Decision Theory

1.4 Utility

George Mason University, Spring 2018

Where do the utilities come from that enter into expected utility calculations and why should we take their maximization as a guide to rational action?

Today we look at an influential attempt to prove that it's possible to represent an agent's preferences numerically as an interval utility function. In demonstrating how to construct such a function, Von Neumann and Morgenstern also prove that maximizing expected utility is rational because the utility of an act is equal to the expected utility of its outcomes, and maximizing utility is just doing what you want to do.

We begin with an intuitive presentation.

First consider a problem with EMV. You're given the *only* ticket for the *one* and *only* lottery your state will ever run. It gives you a one in a million chance of winning \$1,000,000. Since you lose nothing if you fail to win, EMV(ticket)=\$1.

But you'll never win \$1, only either \$1,000,000 or nothing. So what's the connection between these prizes (the values of the outcomes) and the EMV? Why would it be rational to sell your ticket for \$2?

We might want to answer in terms of averages, i.e., what would happen to many people playing this same lottery or to you if you played this lottery year after year. But, by stipulation, this happens just once, so what good are averages? There seems no good way to forge the connection between the value of an act (like keeping the ticket) and the EMV of that act.

We might want to answer in terms of averages, i.e., what would happen to many people playing this same lottery or to you if you played this lottery year after year. But, by stipulation, this happens just once, so what good are averages? There seems no good way to forge the connection between the value of an act (like keeping the ticket) and the EMV of that act.

Now note that expected utility can, in principle, face the same problem. In one-off choices like picking a college, getting married, etc., we can't rely on the law of large numbers. So, what's the relationship between the utility of an act and it's expected utility?

Von Neumann and Morgenstern give us an argument that they are equal if an agent's preferences meet certain structural constraints. That is, they give us an *indirect axiomatic argument* that expected utility is equal to utility. If this is right, then maximizing expected utility is just maximizing utility, and that just means doing what you want to do most.

Von Neumann and Morgenstern's **big idea**: We can measure the strength of a person's preference for a thing on the basis of the risks she is willing to take to receive it.

You're considering where to go on a trip. We know that you prefer Boston to Savannah to Williamsburg. We don't know how much you prefer Boston to Savannah, but, according to Von Neumann and Morganstern, we can figure that out by asking you the following question:

Suppose you were offered a choice between a trip to Savannah and a lottery that will give you a trip to Boston if you "win" and a trip to Williamsburg if you "lose." How great would the chance of winning have to be in order for you to be indifferent between these two choices?

If you prefer Savannah quite a bit more than Williamsburg, you'll demand a fairly high chance at Boston in order to give up the guaranteed trip to Savannah. If you only slightly prefer Savannah to Williamsburg, then you'll accept a much lower chance.

Let's say you say that you need a 75% chance to be indifferent. Then, according to Von Neumann and Morganstern, we should conclude that the trip to Savannah occurs $\frac{3}{4}$ of the way between Williamsburg and Boston on your scale.

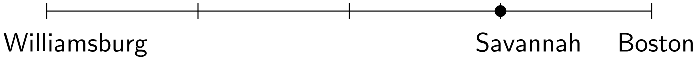

Another way to think about this is that you are ranking not only the basic prizes but also a lottery involving the best and worst trips. You must be indifferent between the lottery and the middle trip.

Let $L(p, x, y)$ represent a lottery that gives you a chance equal to p at a prize x and a chance equal to $1 - p$ at a prize y. Then we can represent your preferences as:

Boston
Savannah, $L(\frac{3}{4}, \textit{Boston}, \textit{Williamsburg})$
Williamsburg

If we assign a utility of 0 to Williamsburg and 1 to Boston, then on this utility scale, Savannah must be assigned $\frac{3}{4}$.

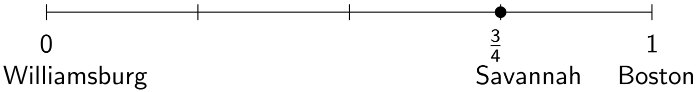

And now we seem to have forged a link between utility and expected utility, for notice that the utility of the lottery is equal to its expected utility.

Indeed, Von Neumann and Morganstern prove that this result holds generally for preferences that meet the structural constraints they set out.

Let's turn to a more rigorous presentation.

Von Neumann and Morgenstern:

If an agent's preferences over outcomes in $\mathcal{O}$ and *lotteries* involving these outcomes satisfy certain structural axioms, then these preferences can be exhibited by an interval utility function that is invariant up to positive linear transformations and where the utility of a lottery is its expected utility.

Let $L(p, x, y)$ designate the *lottery* that gives x with probability p and gives y with probability $1 - p$.

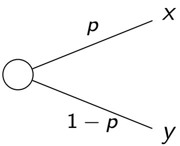

Def 1.3.5. The set of *lotteries* $\mathfrak{L}$ is built up as follows:

- $o \in \mathfrak{L}$ for each outcome $o \in \mathcal{O}$.
- If $L_1 \in \mathfrak{L}$ and $L_2 \in \mathfrak{L}$, then $L(p, L_1, L_2) \in \mathfrak{L}$ for any $p \in \mathbb{R}[0, 1]$.
- Nothing else is in $\mathfrak{L}$.

The Ordering Axioms:

Ax1 (Completeness). $L_1 \succ L_2 \vee L_1 \sim L_2 \vee L_2 \succ L_1$.

Ax2 (Asymmetry). If $L_1 \succ L_2$, then $L_2 \not\succ L_1$.

Ax3 (Negative Transitivity). If $L_1 \not\succ L_2$ and $L_2 \not\succ L_3$, then $L_1 \not\succ L_3$.

These axioms ensure that the agent's preferences can be represented by an ordinal utility function.

