Decision Theory

1.1 Setup

George Mason University, Spring 2018

Faced with a real-world or hypothetical choice situation, the first task of a decision theorist is to set out precisely the relevant actions of the decision maker, the relevant states of the world, the outcome of each act in each state, and the probability of each outcome (when these probabilities are available).

Decision Under Certainty: A choice situation where you know the outcomes of your actions.

Decision Under Risk: A choice situation in which it is possible to assign probabilities to all of the potential outcomes of each action.

Decision Under Ignorance: A choice situation in which it is impossible to assign probabilities to the potential outcomes of one or more of your actions.

For the time being, we will set up the machinery for analyzing decisions under total ignorance. Later on, we will add probabilities to the picture.

In the Oysters example, the situation can be broken down like this:

- **Actions:** eat raw oysters, eat grilled oysters, eat nothing.
- **States:** oysters are fresh, oysters are unfresh.
- **Outcomes:**

  If you eat raw fresh oysters, then you have a delicious meal.

  If you eat raw unfresh oysters, then you get sick.

  If you eat grilled fresh oysters, then you have a decent meal.

  If you eat grilled unfresh oysters, then you have a decent meal.

  If you eat nothing and the oysters are fresh, then you stay hungry.

  If you eat nothing and the oysters are unfresh, then you stay hungry.

We can explicate this informal choice situation using a formal mathematical model.

Def 1.1.1. A *decision model* $\mathcal{M} = \langle \mathcal{A}, \Omega, \mathcal{O}, g \rangle$ consists of a set of actions $\mathcal{A}$, a set of states Ω, a set of outcomes $\mathcal{O}$, and a function $g : \mathcal{A} \times \Omega \to \mathcal{O}$ that maps each action $a \in \mathcal{A}$ and state $s \in \Omega$ to an outcome $o \in \mathcal{O}$.

For instance, a decision model $\mathcal{M}$ of Oysters consists of:

- $\mathcal{A} = \{\text{eat raw}, \text{eat grilled}, \text{eat nothing}\}$
- $\Omega = \{\text{fresh}, \text{unfresh}\}$
- $\mathcal{O} = \{\text{delicious meal}, \text{sick}, \text{decent meal}, \text{hungry}\}$
- $g(\text{eat raw}, \text{fresh}) = \text{delicious meal}$
  $g(\text{eat raw}, \text{unfresh}) = \text{sick}$
  $g(\text{eat grilled}, \text{fresh}) = \text{decent meal}$
  And so forth.

A decision model can be visualized in a table or graph. For example, the model $\mathcal{M}$ for Oysters corresponds to the following *decision matrix*:

	fresh	unfresh
eat raw	delicious meal	sick
eat grilled	decent meal	decent meal
eat nothing	hungry	hungry

This model also corresponds to the following *decision tree* (where boxes are *decision nodes* that call for action and circles are *chance nodes* where nature 'decides' what happens next):

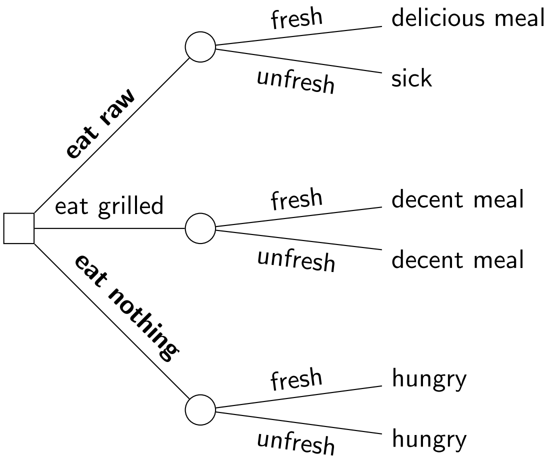

Aside: Decision trees are a particularly nice way to represent a *sequential* choice situation.

Ex. Socks.

In your bedroom, there are two drawers each of which contains exactly one red sock and one blue sock. Getting dressed in the morning, you must first select a sock from the first drawer and then select a sock from the second. You would very much like to wear matching socks. Which socks do you choose?

