Learning Objectives

At the end of the class you should be able to:

- justify the use and semantics of utility
- estimate the utility of an outcome
- build a single-stage decision network for a domain
- compute the optimal decision of a single-stage decision network

Preferences

- Actions result in outcomes
- Agents have preferences over **outcomes**
- Rational agent will do the action that has the best outcome for them
- Sometimes agents don't know the outcomes of the actions, but they still need to compare actions
- Real agents have to act.
  (Doing nothing is (usually) an action).

Preferences Over Outcomes

If o_1 and o_2 are outcomes

- **Weakly preferred:**
  $o_1 \succeq o_2$ means o_1 is at least as desirable as o_2.
- **Indifferent:**
  $o_1 \sim o_2$ means $o_1 \succeq o_2$ and $o_2 \succeq o_1$.
- **Strictly preferred:**
  $o_1 \succ o_2$ means $o_1 \succeq o_2$ and $o_2 \not\succeq o_1$

Lotteries

- Agent may not know the outcomes of actions, but only have a probability distribution of outcomes.
- **Lottery** is a probability distribution over outcomes:

$$[p_1 : o_1, p_2 : o_2, \ldots, p_k : o_k]$$

where the o_i are outcomes and $p_i \geq 0$ such that

$$\sum_i p_i = 1$$

The lottery specifies that outcome o_i occurs with probability p_i.

Properties of rational preferences

- **Completeness:** Agents have to act, and thus they must have preferences:

$$\forall o_1 \forall o_2 \; o_1 \succeq o_2 \text{ or } o_2 \succeq o_1$$

Properties of rational preferences

Transitivity: Preferences must be transitive:

if $o_1 \succeq o_2$ and $o_2 \succ o_3$ then $o_1 \succ o_3$

(Similarly for other mixtures of $\succ$ and $\succeq$.)
Rationale: otherwise $o_1 \succeq o_2$ and $o_2 \succ o_3$ and $o_3 \succeq o_1$.
If they are prepared to pay to get o_2 instead of o_3,
and are happy to have o_1 instead of o_2,
and are happy to have o_3 instead of o_1
$\longrightarrow$ money pump.

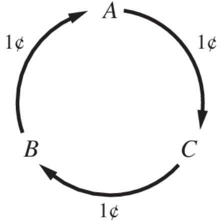

Properties of rational preferences

Monotonicity: An agent prefers a larger chance of getting a better outcome than a smaller chance:

- If $o_1 \succ o_2$ and $p > q$ then

$$[p : o_1, 1 - p : o_2] \succ [q : o_1, 1 - q : o_2]$$

Properties of rational preferences

Decomposability: (no fun in gambling). An agent is indifferent between lotteries that have same probabilities and outcomes. This includes lotteries over lotteries. For example:

$$[p : o_1, 1-p : [q : o_2, 1-q : o_3]]$$
$$\sim \quad [p : o_1, (1-p)q : o_2, (1-p)(1-q) : o_3]$$

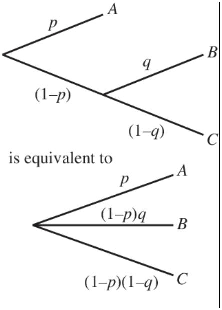

Properties of rational preferences

- Suppose $o_1 \succ o_2$ and $o_2 \succ o_3$. Consider whether the agent would prefer
  - o_2
  - the lottery $[p : o_1, 1 - p : o_3]$

  for different values of $p \in [0, 1]$.
- Plot which one is preferred as a function of p:

