Expected Utility Framework

Preferences

We want to examine the behavior of an individual, called a *player*, who must choose from among a set of outcomes.

- 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.

- 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.

We shall write $x \succ y$ whenever, x is strictly preferred to y and $x \succeq y$ whenever, x is weakly preferred to y. We shall also write $x \sim y$ whenever the player is indifferent between x and y.

Notice the following logical implications:

- $x \succ y \Leftrightarrow (x \succeq y) \wedge \neg(y \succeq x)$,
- $x \sim y \Leftrightarrow (x \succeq y) \wedge (y \succeq x)$, and
- $x \succeq y \Leftrightarrow \neg(y \succ x)$.

Suppose we present the player with two alternatives and ask him to rank them according to some criterion. The assumptions that follow put structure on the player's preferences.

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

A player should also 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 ("or" here is used as in the everyday use indicating exclusion).

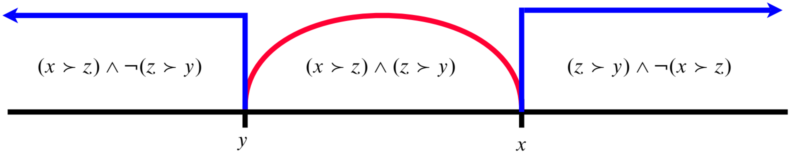

PROPOSITION 1. If the preference relation is asymmetric and negatively transitive, then

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 $x \succ y$ and $y \succ z$, then $x \succ z$.

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

St. Petersburg Paradox

The paradox takes its name from Daniel Bernoulli, one-time resident at St. Petersburg, though the problem was originally proposed by Gabriel Cramer.

A fair coin is tossed at each stage. The pot starts at $2 and is doubled every time a head appears. The first time a tail appears, the game ends and the player wins whatever is in the pot. How much would you pay to play this game?

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

Nicolas Bernoulli (Daniel's cousin) conjectured instead that people will neglect unlikely events (see Prospect Theory).

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.

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

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

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. □

One may wonder if any rational preference ordering $\succeq$ can be represented by some utility function. In general, the answer is no.

PROPOSITION 3. If the set X on which $\succ$ is defined is finite, then $\succ$ admits a numerical representation if, and only if, it is asymmetric and negatively transitive (hence rational). □

Utility Representation

Once a preference is represented by a utility function, then we can formulate the consumer problem as a constrained optimization problem:

$$\max_{x \in X} \; u(x) \text{ such that } p \cdot x \leq w,$$

which may be easily solved analytically or numerically.

Examples of Utility Functions

- Cobb-Douglas utility function: $u(x_1, x_2) = x_1^{\alpha} x_2^{1-\alpha}$ for $\alpha \in (0, 1)$;
- Quasi-linear utility function: $u(x, m) = v(x) + m$;
- Leontief utility function: $u(x_1, x_2) = \min\{x_1, x_2\}$

Utility Representation (Cont.)

If we want well-behaved indifference curves, we need to assume further that preferences are convex and monotonic as follows:

1. Convexity
   for $x, y \in X$, where $x \sim y$, for every $t \in [0,1]$:
   $tx + (1-t)y \succeq x$;

2. Monotonicity
   for $x, y \in X$, where $x \geq y$ and $x \neq y$:
   $x \succeq y$.

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.

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 Framework models uncertain prospects as probability distributions over outcomes. These probabilities are given as part of the description of the outcomes.

- X is a set of outcomes.
- $\wp$ is a set of *probability distributions* over these 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 $supp(p)$; and
2. for each $x \in supp(p)$, a number $p(x) > 0$, with $\sum_{x \in supp(p)} p(x) = 1$.

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. P then is the set of all such lotteries.

Example

Lottery p: Roll a die and if the number that comes up is less than 3, you get $120; otherwise, you get nothing.

Lottery q: Flip a coin and if it comes up heads, you get $100; otherwise, you get nothing.