N.B. Assuming that you can see the color of the socks, this is a decision under certainty.

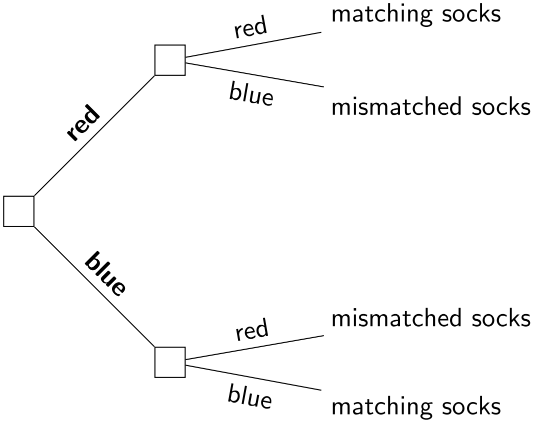

red
blue
red
blue
red
blue
matching socks
mismatched socks
mismatched socks
matching socks

Though the decision-theoretic information for Socks is more naturally expressed using a tree, this information can also be expressed using a matrix (with *strategies* on the left hand side):

	reality (@)
$\langle red, red \rangle$	matching socks
$\langle red, blue \rangle$	mismatched socks
$\langle blue, red \rangle$	mismatched socks
$\langle blue, blue \rangle$	matching socks

In general, we can go back and forth between matrices and trees. These are just two different means for visually organizing the same information in the underlying decision model.

Ex. Gray Day.

You walk out of your house without an umbrella on the way to the nearby bus stop. All of a sudden the sky darkens. If it starts to rain but your bus comes on time, then you will avoid getting too wet. But if your bus is late, then you will get drenched. On the other hand, if you go back inside and grab an umbrella, then you'll have to schlep it around all day. Do you get an umbrella?

How should we model Gray Day?

You walk out of your house without an umbrella on the way to the nearby bus stop. All of a sudden the sky darkens. If it starts to rain but your bus comes on time, then you will avoid getting too wet. But if your bus is late, then you will get drenched. On the other hand, if you go back inside and grab an umbrella, then you'll have to schlep it around all day. Do you get an umbrella?

- **Actions:**
- **States:**
- **Outcomes:**

How should we model Gray Day?

You walk out of your house without an umbrella on the way to the nearby bus stop. All of a sudden the sky darkens. If it starts to rain but your bus comes on time, then you will avoid getting too wet. But if your bus is late, then you will get drenched. On the other hand, if you go back inside and grab an umbrella, then you'll have to schlep it around all day. Do you get an umbrella?

- **Actions:** get umbrella, don't get umbrella
- **States:** rains & bus on time, rains & bus late, *no rain & bus on time, no rain & bus late*
- **Outcomes:**

  If you get the umbrella and bus is on time, then you have to schlep it around all day for no reason.

  If you get the umbrella and the bus is late, then you stay dry but have to schlep it around all day.

  If you don't get the umbrella and the bus is on time, then you stay dry and don't have to schlep it around all day for no reason.

  If you don't get the umbrella and the bus is late, then you get drenched but don't have to schlep it around all day.

How should we represent Gray Day in a formal model?

A decision model $\mathcal{M}$ of Gray Day consists of:

- $\mathcal{A} =$
- $\Omega =$
- $\mathcal{O} =$
- g

How should we represent Gray Day in a formal model?

A decision model $\mathcal{M}$ of Gray Day consists of:

- $\mathcal{A} = \{\text{get umbrella}, \text{don't get umbrella}\}$
- $\Omega =$
  $\{\text{rains \& on time}, \text{rains \& late}, \text{no rain \& on time}, \text{no rain \& late}\}$
- $\mathcal{O} = \{\text{useless shlep}, \text{dry}, \text{useful schlep}, \text{drenched}\}$
- $g(\text{get umbrella}, \text{rains \& on time}) = \text{useless schlep}$
  $g(\text{get umbrella}, \text{rains \& late}) = \text{useful schlep}$
  $g(\text{get umbrella}, \text{no rain \& on time}) = \text{useless schlep}$
  $g(\text{get umbrella}, \text{no rain \& late}) = \text{useless schlep}$
  $g(\text{don't get umbrella}, \text{rains \& on time}) = \text{dry}$
  $g(\text{don't get umbrella}, \text{rains \& late}) = \text{drenched}$
  $g(\text{don't get umbrella}, \text{no rain \& on time}) = \text{dry}$
  $g(\text{don't get umbrella}, \text{no rain \& late}) = \text{dry}$

