Decision Theory

1.5 Paradoxes and Problems

George Mason University, Spring 2018

Here's the question we've been thinking about:
When making a decision under risk, *why* maximize expected utility?

Here's an answer we considered (call it the *law of large numbers* reply):

Answer. In the long run, you will be better off by maximizing EU.

Reply. No real-life decision maker will ever face a decision an infinite number of times. Keynes: "In the long run we are all dead."

Reply. Many decisions are unique, such as the decision to marry a particular partner, the decision to start a particular war, and so on.

So we turned to an *indirect axiomatic argument* to try to show that we should maximize expected utility.

- The aim was to show that for any action $u(a) = EU(a)$.
- We did this by adopting Ax1-Ax7 and proving the Expected Utility Theorem from them.

Recall the Expected Utility Theorem:

Thm 1.3.7. If an agent's preferences over lotteries 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.

Notice that each act $a \in \mathcal{A}$ is itself a lottery whose expected utility is its utility. We forged a connection between utility and expected utility. Resnik: "In choosing an act whose expected utility is maximal an agent is simply doing what he wants to do!"

Now we face some new problems:

Reply. It now seems that an agent whose preferences satisfy Ax1-Ax7 doesn't even need decision theory.

Counter-reply. The maxim 'Maximize EU!' should be understood as 'Have preferences that satisfy the structural constraints Ax1-Ax7 (and then just do what you most prefer to do)!'

This is a common way of understanding what decision theory is really about. It tells us how we ought to structure our preferences if we are to be rational, but once we've done that, it has nothing else to tell us. In that case, we'll just do the action that has the highest EU because it has the highest utility for us. The model, then, is just a way of bringing out how a rational agent would behave. It doesn't itself help us make decisions.

Even on this understanding, there are some sophisticated problems we need to confront. We're going to look at some paradoxes in which expected utility seems to lead us astray and at some responses to those paradoxes.

Ex. Allais Paradox.

You are given a choice between the following two payouts:

A: \$1M.

B: You receive \$5M with 10% probability, \$1M with 89% probability, and nothing with 1% probability.

Do you choose A or B?

Give it some thought and write down your choice.

Next you are next given a choice between the following two payouts:

C: You receive \$5M with 10% probability and nothing with 90% probability.

D: You receive \$1M with 11% probability and nothing with 89% probability.

Do you choose C or D?

Give it some thought and write down your choice.

Now consider this way of rewriting the choices as lotteries:

Do you choose A or B:

A: $L(0.11, \$1M, \$1M)$

B: $L(0.11, L(\frac{10}{11}, \$5M, \$0), \$1M)$

Do you choose C or D:

C: $L(0.11, L(\frac{10}{11}, \$5M, \$0), \$0)$

D: $L(0.11, \$1M, \$0)$

If you chose:

A: \$1M or $L(0.11, \$1M, \$1M)$
and
C: You receive \$5M with 10% probability and nothing with 90% probability or $L(0.11, L(\frac{10}{11}, \$5M, \$0), \$0)$

or

B: You receive \$5M with 10% probability, \$1M with 89% probability, and nothing with 1% probability or $L(0.11, L(\frac{10}{11}, \$5M, \$0), \$1M)$
and
D: You receive \$1M with 11% probability and nothing with 89% probability or $L(0.11, \$1M, \$0)$

then your preferences cannot be represented by a utility function because you violate Better Prizes. If you think $u(A) > u(B)$ then, by Better Prizes, you should think $u(D) > u(C)$. The problem is, Allais showed that most people don't think this, and most people think their preferences are reasonable even on reflection.

Let's look at it another way:

$EU(A) = 1 \times u(\$1\text{M})$.

$EU(B) = 0.1 \times u(\$5\text{M}) + 0.89 \times u(\$1\text{M}) + 0.01 \times u(\$0\text{M})$.

$EU(C) = 0.1 \times u(\$5\text{M}) + 0.9 \times u(\$0\text{M})$.

$EU(D) = 0.11 \times u(\$1\text{M}) + 0.89 \times u(\$0\text{M})$.

$EU(A) - EU(B) = 0.11 \times u(\$1\text{M}) - 0.1 \times u(\$5\text{M}) - 0.01 \times u(\$0\text{M})$.

$EU(D) - EU(C) = 0.11 \times u(\$1\text{M}) - 0.1 \times u(\$5\text{M}) - 0.01 \times u(\$0\text{M})$.

If you choose A over B, then presumably
$EU(A) - EU(B) = EU(D) - EU(C) > 0$.

So if you are an EU-maximizer, then you choose D over C.

