Rational Choice

Von Neumann-Morgenstern Utility Theory

David Emmanuel Gray

Carnegie Mellon University in Qatar

The Decision Matrix

Acts (A) \ States of Affairs (Ω)	ω_1	ω_2	$\dots$	ω_j	$\dots$	ω_n
a_1	$o_{1,1}$	$o_{1,2}$		$o_{1,j}$		$o_{1,n}$
a_2	$o_{2,1}$	$o_{2,2}$		$o_{2,j}$		$o_{2,n}$
$\dots$						
a_i	$o_{i,1}$	$o_{i,2}$		$o_{i,j}$		$o_{i,n}$
$\dots$						
a_m	$o_{m,1}$	$o_{m,2}$		$o_{m,j}$		$o_{m,n}$

Choice Under Risk

In choice under ignorance, the following all hold:

1. There are different outcomes for different states of affairs relevant to the decision,

2. For each combination of action and state of affairs, you know the outcome, and

3. You *do* know how probable each state of affairs is. Let $P = \{p_1, p_2, \ldots, p_j\}$, where $\mathrm{P}(\omega_j) = p_j$ represents the probability that state ω_j occurs.

The Challenge of Rational Choice

How can a ranking of the *outcomes* be used to generate a ranking of the *acts*?

In choice under risk, the most common answer is to rank the acts based their expected utility:

$$v(a_i) = \sum^{n}_{j=1}[p_j \times u(o_{i,j})].$$

The question, of course, is why do it this way. The Von Neumann-Morgenstern theory of cardinal utility provides one explanation.

Background to the Theory

A lottery L is a probability distribution over a finite set of rewards denoted by $R = \{r_1, r_2, \ldots, r_n\}$. In other word, a lottery L is a sequence $\langle p_1, p_2, \ldots, p_n \rangle$, where $p_j \geq 0$ (for $j = 1, 2, \ldots, n$) and $\sum^n_{j=1}[p_j] = 1$. The quantity p_j is just the chance or probability of winning reward r_j.

Background to the Theory

Rewards (R)

Lotteries ($\mathcal{L}$)		r_1	r_2	$\ldots$	r_j	$\ldots$	r_n
	L_1	$p_{1,1}$	$p_{1,2}$		$p_{1,j}$		$p_{1,n}$
	L_2	$p_{2,1}$	$p_{2,2}$		$p_{2,j}$		$p_{2,n}$
	$\ldots$						
	L_i	$p_{i,1}$	$p_{i,2}$		$p_{i,j}$		$p_{i,n}$
	$\ldots$						
	L_m	$p_{m,1}$	$p_{m,2}$		$p_{m,j}$		$p_{m,n}$

Note: This is *not* equivalent to the standard decision matrix for choice under risk!

Background to the Theory

Von Neumann-Morgenstern utility theory introduces one operation for the combining two lotteries into a third lottery: convex combination.

The **convex combination** of two lotteries is denoted by $\oplus$. Fix quantity x, where $0 \leq x \leq 1$. Then the x-convex combination of L_1 and L_2 creates L_3, where:

$L_3 = xL_1 \oplus (1 - x)L_2$, where the p-values for L_3 are

$p_{3,j} = (x \times p_{1,j}) + [(1 - x) \times p_{2,j}]$ (for $j = 1, 2, \ldots, n$).

Example

Suppose one urn has 30 red balls in it. L_1 says that if I pull a red ball out of it, you get QR 100.

Another urn has 10 blue balls in it. L_2 says that if I pull a blue ball of it, you get QR 0.

So there are two prizes and two lotteries over them:

r_1 = QR 100, and $L_1 = \langle 1.00, 0.00 \rangle$, and

r_2 = QR 0. $L_2 = \langle 0.00, 1.00 \rangle$.

Example

I can combine these two lotteries by putting both sets of balls into the same urn, and then make the same deals where a red ball wins QR 100 and a blue ball wins QR 0. In this case, we have a new lottery:

$$L_3 = (0.75)L_1 \oplus (0.25)L_2 = \langle 0.75, 0.25 \rangle.$$

Example

Lotteries ($\mathcal{L}$)	Rewards (R): QR 100	Rewards (R): QR 0
L_1	1.00	0.00
L_2	0.00	1.00
L_3	0.75	0.25

Background to the Theory

You may think of lottery L_3, the result of a convex combination of lotteries L_1 and L_2, as involving a compound chance where, first, a coin (biased x in favor of landing heads) is flipped. If that coin lands heads then lottery L_1 is run, but if it lands tails then lottery L_2 is run.

The Axioms

Von Neumann-Morgenstern utility theory then places three axioms that judgments ($\succ$) over lotteries ought to satisfy:

Axiom 1: Ordering.

Axiom 2: Independence.

Axiom 3: Archimedean or Continuity Condition.

Axiom 1

Ordering: > is a preference relation over lotteries.

This requires that judgments over the lotteries ought to be complete and transitive.

Axiom 2

Independence: Given that $x > 0$, the following holds:

$$L_1 \succ L_2 \text{ if and only if } xL_1 \oplus (1-x)L_3 \succ xL_2 \oplus (1-x)L_3.$$

Informally, this says that taking the convex combination $\oplus$ with a common lottery L_3 does not affect preference concerning L_1 and L_2.

Axiom 2

The motivation behind independence is that the only difference between the convex combinations (from the judgment after the "if and only if" part) is L_1 and L_2, whereas L_3 is the same with the same weight x given to it. As a result, the judgment concerning L_1 and L_2 should really decide the issue.

The idea is that you can safely ignore spots where there are no differences between two lotteries, and instead focus on where they differ.

