Preference modeling

What do we talk about when we talk about "want"?

Suppose an *agent* (= someone who is making a decision) is facing a choice with possible outcomes taken from a set S.

In principle, we could represent the agent as an arbitrary function

$$\begin{aligned} \left\{\begin{matrix}\text{sets of}\\ \text{outcomes}\end{matrix}\right\} &\longrightarrow \{\text{outcomes}\} \\ S &\longmapsto \begin{matrix}\text{the outcome } x \in S\\ \text{which the agent chooses.}\end{matrix} \end{aligned}$$

Remark To simplify, we consider only *revealed preferences*, i.e. preferences you can infer from observing the agent's behavior.

There are two issues with representing agents this way:

(1) *Indifference*. This isn't expressive enough to represent when an agent is indifferent between multiple options. (i.e. there are choices A, B such that for any amount of money $m > 0$, the agent would take $A + \$m$ over B

and B+ $1m over A).

Could fix this by allowing the function to output a subset $B \subseteq S$ of its input S, representing the set of all (equally preferred) best options. But even then, theres a more serious issue:

(2) Too many functions. The agents represented by some functions are "too unreasonable" to be worth considering.

Example Consider the function

$$\{\text{chocolate}, \text{vanilla}\} \longmapsto \text{chocolate}$$

$$\{\text{vanilla}, \text{strawberry}\} \longmapsto \text{vanilla}$$

$$\{\text{chocolate}, \text{strawberry}\} \longmapsto \text{strawberry}$$

Let's suppose the agent is not indifferent between at least two of these, say between strawberry and chocolate. Then there is some amount of money \$m such that the agent would pay \$m to switch from chocolate to strawberry. You can then turn the agent into a money pump:

- Suppose the agent starts with strawberry.
- You offer to swap from strawberry to vanilla for free, and the agent accepts.
- You offer to swap from vanilla to chocolate for free and the agent accepts.
- You offer the agent to swap from chocolate to strawberry for \$m, and the agent accepts.
- Now the agent is back where they started but with \$m less.

Remark This is an example of a Dutch book argument, which demonstrates a strategy for draining certain agents of all their money. We call these agents exploitable, and generally we'll exclude them for purposes of preference modeling.

So we're looking for a more constrained formalism for representing preferences. One option is:

Definition A weak ordering on a set S is a relation $\leq$ on S satisfying

(1) Completeness:
for any $a, b \in S$, either $a \leq b$ or $b \leq a$ (or both)

(2) Transitivity:
for all $a, b, c \in S$, if $a \leq b$ and $b \leq c$ then $a \leq c$.

If S is a set of options, then we model an agent's preferences as a weak ordering on S, where $a \leq b$ means that the agent would either choose b out of the options $\{a, b\}$ or is indifferent between a and b.

Defn If $a, b \in S$ with $a \leq b$ and $b \leq a$ then we say the agent is indifferent between a and b, and write $a \sim b$.
If $a \leq b$ but not $b \leq a$, we write $a < b$ and say b is strictly preferred.

Q: Can the preferences of agent from the example before with S = {chocolate, vanilla, strawberry} be modeled with a weak ordering?

A: No, violates transitivity.

Aside: Some people disagree with the completeness assumption, since we often feel unable to compare two options. We're side-stepping the question by agreeing to focus on revealed preferences. But even for modelling "considered" preferences, one can show that every partial ordering (ordering with only transitivity assumption) can be completed into a weak ordering (which is still transitive).

Example Alice prefers chocolate to vanilla and strawberry, but finds vanilla and strawberry incomparable. You can complete this into either of the following orderings:

- $S \leq V \leq C$
- $V \leq S \leq C$.

A major issue with modeling preferences as weak orderings is that we can't get a sense of degree of preference. For example, I prefer chocolate ice cream a little more than vanilla, but much more than cyanide ice cream.

One option is to compare to money. E.g. I might prefer

$$-\$1M-1\text{¢} \le \text{Cyanide} < -\$1M < \$8 \le V < \$8.01 < \$10 \le C < \$10.01$$

in which case I value vanilla at $\approx \$8$ and chocolate at $\approx \$10$, and cyanide at $\approx -\$1M$ (i.e., you would have to pay me $\ge \$1M$ to eat cyanide ice cream).

Issue: the amount that I value money doesn't increase linearly with number of dollars, e.g. I value \$1B much more than half of \$2B (and similarly for other quantifiable goods: I value 1 private jet much more than half of 2 private jets).

Expected Utility

We'll introduce the abstraction of expected utility. Then we'll discuss assumptions on agent's preferences which imply they can be modelled as selecting actions which maximize expected utility.

Defn A random variable X with values in some

set E is an assignment of probabilities $\mathbb{P}(X=x)$ for each $x \in E$, subject to

(1) $\mathbb{P}(X=x) \in [0,1]$

(2) $\sum_{x \in E} \mathbb{P}(X=x) = 1$

We call X real-valued if $E \subseteq \mathbb{R}$.