$X = \{0, 100, 120\}$

$p = (\frac{2}{3}, 0, \frac{1}{3})$ since $p(0) = \frac{2}{3}$, $p(100) = 0$, and $p(120) = \frac{1}{3}$.

$supp(p) = \{0, 120\}$

$q = (\frac{1}{2}, \frac{1}{2}, 0)$ since $q(0) = \frac{1}{2}$, $q(100) = \frac{1}{2}$, and $q(120) = 0$.

$supp(q) = \{0, 100\}$

Compound Lotteries

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

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

Example

The probability α is the probability of choosing the die lottery. Its complement $(1-\alpha)$ 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.

$\alpha = \frac{1}{4}$ and $1-\alpha = 1-\frac{1}{4} = \frac{3}{4}$

Thus, $(p, q; \frac{1}{4}, \frac{3}{4})$ is the compound lottery where the simple lottery p occurs with probability $\frac{1}{4}$, and the simple lottery q occurs with probability $\frac{3}{4}$.

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.

DEFINITION 5. Let $(p_1, ..., p_K; \alpha_1, ..., \alpha_K))$ 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 $x \in X$,

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

Example

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:

$$\hat{p}(0) = \alpha \cdot p(0) + (1-\alpha) \cdot q(0) = \frac{1}{4} \cdot \frac{2}{3} + \frac{3}{4} \cdot \frac{1}{2} = \frac{13}{24}$$

$$\hat{p}(100) = \alpha \cdot p(100) + (1-\alpha) \cdot q(100) = \frac{1}{4} \cdot 0 + \frac{3}{4} \cdot \frac{1}{2} = \frac{9}{24}$$

$$\hat{p}(120) = \alpha \cdot p(120) + (1-\alpha) \cdot q(120) = \frac{1}{4} \cdot \frac{1}{3} + \frac{3}{4} \cdot 0 = \frac{2}{24}$$

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

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.

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.

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 $\wp$, 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 $\wp$.

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. These assumptions are represented by a set of three axioms:

AXIOM R (RATIONALITY). The strict relation $\succ$ on $\wp$ is asymmetric and negatively transitive.

AXIOM C (CONTINUITY). Let $p, q, r \in \wp$ be such that $p \succ q \succ r$. Then there exists $\alpha, \beta \in (0,1)$, such that $\alpha p + (1-\alpha) r \succ q \succ \beta p + (1-\beta) r$.

Axiom C rules out **lexicographic preferences**. Lexicographic preferences are preferences where one of the outcomes has the highest priority in determining the preference ordering. Let $p = (p_1, p_2)$ and $q = (q_1, q_2)$ be $p, q \in \wp$. We say that p is *lexicographically preferred to* q, and write $p \succ q$ if and only if either $(p_1 > q_1)$ or $(p_1 = q_1$ and $p_2 \geq q_2)$.

The third assumption is 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.

AXIOM I (INDEPENDENCE). The preference ordering $\succ$ on $\wp$ satisfies the *independence axiom* if for all $p, q, r \in \wp$ and any $\alpha \in (0, 1]$, the following holds:

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

This is sometimes called the *Substitution Axiom*. Although Axiom I seems quite reasonable, it might be among the easiest to violate empirically (see Allais' Paradox).

Allais' Paradox

Allais, M.: Le Comportement de l'Homme Rationnel devant le Risque: Critique des Postulats et Axiomes de l' ecole Americaine. *Econometrica*, Vol. 21, No.4. (1953), $503 - 546$.

Problem 1: Choose between

A: $2,500 with probability 0.33, $2,400 with probability 0.66, & $0 with probability 0.01.

B: $2,400 with certainty.

Problem 2: Choose between

C: $2,500 with probability 0.33, & $0 with probability 0.67.

D: $2,400 with probability 0.34, & $0 with probability 0.66.

Allais' Paradox (Cont.)

