Decision Theory, Lecture 1

Rachael Briggs

July 28, 2014

A Simple Puzzle

You are making an omelette. You have cracked five eggs into a bowl. You're not sure whether the sixth egg is still good. Should you crack it into the same bowl, and risk spoiling the omelette, or crack it into a different bowl, and create an extra dish to wash?

Three Key Elements

States fully determine the effects of your actions on the things you care about

- up to the world; you have no direct control over them
- mutually exclusive and jointly exhaustive; exactly one obtains
- sets of states are called **events**
- I'll uses $\mathcal{S}$ for the set of all states and s for individual states

Outcomes are bearers of non-instrumental value

- mutually exclusive and jointly exhaustive: you get exactly one
- I'll use $\mathcal{O}$ for the set of all outcomes and o for individual outcomes

Acts are objects of choice

- formally represented as functions from states to outcomes—each act has finitely many possible outcomes
- the set of notionally possible acts is much richer than the set of available acts
- I'll use $\mathcal{A}$ for the set of all acts and a, b, c, d for individual acts

The Puzzle in Matrix Form

(states are columns; acts are rows; outcomes are cells)

	last egg good	**last egg rotten**
same bowl	6 eggs, 1 dish	0 eggs, 1 dish
new bowl	6 eggs, 2 dishes	5 eggs, 2 dishes

We can add probabilities to the states (representing how likely they are) and utilities to the outcomes (representing how good they are).

	last egg good ($P = 0.8$)	**last egg rotten** ($P = 0.2$)
same bowl	6 eggs, 1 dish ($u = 5$)	0 eggs, 1 dish ($u = -1$)
new bowl	6 eggs, 2 dishes ($u = 4$)	5 eggs, 2 dishes ($u = 3$)

Grand Worlds vs. Small Worlds

The version of the problem I gave was a *small world* puzzle: the outcomes are too coarse-grained to represent *everything* you care about.

- If you don't eat an omelette, you care about the availability of other foods.
- Maybe you're worried about running out of dish detergent.
- Are you planning to wash up, or leave the dishes for your roommate? Either way, there are probably further consequences that matter here.

If the outcomes are too coarse-grained, then the states are too coarse-grained.

- What's in the fridge? Which restaurants are open? Is your roommate cranky?...

If the space of states or outcomes changes, then the space of acts changes too. (Remember, acts are functions from states to outcomes).

You could have a *grand-world* version of the same puzzle, in which acts, states, and outcomes are so fine-grained they represent everything you care about. (The original coarser-grained version of the puzzle is a *small-world* version.)

A Slightly Grander World

(if the omelette doesn't work out, maybe order pizza?)

	last egg good pizzaria open ($P = 0.6$)	**last egg rotten pizzaria open** ($P = 0.15$)	**last egg good pizzaria closed** ($P = 0.2$)	**last egg rotten pizzaria closed** ($P = 0.05$)
same bowl order pizza	6 eggs 1 dish ($u = 5$)	pizza 1 dish ($u = 2$)	6 eggs 1 dish ($u = 5$)	no food 1 dish ($u = -2$)
new bowl order pizza	6 eggs 2 dishes ($u = 4$)	5 eggs 2 dishes ($u = 3$)	6 eggs 2 dishes ($u = 4$)	5 eggs 2 dishes ($u = 3$)
same bowl no order	6 eggs 1 dish ($u = 5$)	no food 1 dish ($u = -2$)	6 eggs 1 dish ($u = 5$)	no food 1 dish ($u = -2$)
new bowl no order	6 eggs 2 dishes ($u = 4$)	5 eggs 2 dishes ($u = 3$)	6 eggs 2 dishes ($u = 4$)	5 eggs 2 dishes ($u = 3$)

Dominance Principles

The plan "same bowl, order pizza" is better than the plan "same bowl, no order". I say:

- Act *a* **strictly dominates** act *b* iff

  for every state *s*, *a* yields a better outcome in *s* than *b* does in *s*.

- Act *a* **weakly dominates** act *b* iff

  for every state *s*, *a* yields at least as good an outcome in *s* as *b* does in *s*, and

  for some state *s*, *a* yields a better outcome in *s* than *b* does in *s*.

There are two principles of dominance. (Naming them gets tricky, because the principle that involves the stronger dominance *relation* is a weaker *principle*. Both of these principles often get called "The Dominance Principle"—and there are other principles that go by the same name.)

Principle of Strict Dominance If *a* strictly dominates *b*, you should prefer *a* to *b*.

Principle of Weak Dominance If *a* weakly dominates *b*, you should prefer *a* to *b*.