Example We can model the outcome of a coin flip as a random variable X w/ values in $\{\text{heads}, \text{tails}\}$ with

$$\mathbb{P}(X = \text{heads}) = \mathbb{P}(X = \text{tails}) = \tfrac{1}{2}.$$

Defn If X is a real-valued R.V. then the expected value of X is

$$\mathbb{E}[X] = \sum_{x \in E} x\, \mathbb{P}(X=x).$$

Example Let X represent the outcome of a die roll (so $E = \{1,2,3,4,5,6\} \subseteq \mathbb{R}$). Then

$$\mathbb{E}[X] = \tfrac{1}{6}\cdot 1 + \tfrac{1}{6}\cdot 2 + \cdots + \tfrac{1}{6}\cdot 6$$
$$= \tfrac{1}{6}(1+2+3+4+5+6) = 3.5.$$

Naive idea: model an agent as having a utility function $u : S \longrightarrow \mathbb{R}$ so that

$$A \precsim B \iff u(A) \leq u(B).$$

Issue: This is basically the same as ranking: these utility values need not capture the notion of wanting B much more than A.

Idea: impose the condition that if $u(B) = 2u(A)$

then the agent is indifferent between getting A with certainty and a 50/50 chance of getting either B or nothing.

More formally:

Defn If E is a set of outcomes, then a lottery L over E is an E-valued random variable.

If $U : E \to \mathbb{R}$ is a utility function, then we extend U to lotteries over E via

$$U(L) := \mathbb{E}[U(L)] = \sum_{x \in E} U(x)\, \mathbb{P}(L = x).$$

We let $\mathcal{L}(E)$ denote the set of lotteries over E.

Example Let $E = \{\text{chocolate, vanilla, strawberry}\}$ and

$$U : \begin{aligned} C &\longmapsto 2 \\ V &\longmapsto 1 \\ S &\longmapsto 0 \end{aligned}$$

, and consider the lottery L which assigns

$$\mathbb{P}(L = C) = \tfrac{1}{4}$$
$$\mathbb{P}(L = V) = \tfrac{1}{2}$$
$$\mathbb{P}(L = S) = \tfrac{1}{4}$$

then $U(L) = \frac{1}{4} \cdot 2 + \frac{1}{2} \cdot 1 + \frac{1}{4} \cdot 0 = 1.$

Rmk Can view outcomes as lotteries: if A is an outcome, we can view it as the lottery that has value A with probability 1. We'll write A for either the outcome or the corresponding lottery.

Defn If $L_1, \dots, L_n$ are lotteries and

$p_1, \dots, p_n \in [0,1]$ with $\sum p_i = 1$, let

$\sum p_i L_i$ be the lottery which first picks a lottery L_i w/ prob p_i, and then selects an outcome from L_i.

Example $L_1 = \frac{1}{3}A + \frac{2}{3}B$

$$L_2 = \frac{1}{4}A + \frac{3}{4}C$$

$$\frac{2}{5}L_1 + \frac{3}{5}L_2 \text{ gives } \begin{cases} A & \text{w.p } \frac{2}{5}\cdot\frac{1}{3} + \frac{3}{5}\cdot\frac{1}{4} \\ B & \text{w.p. } \frac{2}{5}\cdot\frac{2}{3} \\ C & \text{w.p. } \frac{3}{5}\cdot\frac{3}{4} \end{cases}$$

$$= \frac{2}{5}\cdot\frac{1}{3}A + \frac{2}{5}\cdot\frac{2}{3}B + \frac{3}{5}\cdot\frac{1}{4}A + \frac{3}{5}\cdot\frac{3}{4}C$$

Fact In general, can reduce lotteries into simple lotteries like above.

Defn If $u : E \longrightarrow \mathbb{R}$, then define a weak ordering $\precsim_u$ on $\mathcal{L}(E)$ by

$$L \precsim_u M \iff u(L) \leq u(M).$$

This is the ordering induced by the utility fn u.

If $\precsim$ is an ordering on $\mathcal{L}(E)$ we say that $\precsim$ is represented by a utility fn u if $\precsim = \precsim_u$.

Rmk This gives a map $\left\{ \begin{matrix} \text{utility fns} \\ u : E \to \mathbb{R} \end{matrix} \right\} \longrightarrow \left\{ \begin{matrix} \text{weak orderings} \\ \text{on } \mathcal{L}(E) \end{matrix} \right\}$

$$u \longmapsto \precsim_u .$$

Qs. (1) When do two utility functions give the same weak ordering?

(2) Which weak orderings are induced by utility fns?

(3) Which weak orderings on $\mathcal{L}(E)$ are "reasonable" for agents to have.

We'll start w/ (1). It ends up that the answers to (2) and (3) coincide, and this is the vNM representation theorem!

Example The utility functions u and $u+100$ induces the same weak ordering on $\mathcal{L}(E)$.

Example $u(x)$ and $u'(x) = 2u(x)$ induce the same order.

Defn If $u : E \to \mathbb{R}$ is a function, then an affine transformation of u is a function $E \to \mathbb{R}$ of the form

$$x \mapsto \alpha u(x) + \beta$$

for $\alpha, \beta \in \mathbb{R}$. The affine transformation is positive if $\alpha > 0$.

