Introduction to Decision Theory & Expected Value

October 24, 2017

Normative Decision Theory

Decision Theory comes in three different flavors: *Normative* Decision Theory, *Descriptive* Decision Theory, and *Interpretive* Decision Theory. We will be mostly concerned with the first of the three — although, as we'll see, it won't always be easy to disentangle the three.

Normative Decision Theory is concerned with how it's rational to act (especially when you are uncertain about the consequences of your actions). More precisely, Normative Decision Theory attempts to develop a mathematically rigorous account of *Instrumental Rationality:* doing what's best given your ends.

Of particular interest to us will be the notion of **Expected Value**:

Let $L = \{\langle p_1, \$x_1\rangle, \langle p_2, \$x_2\rangle, \dots\}$ be a wager that pays out $\$x_1$ with probability p_1, $\$x_2$ with probability p_2, and so on and so forth. The *expected value* of wager L is the weighted average of its potential payoffs, where the weights correspond the probability of the wager paying out that amount.

$$EV(L) = \sum_i p_i \cdot x_i$$
$$= p_1 \cdot x_1 + p_2 \cdot x_2 + \dots$$

The average of $a_1, \dots, a_n$ is

$$\frac{a_i + \dots a_n}{n} = \sum_{i=1}^{n} (\frac{1}{n}) \cdot a_i$$

Here, the weights — $\frac{1}{n}$ — are all the same. We can get a *weighted* average by changing the weights (just so long as they sum to 1).

Claim: You rationally ought to value wagers in accordance with their expected values — i.e., prefer wagers with higher expected values to ones with less; be indifferent between wagers when they have the same expected values.

Is this right? We will come back to that question very soon.

A Brief Interlude: The Probability Axioms

[Normality]

Every proposition (over which *Pr* is defined) is assigned a probability somewhere between 0 and 1.

$$0 \leq \Pr(X) \leq 1 \quad (1)$$

[Certainty]

Any proposition Ω that is certain to be true (e.g., a logical truth) is assigned probability 1.

$$\Pr(\Omega) = 1 \quad (2)$$

[Additivity]

If propositions X and Y are *mutually exclusive*, then the probability of their disjunction is equal to the sum of their probabilities.

Two propositions are *mutually exclusive* just in case they cannot *both* be true.

$$\text{If } X \& Y \text{ are mutually exclusive, } \Pr(X \vee Y) = \Pr(X) + \Pr(Y) \quad (3)$$

The Overlap Rule

What is the probability of a disjunction when its disjuncts are *not* mutually exclusive?

[OVERLAP]

The probability of a disjunction is equal to the sum of the probabilities of its disjuncts minus the probability its disjuncts' overlap.

$$\Pr(X \vee Y) = \Pr(X) + \Pr(Y) - \Pr(X \wedge Y) \tag{4}$$

This is the intuitive idea behind The Overlap Rule: if the propositions X and Y are not mutually exclusive, then by adding $\Pr(X)$ to $\Pr(Y)$ in order to get $\Pr(X \vee Y)$, we are "double counting" the possibility in which they are *both* true, i.e., $(X \wedge Y)$; to correct for this, we need to subtract out $\Pr(X \wedge Y)$.

Conditional Probability

Let $\Pr(X \mid Y)$ be the probability of X *conditional* on Y.
It is defined as follows:

$$\Pr(X \mid Y) = \frac{\Pr(X \wedge Y)}{\Pr(Y)} \tag{5}$$

$\Pr(X|Y)$ is, roughly, the probability that X is the case on the assumption that Y is the case.

More Rules and Definitions

The Multiplication Rule: If $\Pr(E) > 0$, then

$$\Pr(X \wedge E) = \Pr(X \mid E) \cdot \Pr(E) \tag{6}$$

The Total Probability Rule: If $0 < \Pr(E) < 1$, then

$$\Pr(X) = \Pr(X \mid E) \cdot \Pr(E) + \Pr(X \mid \neg E) \cdot \Pr(\neg E) \tag{7}$$

The Logical Consequence Rule: Suppose that Y logically entails X. Then

$$\Pr(Y) \leq \Pr(X) \tag{8}$$

Statistical Independence. X and Y are said to be *statistically independent* just in case $\Pr(X \mid Y) = \Pr(X)$.

If X and Y are statistically independent, $\Pr(X \wedge Y) = \Pr(X) \cdot \Pr(Y)$.

Proof. From **The Multiplication Rule:**

$$\Pr(X \wedge Y) = \Pr(X \mid Y) \cdot \Pr(Y)$$

And, from the definition of **Statistical Independence:**

$$\Pr(X \mid Y) = \Pr(X)$$

So, $\Pr(X \wedge Y) = \Pr(X) \cdot \Pr(Y)$.

Should You Maximize Expected (Monetary) Value?

There are several problems with this proposal.

1. *Overly Restrictive.* We care about more things than money.
2. *The St. Petersburg Paradox.* Consider the following wager: I will flip a fair coin until it comes up heads; if the first time it comes up heads is the n^{th} toss then I will pay you $\$2^n$. What's the most you'd be willing to pay for this wager? What is its expected monetary value?

3. *Context-sensitivity of Valuing Money.* Doesn't the value of a wager depend on more than merely how much it's expected to pay out? Examples: your total fortune, how much you personally care about money, etc.

4. *Phenomenon of Risk-aversion.* Is it irrational to prefer a sure-thing $\$x$ to a wager whose expected payout is $\$x$?

Lesson: We should move away from "monetary payouts" to "utility".

Decision Problems & *Expected Utility Theory*

Let $\mathbf{A_1}, \mathbf{A_2}, \ldots, \mathbf{A_k}$ be mutually exclusive and mutually exhaustive *acts* (sometimes called "options," or "alternatives").

Let $S_1, S_2, \ldots, S_n$ be a mutually exclusive and mutually exhaustive set of *states*.

Every pair of acts and states $\langle \mathbf{A_i}, S_j \rangle$ corresponds to an *outcome*: $O[\mathbf{A_i}, S_j] = (A_i \wedge S_j)$.

Let $\Pr(S_j)$ be the *probability* that the world is in state S_j, and

Let $u(A_i \wedge S_j)$ be the *subjective degree of value* — or, "utility" — that you assign to the outcome that results from performing act $\mathbf{A_i}$ when state S_j obtains.

Decision Matrix

	S_1	S_2	$\ldots$	S_n
$\mathbf{A_1}$	$O[\mathbf{A_1}, S_1]$	$O[\mathbf{A_1}, S_2]$	$\ldots$	$O[\mathbf{A_1}, S_n]$
$\mathbf{A_2}$	$O[\mathbf{A_2}, S_1]$	$O[\mathbf{A_2}, S_2]$	$\ldots$	$O[\mathbf{A_2}, S_n]$
$\vdots$	$\vdots$	$\vdots$	$\vdots$	$\vdots$
$\mathbf{A_k}$	$O[\mathbf{A_k}, S_1]$	$O[\mathbf{A_k}, S_2]$	$\ldots$	$O[\mathbf{A_k}, S_n]$

Define the **Expected Utility** of an action **A** to be:

$$EU(\mathbf{A}) = \sum_i \Pr(S_i) \cdot u(A \wedge S_i)$$

Further Questions: How should we understand Pr? How should we understand u? Where do they come from, and what do they mean?