The Allais Paradox and Risk-aversion

Ryan Doody

Risk-aversion

The orthodox theory of rational decision-making is *expected utility theory*, according to which there are two important components of decision-making:

1. *Utilities*. How much do you value the various outcomes that might obtain?
2. *Probabilities*. How likely do you think a given act is to realize these outcomes?

Let $h = \{x_1, E_1; x_2, E_2; \ldots x_n, E_n\}$ be a gamble that yields, for each $1 \leq i \leq n$, an outcome x_i if event E_i obtains, and is such that $u(x_1) \leq u(x_2) \leq \cdots \leq u(x_n)$.

Expected Utility

$$EU(h) = \sum_{i=1}^{n} c(E_i) \cdot u(x_i)$$

The value of an act is its *expected utility*, and a rational decision-maker will prefer the act with the highest expected utility.

Many people's preference display risk-aversion in the following sense. Consider a choice between,

(A) \$50 for sure,

(B) A fair coin-flip between \$0 and \$100,

Many people prefer (A) to (B). If such a person is an EU maximizer, then $u(\$50) - u(\$0) > u(\$100) - u(\$50)$.

Objection: This conflates two different reasons why one might prefer (A) to (B); local considerations about how valuable one takes one's outcomes to be vs global considerations about how the gamble's outcomes are structured.

The Allais Paradox

This problem comes from the French economist Maurice Allais, who raised it as a counterexample to Leonard Savage's *Sure-Thing Principle* (which is one of the central axioms underlying Expected Utility Theory).

Roughly, the principle says: if two gambles agree on what happens if one event obtains ($\neg E$), then your ranking of them should depend only on how you rank what would happen if this event doesn't obtain (E).

Sure-Thing Principle

	E	$\neg E$
f	X	Z
g	Y	Z
f^*	X	Z^*
g^*	Y	Z^*

$f \succ g$ if and only if $f^* \succ g^*$

Consider the following two lotteries:

(1_A) a 11% chance of winning \$1,000,000.

(1_B) a 10% chance of winning \$5,000,000.

Which do you prefer? __________

Now consider to other lotteries:

(2_A) a 100% chance of winning \$1,000,000.

(2_B) a 10% chance of winning \$5,000,000, and a 89% chance of winning \$1,000,000.

Which do you prefer? __________

The Allais Preferences: $1_B \succ 1_A$, $2_A \succ 2_B$.

Question: Can you assign utilities to $0, $1,000,000, and $5,000,000 so that your ranking of the lotteries obey Expected Utility Theory?

If you have the Allais Preferences, the answer is: *No*.
Is this, then, a counterexample to Expected Utility Theory?

It's easier to see how this example works if we represent it in a table.

The Allais Paradox

	Tickets 1	2–11	12–100
1_A	$1,000,000	$1,000,000	$0
1_B	$0	$5,000,000	$0
2_A	$1,000,000	$1,000,000	$1,000,000
2_B	$0	$5,000,000	$1,000,000

There is no way to assign utility values to $0, $1,000,000, and $5,000,000 so that 1_B has higher expected utility than 1_A, and that 2_A has higher expected utility than 2_B.

Therefore, these preferences cannot be represented as maximizing expected utility.

Is that right? Is there no way to represent these preferences using EUT? If so, are these preferences irrational?

Risk-weighted Expected Utility Theory

According to *Risk-weighted Expected Utility Theory* (REUT), there are three components of rational decision-making:

1. *Utilities.* How much do you value the various outcomes that might obtain?
2. *Probabilities.* How likely do you think a given act is to realize these outcomes?
3. *Risk-function.* To what extent are you generally willing to accept the risk of something worse in exchange for the possibility of something better.

REUT is a *generalization* of EUT: the two views coincide when $r(p) = p$, for all probabilities p.
The risk function is subject to the following constraints: for all p, $0 \leq r(p) \leq 1$; $r(0) = 0$ and $r(1) = 1$; r is non-decreasing.
So, EUT can be understood as a special case of REUT, which encodes a particular attitude toward risks: it is *risk-neutral*.

Let $h = \{x_1, E_1; x_2, E_2; \ldots x_n, E_n\}$ be a gamble that yields, for each $1 \leq i \leq n$, an outcome x_i if event E_i obtains, and is such that $u(x_1) \leq u(x_2) \leq \cdots \leq u(x_n)$.

Risk-Weighted Expected Utility

$$REU(h) = u(x_1) + r\left(\sum_{i=2}^{n} c(E_i)\right) \cdot (u(x_2) - u(x_1)) + \cdots + r\left(c(E_n)\right) \cdot (u(x_n) - u(x_{n-1}))$$

Example: Consider the choice between (A) and (B), and let's assume that you value money linearly. And suppose that $r(p) = p^2$.

$$\begin{aligned} REU(A) &= 50 \\ REU(B) &= 0 + r(1/2) \cdot (100 - 0) = (1/2)^2 \cdot (100) \\ &= (1/4) \cdot (100) = 25 \end{aligned}$$

Expected Utility. We can rewrite the EU of a gamble, $p \cdot u(x_2) + (1-p) \cdot u(x_1)$, as follows (where x_1 is worse than x_2):

$$u(x_1) + p \cdot (u(x_2) - u(x_1))$$

That's the minimum value of the gamble ($u(x_1)$) plus the amount you might gain above that minimum ($u(x_2) - u(x_1)$) weighted by the probability of realizing that gain (p).

Risk-weighted Expected Utility. Instead of weighting the potential gains by their probabilities, p, REUT weights these potential gains by *a function* of their probabilities, $r(p)$.