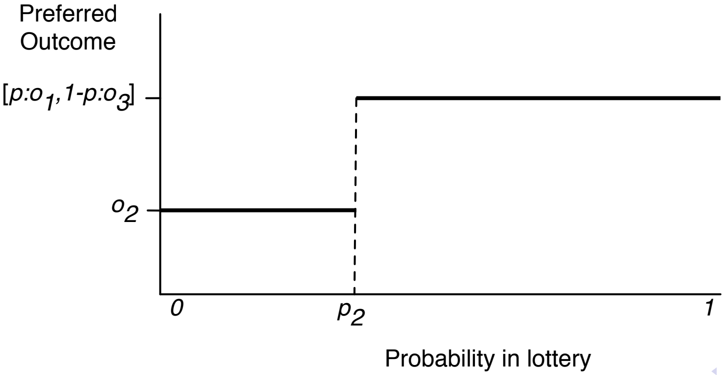

Properties of rational preferences

Continuity: Suppose $o_1 \succ o_2$ and $o_2 \succ o_3$, then there exists some $p \in [0,1]$ such that

$$o_2 \sim [p : o_1, 1 - p : o_3]$$

Properties of rational preferences

Substitutability: if $o_1 \sim o_2$ then the agent is indifferent between lotteries that only differ by o_1 and o_2:

$$[p : o_1, 1 - p : o_3] \sim [p : o_2, 1 - p : o_3]$$

Alternative Axiom for Substitutability

Substitutability: if $o_1 \succeq o_2$ then the agent weakly prefers lotteries that contain o_1 instead of o_2, everything else being equal. That is, for any number p and outcome o_3:

$$[p : o_1, (1-p) : o_3] \succeq [p : o_2, (1-p) : o_3]$$

Rationality

- An agent is defined to be rational if it obeys the completeness, transitivity, monotonicity, decomposability, continuity, and substitutability axioms.
- Rationality also depends on subjective utility (as we will define now)

Utility: What we would like

- We would like a measure of preference that can be combined with probabilities. So that

$$\begin{aligned} & \mathit{value}([p : o_1, 1-p : o_2]) \\ = \; & p \times \mathit{value}(o_1) + (1-p) \times \mathit{value}(o_2) \end{aligned}$$

- What would you prefer ?

$$\$1,000,000 \text{ or } [0.5 : \$0, 0.5 : \$2,000,000]?$$

- We want non-linearity or arbitrary functions: Perceived value of money and actual benefit of money is not linear.

Theorem

If preferences follow the preceding properties, then preferences can be measured by a function

$$utility : outcomes \rightarrow [0,1]$$

such that

- $o_1 \succeq o_2$ if and only if $utility(o_1) \geq utility(o2)$.
- Utility is calculated as:

$$utility([p_1 : o_1, p_2 : o_2, \ldots, p_k : o_k]) = \sum_{i=1}^{k} p_i \times utility(o_i)$$

Many possible utility functions exist

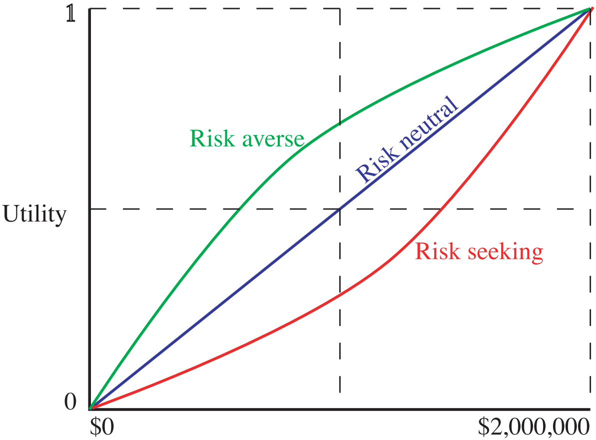

Why? Perceived value, actual value, or both?
Can be generated empirically via querying people
$[p : u_1, 1 - p : u_2]$ with various p