The Allais Paradox plays on the common preference for a good for certain over a risky chance for a more valuable good.

Reply. The outcomes shouldn't be specified solely in terms of money. The outcome in offer B is \$0 *plus serious disappointment*.

Counter-reply. Every objection to the Principle of Maximizing EU might be thwarted by fiddling with the outcomes.

Reply. Bite the bullet. Savage: An agent who chooses A and C is irrational because they violate the *sure-thing principle*.

Imagine a raffle with a hundred tickets:

	Ticket 1	Tickets 2-11	Tickets 12-100
gamble A	\$1M	\$1M	\$1M
gamble B	\$0M	\$5M	\$1M
gamble C	\$0M	\$5M	\$0M
gamble D	\$1M	\$1M	\$0M

The third column should be ignored when deciding between the gambles.

Counter-reply. Why satisfy the sure-thing principle?

Ex. Ellsberg Paradox.

An urn contains 90 balls. You know that 30 of these are yellow. You also know that the remaining 60 balls are either red or blue, but you do not know the proportion. I am about to draw a ball from the urn and I give you a choice between the following two payouts:

A: You receive \$100 if a yellow ball is drawn and \$0 otherwise.

B: You receive \$100 if a red ball is drawn and \$0 otherwise.

Do you choose A or B?

Suppose that I had instead offered a choice between the following two payouts:

C: You receive \$100 if either a red or blue ball is drawn and \$0 otherwise.

D: You receive \$100 if either a yellow or blue ball is drawn and \$0 otherwise.

Do you choose C or D?

Suppose that you assign a conditional probability of p to getting a red ball given that you get a red or blue ball.

$EU(A) = \frac{1}{3} \times u(\$100) + \frac{2}{3} \times u(\$0).$

$EU(B) = \frac{1}{3} \times u(\$0) + \frac{2}{3} \times p \times u(\$100) + \frac{2}{3} \times (1-p) \times u(\$0).$

$EU(C) = \frac{1}{3} \times u(\$0) + \frac{2}{3} \times u(\$100).$

$EU(D) = \frac{1}{3} \times u(\$100) + \frac{2}{3} \times p \times u(\$0) + \frac{2}{3} \times (1-p) \times u(\$100).$

$EU(A) - EU(B) = \frac{1-2p}{3} \times u(\$100) + \frac{2p-1}{3} \times u(\$0).$

$EU(D) - EU(C) = \frac{1-2p}{3} \times u(\$100) + \frac{2p-1}{3} \times u(\$0).$

If you choose A over B, then presumably
$EU(A) - EU(B) = EU(D) - EU(C) > 0.$

So if you are an EU-maximizer, then you choose D over C.

The Ellsberg Paradox plays on the common preference for known risks over unknown risks.

Reply. Bite the bullet. An agent who chooses A and C is irrational because they violate the sure-thing principle (and Better Chances).

	$[\frac{1}{3}]$ Yellow	$[\frac{2}{3} \times p]$ Red	$[\frac{2}{3} \times (1-p)]$ Blue
gamble A	\$100	\$0	\$0
gamble B	\$0	\$100	\$0
gamble C	\$0	\$100	\$100
gamble D	\$100	\$0	\$100

Counter-reply. Why satisfy the sure-thing principle?

Reply. This is a decision under ignorance so the Principle of Maximizing EU does not apply.

Counter-reply. This reply is not open to subjectivists about probability who think that there are no real decisions under ignorance.

Ex. St. Petersburg Paradox.

You are given a choice between the following two payouts:

A: $100.

B: A fair coin is flipped until it lands tails. If the coin lands tails on the first toss, then you receive $2. If the coin lands tails on the second toss, then you receive $4. In general, if the coin lands tails on the nth toss, then you receive $$2^n$.

Do you choose A or B?

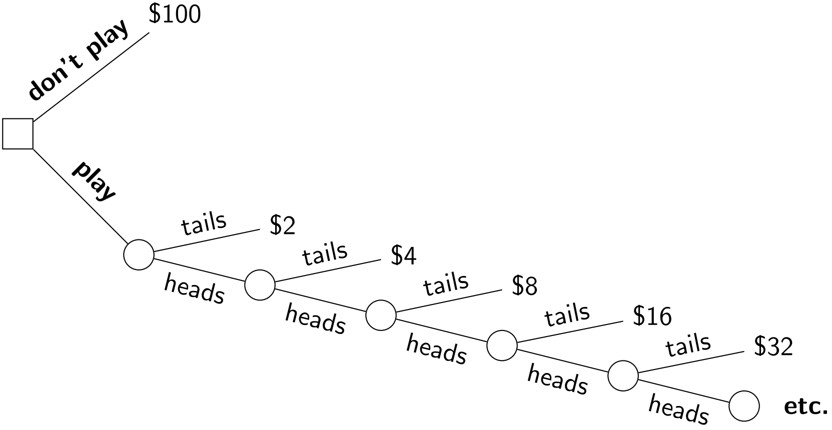

Let $EMV(a)$ designate the *expected monetary value* of $a \in \mathcal{A}$.