By the Principle of Weak Dominance, you should prefer "same bowl, order pizza" to "same bowl, no order".

Normal vs. Extensive Form

My grand(er)-world problem involved two sequential decisions: one about what to do first (crack the egg into the same bowl, or a new bowl) and one about what to do next, if the omelette didn't work out (order pizza or not). The above representation is in *normal form*; we could also represent the decision in *extensive form*, like this.

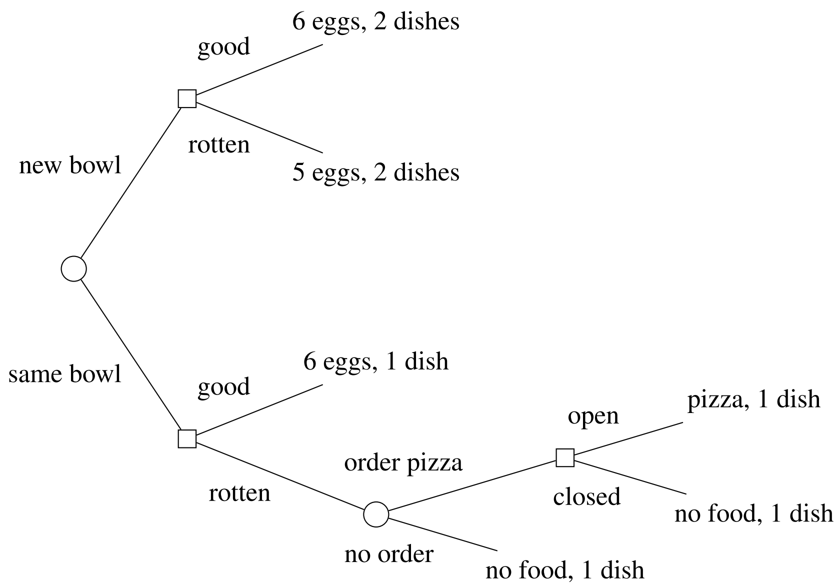

Games

In addition to having individual decisions, where the outcome depends on the actions of one player and of the world, we can have games, where the outcome depends on the actions of more than one player. Two-player games can be represented by matrices, or by trees. Some common two-player games:

The Prisoner's Dilemma

	Cooperate	**Defect**
Cooperate	−1, −1	−10, 0
Defect	0, −10	−5, −5

Centipede

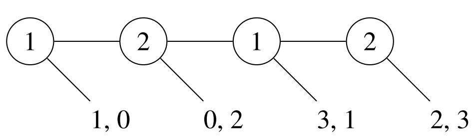

Chicken

	Stay	**Swerve**
Stay	−10, −10	1, −1
Swerve	−1, 1	0, 0

Stag Hunt

	Stag	**Hare**
Stag	3, 3	0, 2
Hare	2, 0	2, 2

Probability and Utility

Back to Savage's puzzle. To figure out what to do in this situation, we need more than the Principle of Weak Dominance. We need:

- A probability function P, which maps a domain $\mathcal{F}$ of events (sets of states) onto closed the interval [0, 1].

  Intuitively, P tells you how likely various possible events are.

- A utility function u, which maps $\mathcal{O}$ onto a set of real numbers.

  Intuitively, u tells you how (non-instrumentally) valuable various outcomes are.

- An expected utility function U, which maps $\mathcal{A}$ onto a set of real numbers.

  Intuitively, U tells you how desirable various acts are, in light of

  - the values of their possible outcomes, and
  - how likely each possible outcome would be, were you to perform the act in question.

  So P and u should suffice to determine U.

More about U:

The expected utility of an act a is gotten by taking a weighted average of the utilities of its possible outcomes. The weighting of an outcome is determined by how likely it is to ensue, if you perform the act in question.

$$U(a) = \sum_{o \in \mathcal{O}} P_a(o)u(o)$$

The greater $U(a)$, the better a is—you should always prefer an act with higher expected utility to an act with lower expected utility.

In the small-world version of the case,

$$U(\textbf{same bowl} = 0.8 \times 5 + 0.2 \times -1 = 3.8$$

$$U(\textbf{new bowl} = 0.8 \times 4 + 0.2 \times 3 = 3.8$$

(so presumably you should be indifferent)

- I'm assuming a *normative* understanding of expected utility theory here: that it's a theory about what you *ought* to prefer. But some authors have a *descriptive* understanding of expected utility theory—they hold that it is an account of how people's actual preferences are structured. Behavioural economics has cast doubt on the descriptive interpretation.