Possible utility as a function of money

y-axis is utility
x-axis is money

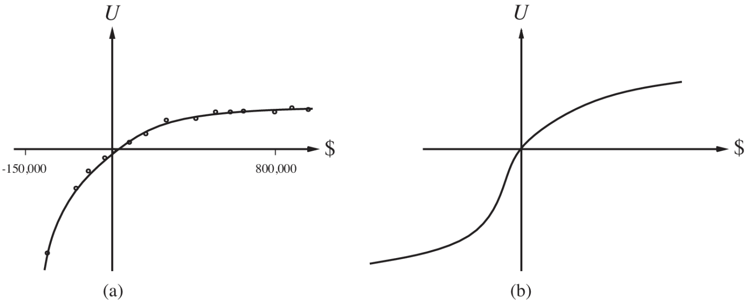

(a) Experimental data (b) Full curve.

Utility can be arbitrary

Someone who really wants a toy worth $30, but who would also like one worth $20:

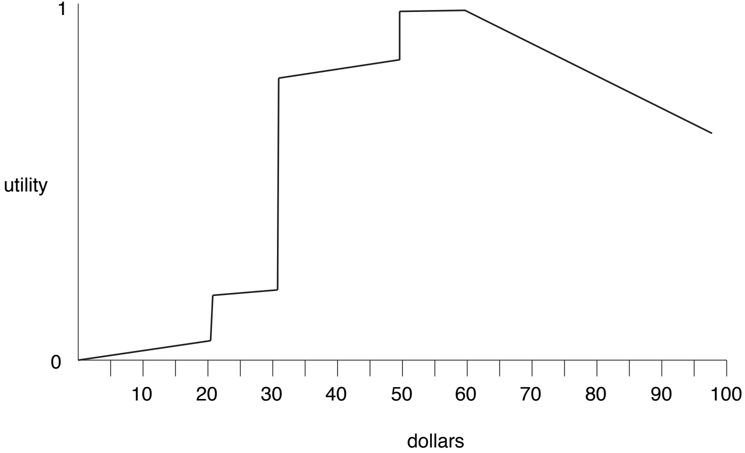

Factored Representation of Utility

- Suppose the outcomes can be described in terms of features $X_1, \ldots, X_n$.
- An **additive utility** is one that can be decomposed into set of factors:

$$u(X_1, \ldots, X_n) = f_1(X_1) + \cdots + f_n(X_n).$$

  This assumes **additive independence**.
- Strong assumption: contribution of each feature doesn't depend on other features.
- Many ways to represent the same utility:
  — a number can be added to one factor as long as it is subtracted from others.

Additive Utility

- An additive utility has a canonical representation:

$$u(X_1, \ldots, X_n) = w_1 \times u_1(X_1) + \cdots + w_n \times u_n(X_n).$$

- If $best_i$ is the best value of X_i, $u_i(X_i{=}best_i) = 1$.
  If $worst_i$ is the worst value of X_i, $u_i(X_i{=}worst_i) = 0$.
- w_i are weights, $\sum_i w_i = 1$.
  The weights reflect the relative importance of features.
- We can determine weights by comparing outcomes.

$$w_1 = u(best_1, x_2, \ldots, x_n) - u(worst_1, x_2, \ldots, x_n).$$

for any values $x_2, \ldots, x_n$ of $X_2, \ldots, X_n$.

Complements and Substitutes

- Often additive independence is not a good assumption.
- Values x_1 of feature X_1 and x_2 of feature X_2 are **complements** if having both is better than the sum of the two.
- Example: on a holiday
  - having a plane booking for a particular day and a hotel booking for the same day are complements: one without the other does not give a good outcome.
- Values x_1 of feature X_1 and x_2 of feature X_2 are **substitutes** if having both is worse than the sum of the two.
- Example: on a holiday
  - Two trips in one day

Generalized Additive Utility

- If there are interactions (e.g., complement or substitute)
- Generalized additive utility can be written as a sum of factors:

$$u(X_1, \ldots, X_n) = f_1(\overline{X_1}) + \cdots + f_k(\overline{X_k})$$

where $\overline{X_i} \subseteq \{X_1, \ldots, X_n\}$.