Claim If $u'(x) = \alpha u(x) + \beta$ is a positive affine transformation, then 'u and 'u' induce the same preference orderings on lotteries.

Pf. Let $L, M \in \mathcal{L}(E)$.

$$L \lesssim_{u'} M \iff \sum_{x\in E} u'(x)\mathbb{P}(L=x) \leq \sum_{x\in E} u'(x)\mathbb{P}(M=x)$$

$$\iff \alpha \sum_{x\in E} u(x)\mathbb{P}(L=x) + \beta \sum_{x\in E} \mathbb{P}(L=x)$$

$$\leq \alpha \sum_{x\in E} u(x)\mathbb{P}(M=x) + \beta \sum_{x\in E} \mathbb{P}(M=x)$$

$$\iff \alpha u(L) + \beta \leq \alpha u(M) + \beta$$

$$\iff u(L) \leq u(M) \iff L \lesssim_u M \qquad \square$$

Rmk This implies that comparing utilities between agents is *not* sensible. Alice and Bob might have identical preferences, but we could model them as giving very *different utility values* to particular outcomes.

Soon, we'll prove the von Neumann-Morgenstern representation theorem, which is that "reasonable" agents have preferences representable by utility functions.

The vNM axioms capture what "reasonable" means.

Defn Let $\preceq$ be a weak ordering on $\mathcal{L}(E)$.

$\preceq$ satisfies:

(1) *continuity* if for all lotteries $L \preceq M \preceq N$, there is some $p \in [0,1]$ such that

$$M \sim pL + (1-p)N.$$

(2) *independence* if for all $L, M, N \in \mathcal{L}(E)$ and $p \in (0,1]$:

$$L \preceq M \iff pL + (1-p)N \preceq pM + (1-p)N.$$

Rmks Independence is the more controversial assumption. For example:

	1%	34%	65%
L_1	\$0	\$2500	\$2400
L_2	\$2400	\$2400	\$2400

many people prefer $L_2 \succeq L_1$.

	1%	34%	65%
L_3	\$0	\$2500	\$0
L_4	\$2400	\$2400	\$0

many people choose $L_3 \succeq L_4$.

These two choices together violate independence.

Claim Agents who violate independence can be money pumped.

Argument Suppose Alice prefers $A \succ B$ but

$$pA + (1-p)C \preceq pB + (1-p)C.$$

Suppose you have $pA + (1-p)C$. I offer you to swap to $pB + (1-p)C$. If it comes out B I offer to swap A for B for 1¢. In the end, you got the lottery $pA + (1-p)C$ which you had at the start, but paid 1¢ for it.

Rmk

- This is not a repeatable money pump
- Later in the course, we might talk about reasonable ways of decision-making which don't fall for this money pump, by refusing to swap A after the lottery.

Thm (von Neuman-Morgenstern representation thm)

A weak order $\preceq$ on $\mathcal{L}(E)$ is induced by a utility function iff $\preceq$ satisfies continuity and independence.

Pf [Given by Lucy.]

Rmk If you satisfy the vNM conditions, then the theorem says you can be modelled as maximizing a utility fn in expectation.

Corollary

If a weak ordering on $\mathcal{L}(E)$ is induced by a utility function $U: E \to \mathbb{R}$, then U is unique up to positive affine transformation.

Pf. We've already seen that positive affine transformations of U induce the same weak order. Thus, it suffices to show that if $U, U' : E \longrightarrow \mathbb{R}$ induce the same order $\lesssim = \lesssim_U = \lesssim_{U'}$ on $\mathcal{L}(E)$ then there are $\alpha, \beta \in \mathbb{R}$, $\alpha > 0$ with $U'(x) = \alpha U(x) + \beta$.

Order the outcomes in E as

$$A_1 \lesssim A_2 \lesssim \cdots \lesssim A_n.$$

If $A_1 \sim A_n$, then U and U' are both constant, and we can take $\alpha = 1$, $\beta = U'(A) - U(A)$ for any $A \in E$.

Otherwise, $A_1 < A_n$. We can solve the system of 2 equations in 2 variables α, β

$$\begin{cases} U'(A_1) = \alpha U(A_1) + \beta \\ U'(A_n) = \alpha U(A_n) + \beta \end{cases}$$

to get $\alpha = \dfrac{U'(A_n) - U'(A_1)}{U(A_n) - U(A_1)}$, $\beta = [\text{something}]$.

Since $A_1 < A_n$, $\alpha > 0$.

By the vNM theorem, $\lesssim$ satisfies continuity. Thus for any $A \in E$ there is $p \in [0,1]$ s.t.

$A \sim pA_1 + (1-p)A_n$.

Using the α, β from above, we then have

$$\begin{aligned} u'(A) &= pu'(A_1) + (1-p)u'(A_n) \\ &= p(\alpha u(A_1) + \beta) + (1-p)(\alpha u(A_n) + \beta) \\ &= \alpha(pu(A_1) + (1-p)u(A_n)) + \beta \\ &= \alpha u(A) + \beta \end{aligned}$$

as desired.