Choice Under Uncertainty

Lotteries

Without uncertainty, there is no need to distinguish between a consumer's choice between alternatives and the resulting outcome. A consumption bundle is the choice and it is also the outcome providing utility.

With uncertainty, the set of alternatives and the set of outcomes could be different. For example, a consumer could purchase an insurance policy and the outcome depends on whether she has an accident.

Denote the set of possible outcomes (or consequences) by C. For simplicity, we assume that the set of outcomes is finite, so $C = \{1, 2, ..., N\}$.

We use the term *lottery* to refer to a risky alternative received by or chosen by the decision maker.

Def. A simple lottery L is a list $L = (p_1, ..., p_N)$, where p_n is interpreted as the probability of outcome n occurring.

Geometrically, a lottery is a point in the $(N-1)$ dimensional simplex in $\Re^N$, $\Delta^{N-1} = \{p \in \Re^N_+ : \sum_{n=1}^N p_n = 1\}$.

Note: We usually reserve the letter p to denote prices. Only in the section on choice under uncertainty does p refer to a probability.

A *compound lottery* is a lottery in which the outcomes are themselves lotteries. For example, a decision to go to grad school might yield a professor job at a top-tier department with probability 0.1, a professor job at a second-tier department with probability 0.4, a professor job at a third-tier department with probability 0.2, a job at a consulting firm or government agency with probability 0.2, and a job driving a taxi with probability 0.1. Each of these jobs themselves are lotteries that offer various earnings profiles with various probabilities (and require future choices!).

Def. Given K simple lotteries $L_k = (p_1^k, ..., p_N^k)$ for $k = 1, ..., K$, and given probabilities $\alpha_k \geq 0$ with $\sum_{k=1}^K \alpha_k = 1$, the compound lottery $(L_1, ..., L_K; \alpha_1, ..., \alpha_K)$ is the risky alternative yielding the lottery L_k with probability α_k, for $k = 1, ..., K$.

Any compound lottery can be reduced to a corresponding simple lottery that generates the same probability distribution over outcomes in C. The probability of outcome n is given by $p_n = \alpha_1 p_n^1 + \alpha_2 p_n^2 + ... + \alpha_K p_n^K$. The term $\alpha_k p_n^k$ is the probability that lottery k is selected in the compound lottery **and** that outcome n is selected in lottery k.

More complicated compound lotteries are possible, but they can all be reduced to a unique simple lottery.

Example: $N = 3$ and $K = 4$.

L_1 is the degenerate simple lottery yielding outcome 1 for sure, $L_1 = (1, 0, 0)$. $L_2 = (\frac{1}{3}, \frac{1}{3}, \frac{1}{3})$, $L_3 = (\frac{1}{2}, 0, \frac{1}{2})$, $L_4 = (\frac{1}{2}, \frac{1}{2}, 0)$.

$\alpha = (\alpha_1, \alpha_2, \alpha_3, \alpha_4) = (\frac{1}{6}, \frac{1}{6}, \frac{1}{3}, \frac{1}{3})$.

Then

$$p_1 = \frac{1}{6} \cdot 1 + \frac{1}{6} \cdot \frac{1}{3} + \frac{1}{3} \cdot \frac{1}{2} + \frac{1}{3} \cdot \frac{1}{2} = \frac{5}{9}$$

$$p_2 = \frac{1}{6} \cdot 0 + \frac{1}{6} \cdot \frac{1}{3} + \frac{1}{3} \cdot 0 + \frac{1}{3} \cdot \frac{1}{2} = \frac{2}{9}$$

$$p_3 = \frac{1}{6} \cdot 0 + \frac{1}{6} \cdot \frac{1}{3} + \frac{1}{3} \cdot \frac{1}{2} + \frac{1}{3} \cdot 0 = \frac{2}{9}$$

Preferences Over Lotteries

We will impose assumptions or axioms that preferences over lotteries must satisfy, ultimately leading to utility functions.

Axiom (Consequentialism)–For any risky alternative that might be a compound lottery, only the reduced simple lottery specifying probabilities of final outcomes (consequences) is relevant for the decision maker's (DM's) preferences.

Note: One can imagine situations in which the process matters and not just the final outcome. Machina's ice cream cone example.

Note: One can also imagine that DM's can have a difficult time computing the reduced lottery, so that systematic departures from consequentialism are possible.

If preferences satisfy consequentialism, then we can take the set of risky alternatives to be the set of simple lotteries over C, which we denote by $\mathcal{L}$. The DM's preferences are captured by a (weak) preference relation $\succeq$ on $\mathcal{L}$.