Ax4 (Continuity). If $L_1 \succ L_2$ and $L_2 \succ L_3$, then there is $p \in \mathbb{R}[0,1]$ such that $L_2 \sim L(p, L_1, L_3)$.

Ax5 (Better Prizes). $L_1 \succ L_2$ if and only if $L(p, L_1, L_3) \succ L(p, L_2, L_3)$, and $L_1 \succ L_2$ if and only if $L(p, L_3, L_1) \succ L(p, L_3, L_2)$.

Ax6 (Better Chances). If $L_1 \succ L_2$, then $p > q$ if and only if $L(p, L_1, L_2) \succ L(q, L_1, L_2)$.

Ax7 (Reduction of Compound Lotteries).
$L(p, L(q, L_1, L_2), L(r, L_1, L_2)) \sim L(pq + (1-p)r, L_1, L_2)$.

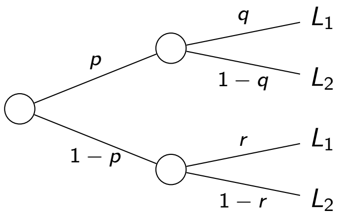

Ax4 (Continuity). If $L_1 \succ L_2$ and $L_2 \succ L_3$, then there is $p \in \mathbb{R}[0, 1]$ such that $L_2 \sim L(p, L_1, L_3)$.

Ax5 (Better Prizes). $L_1 \succ L_2$ if and only if $L(p, L_1, L_3) \succ L(p, L_2, L_3)$, and $L_1 \succ L_2$ if and only if $L(p, L_3, L_1) \succ L(p, L_3, L_2)$.

Ax6 (Better Chances). If $L_1 \succ L_2$, then $p > q$ if and only if $L(p, L_1, L_2) \succ L(q, L_1, L_2)$.

Ax7 (Reduction of Compound Lotteries).
$L(p, L(q, L_1, L_2), L(r, L_1, L_2)) \sim L(pq + (1 - p)r, L_1, L_2)$.

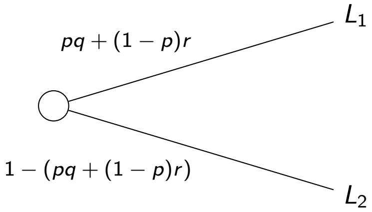

Thm 1.3.7 (Expected Utility Theorem). If an agent's preferences satisfy Ax1-Ax7, then there is a utility function $u : \mathfrak{L} \to \mathbb{R}[0,1]$ such that:

(i) $u(L_1) > u(L_2)$ if and only if $L_1 \succ L_2$.

(ii) $u(L(p, L_1, L_2)) = p \times u(L_1) + (1 - p) \times u(L_2)$.

(iii) Any u' satisfying (i) and (ii) is a positive linear transformation of u.

Partial Proof of Thm 1.3.7. Assume that an agent's preferences satisfy Ax1-Ax7. We construct a utility function u satisfying (i) and (ii).

Let $B \in \mathcal{O}$ designate a *best* outcome—that is, $B \succcurlyeq o$ for all $o \in \mathcal{O}$.

Let $W \in \mathcal{O}$ designate a *worst* outcome—that is, $o \succcurlyeq W$ for all $o \in \mathcal{O}$.

Assume that $B \succ W$.

Let $u(B) = 1$ and $u(L) = 1$ for any lottery $L \in \mathfrak{L}$ such that $B \sim L$.

Let $u(W) = 0$ and $u(L) = 0$ for any lottery $L \in \mathfrak{L}$ such that $W \sim L$.

Now consider any non-extreme $L \in \mathfrak{L}$ such that $B \succ L$ and $L \succ W$.

By Ax4, $L \sim L(p, B, W)$ for some $p \in \mathbb{R}[0, 1]$ (moreover, p is unique).

Let $u(L) = p$.

The utility function $u : \mathfrak{L} \to \mathbb{R}[0,1]$ constructed in this way satisfies (i).

By Ax6, $u(L_1) > u(L_2)$ if and only if $L(u(L_1), B, W) \succ L(u(L_2), B, W)$.

$L_1 \sim L(u(L_1), B, W)$.

$L_2 \sim L(u(L_2), B, W)$.

By Ordering Axioms, if $L_1 \sim L_3$ and $L_2 \sim L_4$, then $L_3 \succ L_4$ if and only if $L_1 \succ L_2$.

Thus, $u(L_1) > u(L_2)$ if and only if $L_1 \succ L_2$.

The utility function $u : \mathfrak{L} \to \mathbb{R}[0,1]$ also satisfies (ii).

Using the Ordering Axioms and Ax5, it is fairly straightforward to show the following:

Lem 1.3.1 (Substitution of Lotteries). If $L_1 \sim L(p, L_2, L_3)$, then both $L(q, L_1, L_4) \sim L(q, L(p, L_2, L_3), L_4)$ and $L(q, L_4, L_1) \sim L(q, L_4, L(p, L_2, L_3))$.

$L(p, L_1, L_2) \sim L(u(L(p, L_1, L_2)), B, W)$.

$L_1 \sim L(u(L_1), B, W)$

$L_2 \sim L(u(L_2), B, W)$

By Substitution of Lotteries, $L(p, L_1, L_2) \sim L(p, L(u(L_1), B, W), L(u(L_2), B, W))$.

By Ax7, $L(p, L(u(L_1), B, W), L(u(L_2), B, W)) \sim L(p \times u(L_1) + (1-p) \times u(L_2), B, W)$.

By Ordering Axioms, $L(u(L(p, L_1, L_2)), B, W) \sim L(p \times u(L_1) + (1-p) \times u(L_2), B, W)$.

By Better Chances, $u(L(p, L_1, L_2)) = p \times u(L_1) + (1-p) \times u(L_2)$.