makers behave *as if* they are maximizing expected utility when in fact they are simply choosing on the basis of their preferences.**

4 Risk Aversion

Now that we know that we can represent preferences with real numbers, what else can we say about them that might be helpful? In many settings, individuals seem to display aversion to risk. Let us formalize this notion in terms of lotteries.

Suppose $X \subset \mathbb{R}$. That is, the outcome is a continuous random variable on some subset of the real line. Although our results thus far derived the expected utility representation for the set of finite outcomes, they can be extended to the infinite domain at the expense of technical complications in the proofs (which is why we did not do it). So, suppose outcome is a continuous random variable x . You could think of x in terms of money. Although not necessary, it will be helpful to understand the results. Let $u(\cdot)$ be a utility function defined on sure amounts of money. This is different from the vNM utility function $U(\cdot)$ which is defined on lotteries. We shall call $u(\cdot)$ the *Bernoulli utility function*.

For any lottery p , let p_E represent the expected value of p , or:⁸

$$p_E = \sum_X p(x)x.$$

That is, the expected value of a lottery is the outcome (real number) multiplied by the probability that it occurs assigned to it by the lottery, summed over all possible outcomes. The definition in the continuous case is analogous. Note that this does not use $u(x)$ but the value of the outcome x , which is a real number. To see the difference, note that:

$$U(p) = \sum_X p(x)u(x).$$

That is, while u is defined on outcomes, U is defined on lotteries over these outcomes. Note that p_E , like $U(p)$, is a real number.

For any outcome x , let the lottery $\delta(x)$ be the simple lottery that yields x with certainty. For each lottery p then, $\delta(p_E)$ then is the degenerate lottery that yields the expected value

⁸This is the discrete case, but the definition of the continuous case is analogous.

of the lottery p for sure. The attitude toward risk relates the utility of getting a sure-thing outcome (i.e. an outcome from such a degenerate lottery) to taking a risky gamble of the original lottery.

The expected value of a lottery p , p_E , is a real number, and $u(p_E)$ is the utility associated with this number. To see this, note that $U(\delta(p_E)) = (1)u(p_E) = u(p_E)$. We shall compare this utility with the expected utility of the lottery: $U(p)$. Here is the main result which we shall not prove.

PROPOSITION 6. *For all lotteries $p \in \mathcal{P}$, it is the case that $\delta(p_E) \succeq p$ if, and only if, $u(\cdot)$ is concave; it is the case that $p \succeq \delta(p_E)$ if, and only if, u is convex; and it is the case that $\delta(p_E) \sim p$ if, and only if, u is linear.*

A player who prefers to get the expected value of a lottery for sure instead of taking a risky gamble is said to be *risk averse*. The lack of willingness to take a risky gamble implies a concave utility function $u(\cdot)$ because it requires that $u(p_E) > U(p)$. An example of a concave utility function is $u(x) = \ln(x)$.

If a player prefers to take the risky gamble instead of the expected value of the lottery for sure, then he is *risk acceptant*. The willingness to take a risky gamble implies a convex utility function $u(\cdot)$ because it requires that $u(p_E) < U(p)$. An example of a convex utility function is $u(x) = x^2$.

Finally, a player is *risk neutral* if he is indifferent between the sure thing and the risky gamble. This implies a linear utility function $u(\cdot)$ because it requires that $u(p_E) = U(p)$. An example of a linear utility function is $u(x) = x$.

The three risk attitudes are shown in Figure 2.⁹ Consider the concave utility function (*Risk Averse*) and the two outcomes x and y . Take some $\alpha \in (0,1)$ and consider the lottery $p = \alpha\delta(x) + (1 - \alpha)\delta(y)$. That is, an α chance of x and an $1 - \alpha$ chance of y . The expected value of this lottery is then $p_E = \alpha x + (1 - \alpha)y$.

What does the player prefer: the sure thing $\delta(p_E)$ or the lottery p ? To decide, compare the expected utilities of the two. For $\delta(p_E)$, the utility is

$$U(\delta(p_E)) = u(p_E) = u(\alpha x + (1 - \alpha)y),$$

while the expected utility of p is

$$U(p) = \alpha u(x) + (1 - \alpha)u(y).$$

Because the function $u(\cdot)$ is concave, $u(\alpha x + (1 - \alpha)y) > \alpha u(x) + (1 - \alpha)u(y)$. Thus, the player must be risk averse because the expected utility of the sure thing exceeds the expected utility of the risky gamble. That is, this player must prefer to take the expected value of the lottery than play the lottery.