$EMV(A) = \$100$.

$EMV(B) = \frac{1}{2} \times \$2 + \frac{1}{4} \times \$4 + \frac{1}{8} \times \$8 + ... = \$1 + \$1 + \$1 + ... = \$\infty$.

Reply. There are diminishing returns to money.

Counter-reply. The St. Petersburg Paradox can be reframed in terms of utilities. If the coin lands tails on the first toss, then you receive a prize worth 2 utiles, and so forth.

Reply. There is, or should be, an upper bound on utility.

Counter-reply. This upper bound is ad hoc.

Reply. Jeffrey: "Anyone who offers to let the agent play the St. Petersburg game is a liar, for he is pretending to have an indefinitely large bank."

Counter-reply. All sorts of hypothetical prizes can be allowed.

Ex. Moscow Game.

You are given a choice between the following two payouts:

A: A fair coin is flipped until it lands tails. If the coin lands tails on the first toss, then you receive \$2. If the coin lands tails on the second toss, then you receive \$4. In general, if the coin lands tails on the nth toss, then you receive $\$2^n$.

B: A biased coin that lands tails with probability 0.4 is flipped until it lands tails. If the coin lands tails on the first toss, then you receive \$2. If the coin lands tails on the second toss, then you receive \$4. In general, if the coin lands tails on the nth toss, then you receive $\$2^n$.

Do you choose A or B?

Intuitively, B is preferable.

However, $EMV(A) = EMV(B) = \infty$.

Ex. Leningrad Game.

You are given a choice between the following two payouts:

A: A fair coin is flipped until it lands tails. If the coin lands tails on the first toss, then you receive \$2. If the coin lands tails on the second toss, then you receive \$4. In general, if the coin lands tails on the nth toss, then you receive $\$2^n$.

B: A biased coin that lands tails with probability 0.4 is flipped until it lands tails. If the coin lands tails on the first toss, then you receive \$2. If the coin lands tails on the second toss, then you receive \$4. In general, if the coin lands tails on the nth toss, then you receive $\$2^n$. However, if the coin lands tails on the third toss, then you receive \$8 *and* get to play the St. Petersburg Game (which is just the game in A).

One solution to the Moscow Paradox is to argue that the appropriate comparison is not EU but Relative EU, i.e., compare the EMV at each toss. In this case, it seems clear that it's better to play A than B in the Moscow game. However, B in Leningrad hijacks this solution.

Ex. Two Envelope Paradox.

You are offered a choice between two envelopes A and B. You know that one of these envelopes contains twice as much money as the other, but you do not know how much money is in either envelope. You pick A. But right before you open this envelope, you are offered the opportunity to switch to B. Do you switch?

If you are an EU-Maximizer, then it seems that you should switch.

Suppose that there is $\$S$ in envelope A.

$EMV(\text{A}) = \$S.$

$EMV(\text{B}) = \frac{1}{2} \times \$\frac{1}{2}S + \frac{1}{2} \times \$2S = \$\frac{5}{4}S.$

$EMV(\text{B}) > EMV(\text{A}).$

But now suppose I offer you the opportunity to switch back...

Reply. The Two Envelope Paradox requires that there is an infinite amount of money in the world. If there is only $\$T$ available, then the envelope with more money can contain no more than $\$\frac{2}{3}T$. If this envelope contains $\$\frac{2}{3}T$, then the other envelope must contain $\$\frac{1}{3}T$.

Counter-reply. We can work with utilities instead.

Ex. Newcomb's Paradox.

You are standing in front of a table on top of which are two boxes A and B. Box A is transparent and you can see that it contains \$1000. Box B is opaque and contains either \$1M or nothing. You are offered the choice between taking only box B or both of these boxes. But before selecting, you are also given one extra piece of information. At some point in the past, a prophetic being called *The Predictor* predicted what you will do. If The Predictor predicted that you will take only box B, then \$1M was placed inside this box. If The Predictor predicted that you will take both boxes, then box B was left empty. The Predictor is almost always right. So what do you select?