- What if two acts have the same expected utility; does expected utility theory tell you to be indifferent between them? If so, it seems to conflict with the Principle of Weak Dominance. Suppose you are about to throw an infinitely thin dart at a dartboard. You can take one of two options:

	hit centre point ($P = 0$)	**hit other point** ($P = 1$)
Bet 1	\$100	\$100
Bet 2	\$−100	\$100

  These bets have the same expected utility, but Bet 1 weakly dominates Bet 2.

- A purported counterexample to expected utility theory from Bernard Williams: An aunt has one golden retriever puppy, which she can give to one of her two nieces. (Assume that both nieces are in a good position to take care of the puppy.) She is indifferent between giving the puppy to Niece 1 and giving it to Niece 2, but she prefers flipping a fair coin, and giving the puppy to Niece 1 if the coin lands heads and to Niece 2 if the coin lands tails.

More about P:

P is a probability function. A reminder of what this means formally: P's domain $\mathcal{F}$ is a set of sets of events containing the trivial event $\mathcal{S}$, and closed under complementation and countable intersection. P obeys the three Kolmogorov axioms.

1. (Necessity) $P(\mathcal{S}) = 1$

   The necessary event gets probability 1.

2. (Zero) For all $E \in \mathcal{F}$, $P(E) \geq 0$

   Every event has a probability of at least 0.

3. (Finite Additivity) If $E_1 \cap E_2 = \emptyset$, then $P(E_1 \cup E_2) = P(E_1) + P(E_2)$

   If two events are incompatible, then the probability of their disjunction is the sum of their probabilities.

P gives us a two-place notion of conditional probability $P(E|H)$ (read: "the probability of E given H"), which obeys the

Ratio Formula

$$P(E|H) = \frac{P(E \cap H)}{P(H)}$$

Interpretations of Probability In addition to formal features of P, we care about the philosophical interpretation of P. Here's the usual taxonomy of accounts. (The truth may wind up being some hybrid of these.) Whichever interpretation we choose, we'll need to explain why it is consistent with the Kolmogorov axioms.

Classical The probability of E is the following ratio:

$$\frac{\text{number of } E \text{ possibilities}}{\text{total number of possibilities}}$$

Frequentist Probabilities are always relative to a reference class. The probability of E (among the Rs) is the following ratio:

$$\frac{\text{number } R \text{ events that are } E \text{ events}}{\text{total number of } R \text{ events}}$$

Propensity Probabilities are long-run frequencies predicted by theoretical laws, or graded, single-case dispositions predicted by laws.

Subjective Probabilities are degrees of belief that an event will occur.

Evidential Probabilities measure the degree to which the hypothesis that a given event will occur is supported by the evidence.

More About u:

u is meant to measure something about the values of outcomes. But utilities are not monetary values.

- Suppose I need $10 for a bus ticket home. I will (reasonably) prefer a 50-50 gamble between $10 and $0 to a sure $7, and I will (reasonably) prefer a sure $10 to a 50-50 gamble between $26 and nothing.
- It's psychologically unrealistic to suppose that I (should) value gaining $200 twice as much as I value $100. The more money I already have, the less I (reasonably) care about additional money. Money, like many goods, has *diminishing marginal utility*.

Rather, u should be thought of as summarising information about preference strength. How much information does u give us? That is: which features of u represent reality?

- Another way of asking this question: what are the *allowable transformations* of the utility scale? (Moh's scale; length in inches vs. length in centimetres).
- Order matters: $u(o_1) > u(o_2)$ iff $o_1 \succ o_2$.

  (Note: I will use straight $\geq$, $>$, and $=$ for comparing magnitudes, and curly $\succsim$, $\succ$, and $\sim$ to denote preference.)

- Relative distance matters: $|u(o_1) - u(o_2)| > |u(o_3) - u(o_4)|$ iff the preference between o_1 and o_2 is stronger than the preference between o_3 and o_4.

  If relative distance did not matter, then we could make decision theory give different advice depending on which of two equally good representations of a situation we used.

- According to most decision theorists, the position of 0 and the size of the unit do not matter.

  Another way to say this: the set of allowable transformations is the set of *positive linear transformations*:

  $u^*(o) = xu(o) + y$

  (where x and y are real numbers with $x > 0$)

  Another way to say this: utility is measured by a linear scale.

Other Things With Utilities (besides acts and outcomes)

Lotteries Equivalence classes of acts that offer the same outcome values at the same probabilities.

Bets Assignments of probabilities to *changes* in utility.

Goods Items like money or wheat which can be had in greater or lesser quantities, and which contribute to overall outcome utilities.