Conversely, a risk acceptant player prefers to gamble than get the expected outcome of the gamble. Correspondingly, the expected utility of taking the expected value of the lottery with certainty must be less than the expected utility of playing the lottery.

A risk neutral player is indifferent between the two.

To illustrate this with an example, consider the interval $X = [-1, 1]$, and a player with the linear utility function $u(x) = ax + b$, with $a > 0$ (so we get the strictly increasing property

⁹The figure may be a little misleading because it gives the impression that x and y must be such that $u(x)$ and $u(y)$ lie exactly on the 45 degree line. This is not the case: pick any x and y and connect them with a straight line. By concavity, the graph of the risk averse utility function will lie above this segment.




Figure 2: Concavity, Risk Aversion, and Certainty Equivalents. Denote the expected utility of p by $U(p) = \alpha u(x) + (1 - \alpha)u(y)$; the expected value of p by $p_E = \alpha x + (1 - \alpha)y$; the certainty equivalent of p by $C(p)$; and the utility of the expected value of p by $u(p_E)$.

discussed above). Does this player prefer getting .5 for sure or does he prefer to gamble on the lottery p , in which $p(-1) = .25$ and $p(1) = .75$?

$$u(.5) = .5a + b$$

$$U(p) = (.25)[(-1)a + b] + (.75)[(1)a + b] = .5a + b$$

So this utility function represents a player that is indifferent between the sure thing and the gamble on p . As we noticed above, players whose preferences are represented by linear utility functions are risk neutral.

Consider now the utility function $u(x) = (ax + b)^r$ and two players, player A, with $r = 1/2$ and player B, with $r = 2$. Let's examine the same gamble on p as before and let, for simplicity,

$$a = b = 2.$$

$$\begin{aligned} u_A(.5) &= \sqrt{.5a + b} = \sqrt{3} \\ U_A(p) &= (.25)\sqrt{-a + b} + (.75)\sqrt{a + b} = 1.5 \end{aligned}$$

Since $\sqrt{3} > 1.5$, player A must be preferring the sure thing to the gamble. That is, this player is risk averse. It is not difficult to show that u_A is concave.

$$\begin{aligned} u_B(.5) &= (.5a + b)^2 = 9 \\ U_B(p) &= (.25)(-a + b)^2 + (.75)(a + b)^2 = 12 \end{aligned}$$

In this case, player B must be preferring the gamble to the sure thing (since $12 > 9$). That is, this player is risk acceptant. Again, it is not hard to show that u_B is convex.

More generally, for any $r < 1$, the utility function $u(x) = (ax + b)^r$ will be concave and thus useful to represent a risk averse player. For any $r > 1$, the utility function will be convex and thus useful to represent a risk acceptant player. For $r = 1$, the utility function is linear and can be used to represent a risk neutral player.

Sometimes it will be useful to find an outcome, whose utility is the same as the expected utility of some lottery. In other words, we go backwards: we start with some utility and we wish to find the outcome that produces it. Consider two outcomes $x, y \in X$ and the utility function $u : X \rightarrow \mathbb{R}$. Let p be some lottery such that x occurs with probability α and y occurs with probability $1 - \alpha$. From the intermediate value theorem of calculus, we know that for every $\alpha \in [0, 1]$, there is some value z with $u(z) = \alpha u(x) + (1 - \alpha)u(y) = U(p)$. From the expected utility representation, we know that for any such z , it is the case that $\delta(z) \sim p$. To see this, note that $U(\delta(z)) = u(z) = U(p)$. This z is called the *certainty equivalent* of p .

Almost all utility functions that we shall deal with in this course will be continuous and strictly increasing. The following proposition then gives us a useful result that states that this certainty equivalent always exists and is unique.

PROPOSITION 7. *If $u(\cdot)$ is continuous and strictly increasing, and if X is an interval of $\mathbb{R}$, then every $p \in \mathcal{P}$ has a unique certainty equivalent.*

Referring back to Figure 2, the certainty equivalent of p , denoted by $C(p)$, yields the same utility as the expected utility of p . Since the player is risk averse, $C(p) \leq p_E$.