Fallacies and paradoxes

Humans are not internally consistent rational agents...
or are they?

Allais Paradox (1953)

What would you prefer:

A: %80 chance of $\$4,000$

B: %100 chance of $\$3,000$

What would you prefer:

C: %20 chance of $\$4,000$

D: %25 chance of $\$3,000$

Most people like B over A, and C over D, which isn't internally consistent.

Framing Effects [Tversky and Kahneman]

- A disease is expected to kill 600 people. Two alternative programs have been proposed:

  Program A: 200 people will be saved

  Program B: probability 1/3: 600 people will be saved
  probability 2/3: no one will be saved

  Which program would you favor?

A disease is expected to kill 600 people. Two alternative programs have been proposed:

Program C: 400 people will die

Program D: probability 1/3: no one will die
probability 2/3: 600 will die

Which program would you favor?

Tversky and Kahneman: 72% chose A over B. 22% chose C over D.

Framing Effects [Tversky and Kahneman]

- A disease is expected to kill 600 people. Two alternative programs have been proposed:

  Program A: 200 people will be saved

  Program B: probability 1/3: 600 people will be saved

  probability 2/3: no one will be saved

  Which program would you favor?

A disease is expected to kill 600 people. Two alternative programs have been proposed:

Program C: 400 people will die

Program D: probability 1/3: no one will die

probability 2/3: 600 will die

Which program would you favor?

Tversky and Kahneman: 72% chose A over B.
22% chose C over D.

Framing Effects

- Suppose you had bought tickets for the theatre for \$50. When you got to the theatre, you had lost the tickets. You have your credit card and can buy equivalent tickets for \$50. Do you buy the replacement tickets on your credit card?
- Suppose you had \$50 in your pocket to buy tickets. When you got to the theatre, you had lost the \$50. You have your credit card and can buy equivalent tickets for \$50. Do you buy the tickets on your credit card?

Prospect Theory

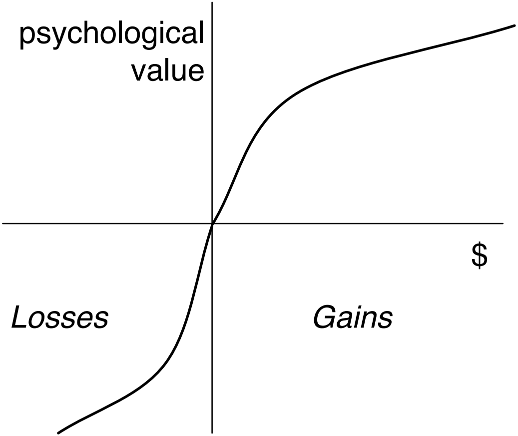

- In mixed gambles, loss aversion causes extreme risk-averse choices

Reference Points

Consider Anthony and Betty:

- Anthony's current wealth is $1 million.
- Betty's current wealth is $4 million.

They are both offered the choice between a gamble and a sure thing:

- Gamble: equal chance to end up owning $1 million or $4 million.
- Sure Thing: own $2 million

What does expected utility theory predict?
What does prospect theory predict?
Is this actually rational?

The Ellsberg Paradox

Box contains 1/3 red balls, 2/3 either black or yellow (unknown proportion)

A: \$100 for red

B: \$100 for black

What would you prefer:

C: \$100 for red or yellow

D: \$100 for black or yellow

If red is greater than black, most people like A over B, and D over C, which isn't internally consistent, perhaps due to ambiguity aversion.

Predictor Paradox

Two boxes:

Box 1: contains \$10,000

Box 2: contains either \$0 or \$1m

- You can either choose both boxes or just box 2.
- The "predictor" has put \$1m in box 2 if he thinks you will take box 2 and \$0 in box 2 if he thinks you will take both.
- The predictor has been correct in previous predictions.
- Do you take both boxes or just box 2?