How can we represent Gray Day in *decision matrix*?

	r & ot	r & late	no r & ot	no r & late
get umbrella	useless s	useful s	useless s	useless s
don't get umbrella	dry	drenched	dry	dry

Since the outcomes in the two right-hand most columns are the same, we can collapse into a single state of the world where it's not raining.

	r & ot	r & late	no rain
get umbrella	useless s	useful s	useless s
don't get umbrella	dry	drenched	dry

Why wouldn't we want to collapse the left-hand and right-hand columns?

Given the intricate nature of many choice situations, it is often a highly non-trivial matter to specify an appropriate formal decision model.

In Pharmaceutical Company, should the outcomes be specified in purely monetary terms? Should they incorporate human suffering and/or the company's goodwill?

In Love vs. Work, what possible states of the world are even relevant?

Unfortunately, there are often many plausible candidate decision models of a particular choice situation, and there is no hard and fast rule for selecting a single model in this set. Moreover, the cost of working with an inappropriate model can be enormous.

Applied decision theory is hard.

Our primary aim in this course is not to learn how to accurately model complex real-life choice situations, but rather to learn how to work with the abstract decision models themselves.

This being said, we will dabble in applied theory. So it is worth briefly mentioning some rough modeling principles.

1. The states in Ω should be mutually exclusive and exhaustive. Also, do not use a state space that is unnecessarily fine-grained.

	extremely fresh	fresh	unfresh
eat raw	delicious meal	delicious meal	sick
eat grilled	decent meal	decent meal	decent meal
eat nothing	hungry	hungry	hungry

You can always add a catch-all 'None of the previous states obtain' to Ω but it might be difficult to associate outcomes with acts in this state.

2. Try to avoid specifying states in Ω that are causally dependent on actions in $\mathcal{A}$, especially when probabilities are unavailable.

Ex. Superbowl.

Who do you think will win the Super Bowl? If you correctly guess that the New England Patriots will win, then you receive \$100. If you correctly guess that the Philadelphia Eagles will win, then you receive \$200.

	You win the bet	You lose the bet
bet on Patriots	win \$100	win \$0
bet on Eagles	win \$200	win \$0

2. Try to avoid specifying states in Ω that are causally dependent on actions in $\mathcal{A}$, especially when probabilities are unavailable.

Ex. Superbowl.

Who do you think will win the Super Bowl? If you correctly guess that the New England Patriots will win, then you receive \$100. If you correctly guess that the Philadelphia Eagles will win, then you receive \$200.

	Patriots wins	Eagles wins
bet on Patriots	win \$100	win \$0
bet on Eagles	win \$0	win \$200

In order to use a decision model $\mathcal{M}$ to prescribe or evaluate a particular choice, we must supplement $\mathcal{M}$ with information regarding the decision maker's *preferences*.

Many philosophers think that rationality does not mandate any single preference. Hume: "'Tis not contrary to reason to prefer the destruction of the whole world to the scratching of my finger."

But philosophers think that the preference profile of a rational agent will have a nice structure. In general, if a decision maker's preferences meet certain structural constraints (more on these later in the course), then aspects of her preferences can be usefully exhibited by *utility functions* that associate real numbers with outcomes in $\mathcal{O}$.

Time to introduce preference relations over $\mathcal{O}$.

$o_1 \succ o_2$ designates that o_1 is preferred to o_2.

$o_1 \sim o_2$ designates that o_1 and o_2 are preferred equally.

$o_1 \succcurlyeq o_2$ designates that o_1 is at least as preferred as o_2.