Axiom (Continuity)–For any $L, L', L'' \in \mathcal{L}$, the following sets are closed subsets of $[0, 1]$:

$$\begin{aligned} \{\alpha &\in [0,1] : \alpha L + (1-\alpha)L' \succeq L''\} \\ \{\alpha &\in [0,1] : L'' \succeq \alpha L + (1-\alpha)L'\} \end{aligned}$$

Continuity says that if one lottery is strictly preferred to another, then it remains strictly preferred if we make sufficiently small changes to the probabilities.

Suppose the outcomes are {go out and have fun, go out and have an accident, stay at home}. $L = (1, 0, 0), L' = (0, 1, 0), L'' = (0, 0, 1)$. Then someone not willing to take any risk might have preferences satisfying $\{\alpha \in [0, 1] : L'' \succeq \alpha L + (1-\alpha)L'\} = [0, 1)$, which is not a closed set. Not continuous.

Under continuity, then the preference relation $\succeq$ can be represented by a utility function $U : \mathcal{L} \to \Re$, such that $L \succeq L'$ if and only if $U(L) \geq U(L')$.

Axiom (Independence)–For any $L, L', L'' \in \mathcal{L}$ and $\alpha \in (0, 1)$, we have

$$L \succeq L' \text{ iff } \alpha L + (1 - \alpha)L'' \succeq \alpha L' + (1 - \alpha)L''.$$

In other words, if L is preferred to L', then introducing a fixed probability of a third option should not change the preference.

Independence is a reasonable axiom, but there may be circumstances in which it fails. We will discuss this later but for now consider this example of disappointment due to Machina. $C = \{\text{trip to Venice, see a beautiful movie about Venice, stay home}\}$. If you prefer $(0, 1, 0)$ to $(0, 0, 1)$, then independence requires $(.99, .01, 0) \succeq (.99, 0, .01)$. But if you do not win the trip to Venice, will you really want to sit through the movie?

Def. The utility function $U : \mathcal{L} \rightarrow \Re$ has an *expected utility* form if there is an assignment of numbers to each outcome, $(u_1, ..., u_N)$, such that for every simple lottery $L = (p_1, ..., p_N)$, we have

$$U(L) = \sum_{n=1}^{N} p_n u_n.$$

A utility function with the expected utility form is called a von Neumann-Morgenstern (or v.N-M) expected utility function.

Note: Expected utility is an appropriate terminology, because the utility of each outcome is well-defined, and the utility of a lottery is the expectation of the utility of the outcome.

Proposition (Linearity): $U : \mathcal{L} \rightarrow \Re$ has an expected utility form if and only if it is linear. That is,

$$U\left(\sum_{k=1}^{K} \alpha_k L_k\right) = \sum_{k=1}^{K} \alpha_k U(L_k)$$

holds for any compound lottery $(L_1, ..., L_K; \alpha_1, ..., \alpha_K)$.

proof sketch $\Longrightarrow$ If U has the expected utility form, we can write

$$\begin{aligned}
U\left(\sum_{k=1}^{K} \alpha_k L_k\right) &= \sum_{n=1}^{N}\left(\sum_{k=1}^{K} \alpha_k p_n^k\right) u_n \\
&= \sum_{k=1}^{K}\sum_{n=1}^{N} \alpha_k p_n^k u_n \\
&= \sum_{k=1}^{K} \alpha_k \sum_{n=1}^{N} p_n^k u_n = \sum_{k=1}^{K} \alpha_k U(L_k).
\end{aligned}$$

proof sketch $\Longleftarrow$ If

$$U\left(\sum_{k=1}^{K}\alpha_k L_k\right)=\sum_{k=1}^{K}\alpha_k U(L_k)$$

holds for any compound lottery $(L_1,...,L_K;\alpha_1,...,\alpha_K)$, consider the simple lottery L as a degenerate compound lottery over certain outcomes, $L=(L^1,...,L^N;p_1,...,p_N)$, where $L^n=(0,...,0,1,0,...,0)$. We have

$$U(L)=U(\sum_{n=1}^{N}p_n L^n)=\sum_{n=1}^{N}p_n U(L^n).$$

Just define $u_n=U(L^n)$, and we are done.

Proposition: Suppose that $U : \mathcal{L} \to \Re$ is a v.N-M expected utility function for the preference $\succsim$. Then $\widetilde{U}$ is also a v.N-M expected utility function for the preference $\succsim$ if and only if there are scalars $\beta > 0$ and γ such that $\widetilde{U}(L) = \beta U(L) + \gamma$ holds for every $L \in \mathcal{L}$.

Note: When making choices under uncertainty, v.N-M expected utility pins down the cardinal utility function over outcomes. Statements like "the utility of outcome 1 exceeds the utility of outcome 2 by more than the difference in utility between outcome 3 and outcome 4" make sense. To see this, $u_1 - u_2 > u_3 - u_4$ is equivalent to $\frac{1}{2}u_1 + \frac{1}{2}u_4 > \frac{1}{2}u_2 + \frac{1}{2}u_3$, which is equivalent to the statement that the lottery $(\frac{1}{2}, 0, 0, \frac{1}{2})$ is preferred to the lottery $(0, \frac{1}{2}, \frac{1}{2}, 0)$.