Problem 1 indicates that

$u(\$2,400) > 0.33u(\$2,500) + 0.66u(\$2,400)$ or $0.34u(\$2,400) > 0.33u(\$2,500)$.

Problem 2 indicates that

$0.34u(\$2,400) < 0.33u(\$2,500)$. $\rightarrow\leftarrow$

The Von Neumann-Morgenstern Utility Theorem

- Recall that we were able to prove that we can represent rational preferences with numbers (under some conditions). We now want to see whether we can do "the same" for the lotteries.

- The question is whether given preferences over lotteries we can guarantee that we can find numbers that would make this calculation work such that the ranking of the expected utilities of two lotteries will be the same as their preference ordering. Let's first formally define the function we are interested in using.

DEFINITION 6. Let X denote a finite set of outcomes. The utility function $U : \wp \to \mathbb{R}$ has the **expected utility form** if there is a (Bernoulli) utility (payoff) function $u : X \to \mathbb{R}$ that assigns real numbers to outcomes such that for every simple lottery $p \in \wp$, we have

$$U(p) = \sum_{x \in X} \text{p(x)u(x)}.$$

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

- Although this is called the vNM expected utility representation, it is much older, going back to Gabriel Cramer and Daniel Bernoulli in the 18^{th} century.
- Nowadays, vNM is used for expected utility functions, and Bernoulli usually refers to payoffs for certain outcomes.

The *Expected Utility Theorem* first proved by von Neumann in 1944 is the cornerstone of game theory. It states that if the decision-maker's preferences over lotteries satisfy Axioms R, C and I, then these preferences are representable with a function that has the expected utility form.

We know that when we are dealing with rational preferences over certain outcomes, we can 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. Now we are going to see how to represent the preferences over these lotteries with numbers that are produced by calculating expected utilities of these lotteries.

Example

Suppose that each day I leave home to come to my office, I face three possible outcomes: getting to work safely, having a minor accident and having a major accident. Suppose that I get to work either driving a scooter or driving a car. Each mode of transportation is associated with a lottery over these three outcomes. So, suppose the following two lotteries summarize all of these probabilities: $d = (0.87, 0.12, 0.01)$ if I drive the car, and $r = (0.94, 0.04, 0.02)$ if I ride the scooter. Since I ride the scooter every day, it must be the case that $r \succ d$.

The question now is, can we find numbers (u_1, u_2, u_3) such that when we assign them to the outcomes (safe, minor accident, major accident) and calculate the expected utilities, we would get $U(r) > U(d)$ (that is, preserve the preference ordering)?

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 the highest expected utility because this is his most preferred lottery.

THEOREM 1 (EXPECTED UTILITY THEOREM). *A preference relation* $\succ$ *on the set* $\wp$ *of simple lotteries on* X *satisfies Axioms R, C and I if, and only if, there exists a function that assigns a real number to each outcome,* $u : X \to \mathbb{R}$ *such that for any two lotteries* $p, q \in \wp$*, the following holds:*

$$p \succ q \Leftrightarrow \sum_{x \in X} p(x)u(x) > \sum_{x \in X} q(x)u(x).$$

- 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, that is, the number that $u(x)$ assigns to outcome x.
- The theorem makes two claims. First, it states that only if the preference relation $\succ$ can be represented by a vNM utility function, then must satisfy Axioms R, C and I. This is the necessity (only if) part of the claim.
- Second, and more importantly, the theorem states that if the preference relation satisfies Axioms R, C and I, then must be representable by a vNM utility function. This is the sufficiency (if) part of the claim.

Final Remarks

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 choosing on the basis of their preferences.

- Yet, there are some serious and valid criticisms of Expected Utility Theory.
- For one, the Allais' Paradox indicates that people do not have an intuitive feel for probabilities so they do not recognize independence when they deal with compound lotteries.
- Another problem is the way people interpret information which seems to be affected by how this information is presented to them (see Prospect Theory).