In Oysters, delicious meal $\succ$ decent meal $\succ$ hungry $\succ$ sick.

Note that these relations are interdefinable:

$o_1 \succ o_2$ if and only if $o_1 \succcurlyeq o_2$ and $o_1 \not\sim o_2$.

$o_1 \sim o_2$ if and only if $o_1 \succcurlyeq o_2$ and $o_2 \succcurlyeq o_1$.

$o_1 \succcurlyeq o_2$ if and only if $o_1 \succ o_2$ or $o_1 \sim o_2$.

Def 1.1.2. A function $u : \mathcal{O} \rightarrow \mathbb{R}$ is an *ordinal utility function* just in case it satisfies the following conditions:

(i) $u(o_1) > u(o_2)$ if and only if $o_1 \succ o_2$.

(ii) $u(o_1) = u(o_2)$ if and only if $o_1 \sim o_2$.

(iii) $u(o_1) \geq u(o_2)$ if and only if $o_1 \succcurlyeq o_2$.

(**N.B.** condition (i) implies (ii) and (iii).)

An ordinal utility function captures the *ordering* of the agent's preferences over the outcomes in $\mathcal{O}$.

It needn't capture the *strength* or *intensity* of the agent's preferences.

Ordinal utility functions are invariant up to *positive monotone transformations* (functions $t : \mathbb{R} \rightarrow \mathbb{R}$ where $t(x) \geq t(y)$ if and only if $x \geq y$). That is, if you apply a positive monotone transformation t to an ordinal utility function u, then the resulting composite function $t \circ u$ is still an ordinal utility function.

For example, the following function $u_1 : \mathcal{O} \to \mathbb{R}$ is an ordinal utility function:

$u_1(\text{delicious meal}) = 3$
$u_1(\text{decent meal}) = 2$
$u_1(\text{hungry}) = 0$
$u_1(\text{sick}) = -10$

Substituting these values for outcomes in the matrix for Oysters yields

	fresh	spoiled
eat raw	3	-10
eat grilled	2	2
eat nothing	0	0

The new function $u_2 : \mathcal{O} \mapsto \mathbb{R}$ obtained from u_1 by applying the positive monotone transformation $t(x) = x^3$ is also an ordinal utility function:

$u_2(\text{delicious meal}) = 27$
$u_2(\text{decent meal}) = 8$
$u_2(\text{hungry}) = 0$
$u_2(\text{sick}) = -1000$

Substituting these values for outcomes in the matrix for Oysters yields

	fresh	spoiled
eat raw	27	-1000
eat grilled	8	8
eat nothing	0	0

Again: ordinal utility functions needn't encode information about the strength of an agent's preferences.

The utility function u_1 suggests that the preference for a decent meal over being hungry is *stronger* than the preference for a delicious meal over a decent meal.

The utility function u_2 suggests that the preference for a decent meal over being hungry is *weaker* than the preference for a delicious meal over a decent meal.

Def 1.1.3. A function $u : \mathcal{O} \to \mathbb{R}$ is an *interval utility function* just in case $|u(o_1) - u(o_2)|$ reflects the relative strength of the agent's preference between $o_1, o_2 \in \mathcal{O}$.

Interval utility functions are invariant up to *positive linear transformations* (functions $t : \mathbb{R} \to \mathbb{R}$ where $t(x) = ax + b$ for $a > 0$). That is, if you apply a positive linear transformation t to an interval utility function u, then the resulting composite function $t \circ u$ is still an interval utility function.

Note that interval utility functions are ordinal utility functions but the converse is not true. For an agent's preferences to be representable with an interval utility function, her preference profile must satisfy even more constraints than are necessary for her preferences to be representable with an ordinal utility function.

We are all set up to start recommending and evaluating choices under ignorance. Keep in mind that we are interested in *rational* decisions that needn't lead to the best outcome when all is said and done.

Right Decision: A decision is *right* when there is no alternative action that leads to a better outcome in the actual world.

	fresh (@)	unfresh
eat raw	5	-10
eat grilled	2	2
eat nothing	0	0

Rational Decision: ???