Expected utility is a cardinal notion, like temperature. It can be rescaled, but for a given preference relation, specifying the utility of two outcomes pins down the entire function. Unlike temperature which can be objectively measured, utility functions vary across individuals.

The Expected Utility Theorem

If preferences can be represented by a v.N-M expected utility function, then they must be continuous because linear functions are continuous.

Preferences must also satisfy the independence axiom, again because of the linear structure. That is, it is clear that

$$\sum_{n=1}^{N} p_n u_n \geq \sum_{n=1}^{N} p'_n u_n$$

holds if and only if we have

$$\alpha \sum_{n=1}^{N} p_n u_n + (1-\alpha) \sum_{n=1}^{N} p''_n u_n \geq \alpha \sum_{n=1}^{N} p'_n u_n + (1-\alpha) \sum_{n=1}^{N} p''_n u_n$$

The Expected Utility Theorem shows that the converse is also true.

Theorem: Suppose that $\succsim$ on the space of lotteries $\mathcal{L}$ satisfies the continuity and independence axioms. Then $\succsim$ can be represented by a v.N-M expected utility function. That is, there is an assignment of numbers to each outcome, $(u_1, ..., u_N)$, such that for any two simple lotteries $L = (p_1, ..., p_N)$ and $L' = (p'_1, ..., p'_N)$, we have

$$L \succsim L' \text{ if and only if } \quad \sum_{n=1}^{N} p_n u_n \geq \sum_{n=1}^{N} p'_n u_n.$$

Intuition for the Theorem with $N = 3$:

Continuity implies that preferences over lotteries can be represented by a utility function, so we can draw indifference curves in the lottery simplex (triangle).

Independence implies that the indifference curves must be linear. If $L \sim L'$ then $\alpha L + (1-\alpha)L \sim \alpha L' + (1-\alpha)L$, so any point on the segment connecting L and L' must be on the same indifference curve as L.

Independence implies that the indifference curves must all be parallel. Consider the triangle formed by $L \sim L'$ and a third lottery L''. By independence, $\frac{1}{3}L + \frac{2}{3}L'' \sim \frac{1}{3}L' + \frac{2}{3}L''$, so the lottery$\frac{2}{3}$ of the way from L to L'' is indifferent to the lottery $\frac{2}{3}$ of the way from L' to L''. This indifference curve must be parallel to the line connecting L and L'.

Parallel straight line indifference curves admit a v.N-M representation. (Normalize the best outcome to have utility 1 and the worst to have utility 0. The slope pins down the utility of the remaining outcome.)

Discussion of Expected Utility

1. Analytically easy. Normatively appealing–violations are mistakes that should be corrected. What would you replace it with?

2. The Allais Paradox. Outcomes are prizes: $C = \{\$2,500,000; \$500,000; \$0\}$. Most people strictly prefer $L_1 = (0, 1, 0)$ to $L_1' = (0.10, 0.89, 0.01)$. However, most people also strictly prefer $L_2' = (0.10, 0, 0.90)$ to $L_2 = (0, 0.11, 0.89)$. These preferences are inconsistent with expected utility.

$$\begin{aligned} L_1 &\succ L_1' \text{ implies } u_{.5} > 0.10u_{2.5} + 0.89u_{.5} \\ \text{which implies } 0.11u_{.5} &> 0.10u_{2.5} + 0.01u_0 \end{aligned}$$

$$\begin{aligned} L_2' &\succ L_2 \text{ implies } 0.10u_{2.5} + 0.90u_0 > 0.1 \\ \text{which implies } 0.11u_{.5} &< 0.10u_{2.5} + 0.01u_0 \end{aligned}$$

Does this show that people are not expected utility maximizers? Maybe they overstate the importance of small risks and would prefer L_1' upon reflection. Maybe they would feel enormous regret if they could have guaranteed $\$500,000$ but chose L_1' and wound up with zero–the utility of a lottery depends on the other choices available. Maybe we do not trust the referee who is running the lottery.

3. Induced Preferences. Some apparent departures from expected utility are due to the decision problem not being fully modeled. For example, you are invited to a dinner party that either serves fish (outcome F) or meat (outcome M). Suppose that you are indifferent between fish and meat, $u_F = u_M$.

Then a v.N-M utility function should assign the same utility to all lotteries. However, you strictly prefer the lotteries $(1, 0)$ and $(0, 1)$ to $(\frac{1}{2}, \frac{1}{2})$, because you can bring a bottle of white wine if fish is being served and red wine if meat is being served.

Outcomes should be modeled as including the kind of wine brought to the party and the dish being served.