Example

Consider the following two lotteries:

Lotteries (L)	Rewards (R): QR 100	QR 50	QR 0
A	0.45	0.25	0.35
B	0.60	0.25	0.15

Example

Independence says you can essentially ignore where the two lotteries are the same:

	Rewards (R)		
Lotteries (L)	QR 100	[redacted]	QR 0
A	0.45	[redacted]	0.35
B	0.60	[redacted]	0.15

Example

This can be seen by taking apart A and B by removing their common reward. Put that common reward into its own lottery (L_3). The remainder of A then becomes L_1 while the remainder of B becomes L_2. See the next slide for the table. Notice that A is the 0.75-convex combination of L_1 and L_3, whereas B is the 0.75-convex combination of L_2 and L_3.

Example

Lotteries (L)	QR 100	QR 50	QR 0
A	0.45	0.25	0.35
B	0.60	0.25	0.15
L_1	0.60	0.00	0.40
L_2	0.80	0.00	0.20
L_3	0.00	1.00	0.00

Rewards (R)

Example

Since $A = (0.75)L_1 \oplus (0.25)L_3$, and $B = (0.75)L_2 \oplus (0.25)L_3$, independence says that a judgment between A and B should reduce to a judgment concerning L_1 and L_2.

Axiom 2

Put slightly differently, given each convex combination, *either* the first part happens, causing L_1 or L_2 to occur, *or* the second part happens, causing L_3 to occur no matter what.

Now if the first part happens, then the judgment over L_1 and L_2 says which is better. But if the second part happens, L_3 occurs no matter what, so the result is indifference. Consequently, a strict preference over L_1 and L_2 should settle the issue.

Axiom 3

The Archimedean or Continuity Condition:

If $L_1 \succ L_2 \succ L_3$, then there exists some x and y, where $0 < x, y < 1$, such that the following holds:

$$xL_1 \oplus (1 - x)L_3 \succ L_2 \succ yL_1 \oplus (1 - y)L_3.$$

This is a technical condition to allow the use of real numbers to provide magnitudes for cardinal utilities.

Axiom 3

There are two important implications of this axiom:

1. There is no lottery L_1 *so* good that combining even a tiny chance of getting it with a worse lottery L_3 that will cause this combination to be better than L_2 (a lottery worse than L_1 but better than L_3).

2. There is no lottery L_3 *so* bad that combining even a tiny chance of getting it with a better lottery L_1 that will cause this combination to be worse than L_2 (a lottery better than L_3 but worse than L_1).

Example

Lotteries ($\mathcal{L}$)	Rewards (R): QR 3,000	QR 30	Death
L_1	1.00	0.00	0.00
L_2	0.00	1.00	0.00
L_3	0.00	0.00	1.00

Example

Obviously $L_1 \succ L_2 \succ L_3$. As a result, the Archimedean or continuity condition says that there must be some x (probably really close to 1) such that you judge

$$xL_1 \oplus (1 - x)L_3 > L_2.$$

In other words, you should strictly prefer a small risk of death in getting QR 3,000 over QR 30 for sure.

The Theorem

The Von Neumann-Morgenstern Theorem: $\succ$ satisfies axioms 1, 2, and 3 over lotteries if and only if there exists some cardinal utility function u on rewards such that following holds:

$$L_1 \succcurlyeq L_2 \text{ if and only if } v(r_1) \geq v(r_2),$$

$$\text{where } v(L) = \Sigma^n_{j=1}[p_j \times u(r_j)].$$

Furthermore: utility function u is an *interval* scale. That is, u is unique under positive affine transformations: any utility function u'—where $u'(x) = [\alpha \times u(x)] + \beta$ for any $\alpha > 0$ and any β—generates the same judgments over lotteries as utility function u.

Implications of the Theorem

There are two important implications of the theorem:

1. It provides a representation theorem for choice under risk (since choice under risk essentially involves choice between different lotteries), and

2. It provides a method for constructing an interval utility function over outcomes.

Representation Theorem

Lemma (for Choice Under Risk): $\succ$ satisfies axioms 1, 2, and 3 over acts* if and only if there exists some cardinal utility function u on outcomes such that following holds:

$a_1 \succcurlyeq a_2$ if and only if $v(a_1) \geq v(a_2)$,

where $v(a_i) = \sum_{j=1}^{n} [p_j \times u(o_{i,j})]$.

Furthermore: utility function u is an *interval* scale. That is, u is unique under positive affine transformations: any utility function u'—where $u'(x) = [\alpha \times u(x)] + \beta$ for any $\alpha > 0$ and any β—generates the same judgments over acts as utility function u.

*For axioms 2 and 3 to apply, ⊕ must be redefined over acts. I leave the details as an exercise.

Constructing Interval Utilities

This can also be used to create an interval utility function u over outcomes. Begin this by identifying the best and worst possible outcomes:

o^* = the best possible outcome, and
o_* = the worst possible outcome.

Then assign utility values for each of these:

$$u(o^*) = 1.00, \text{ and}$$
$$u(o_*) = 0.00.$$

Constructing Interval Utilities

Now set up lotteries for each of these outcome:

$L^* = \text{get } o^*$ for sure, and
$L_* = \text{get } o_*$ for sure.

Now for each outcome o_i, I ask, what point are you indifferent between the following:

1. o_i for sure, or
2. $xL^* \oplus (1 - x) L_*$.

Finally, assign $u(o_i) = x$.

Example

Now it may be possible to construct an interval-valued utility function for the outcomes of Pascal's wager. Recall that the following judgments hold for these outcomes:

Heaven > Benefits of Atheism > Burdens of Belief > Hell.

u(Heaven) =

u(Benefits of Atheism) =

u(Burdens of Belief) =

u(Hell) =

Next Class…

We will discuss what it means to maximize expected utility along with more arguments addressing why it is rational to do so in choice under risk.