Slide 1

Making Simple Decisions

Outline

- Combining Beliefs and Desires Under Uncertainty
- Utility Theory
- Decision Networks
- Value of Information

Slide 2

Combining Beliefs and Desires Under Uncertainty

- Agent's preferences between world states are captured by a utility function (assigns a single number $U(S)$ to express the desirability of a state S).
- Expected utility of an action A

$$\mathrm{EU}(A|E) = \sum_i P(Result_i(A)|E, Do(A))U(Result_i(A))$$

where $Result_i(A)$ are all the possible outcomes of A.

Slide 3

- Principle of maximum expected utility (MEU):
  a rational agent should choose an action that maximizes the agent's expected utility.
- If an agent maximizes a utility function that correctly reflects the performance measure by which the behavior is being judged, then it will achieve the highest possible performance score if we average over the possible environments in which the agent could be placed.
- One-shot decisions/sequential decisions

Slide 4

Utility Theory

- Why MEU reasonable?
  Do utility functions always exists?
- Lotteries
  $L = [p, A; 1-p, B]$
  A $([1, A])$
- Preferences between lotteries or states
  $A \succ B$
  $A \sim B$
  $A \succeq B$

Slide 5

Axioms of utility theory

- Orderability: $(A \succ B) \vee (B \succ A) \vee (A \sim B)$
- Transitivity: $(A \succ B) \wedge (B \succ C) \Rightarrow (A \succ C)$
- Continuity: $A \succ B \succ C \Rightarrow \exists p [p, A; 1-p, C] \sim B$
- Substitutability:
  $A \sim B \Rightarrow [p, A; 1-p, C] \sim [p, B; 1-p, C]$
- Monotonicity:
  $A \succ B \Rightarrow (p \geq q \Leftrightarrow [p, A; 1-p, B] \succeq [q, A; 1-q, B]$
- Decomposability: $[p, A; 1-p, [q, B; 1-q, C]] \sim$
  $[p, A; (1-p)q, B; (1-p)(1-q), C]$

Slide 6

The existence of a utility function follows from the axioms:

- Utility principle:
  If an agent's preferences obey the axioms of utility theory, then there is a real-valued function U that operates on states such that
  $U(A) > U(B) \Leftrightarrow A \succ B$
  $U(A) = U(B) \Leftrightarrow A \sim B$
- Maximum Expected Utility principle:
  The utility of a lottery

$$U([p_1, S_1; \ldots; p_n, S_n]) = \sum_i p_i U(S_i)$$

Slide 7

Utility functions

- The utility of money
  monotonic preference for money
  money: value function/ordinal utility measure
- Expected monetary value (EMV)
  Example. Win $1M\mathrm{FIM}$/flip a coin for $3M\mathrm{FIM}$ or $0\mathrm{FIM}$
  $0.5 \times 0\mathrm{FIM} + 0.5 \times 3000000\mathrm{FIM} = 1500000\mathrm{FIM}$
  $1 \times 1000000\mathrm{FIM} = 1000000\mathrm{FIM}$

  Expected utilities:

$$\begin{aligned} \mathrm{EU}(Accept) &= 0.5U(S_k) + 0.5U(S_{k+3000000}) \\ \mathrm{EU}(Decline) &= U(S_{k+1000000}) \end{aligned}$$

Slide 8

- Actual utility functions: logarithmic form
- Risk-averse
  $U(S_L) < U(S_{\mathrm{EMV}(L)})$
- Risk-seeking
- Certainty equivalent for a lottery:
  $400 for a gamble that gives $1000 half the time
- Insurance premium:
  difference between EMV and certainty equivalent for a lottery.
- Risk-neutral: linear curve

Slide 9

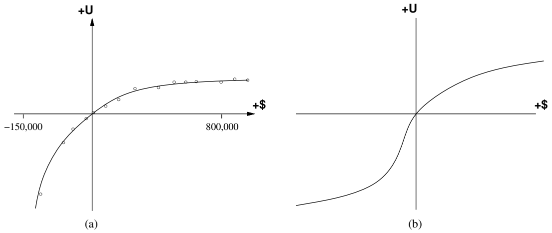

Slide 10

Utility scales and assessment

- Agents using $U(S)$ and $U'(S) = k_1 + k_2 U(S)$, respectively, behave identically where k_1 and k_2 are constants and k_2 positive.
- Normalized utilities:
  best possible prize: $u_\top = 1$,
  worst possible catastrophe: $u_\perp = 0$
  intermediate outcomes S: $U(S) = p$ where
  the agent is indifferent between S and $[p, \top; (1-p), \perp]$.
- value on human life
  micromort (1/1000000 chance of death): $20

Slide 11

Multiattribute utility functions

- Representation theorems:

$$U(x_1, \ldots, x_n) = f[f_1(x_1), \ldots, f_n(x_n)]$$

- Dominance: narrows down the choices.

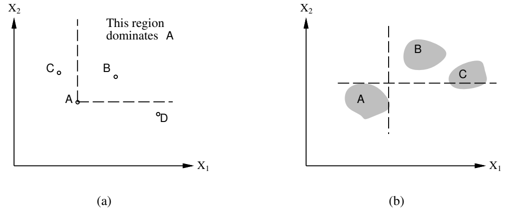

Slide 12

Decision Networks

(Influence Diagrams)

- Chance nodes (ovals)
  random variables with CPTs
- Decision nodes (rectangles)
- Utility nodes (diamonds)
  agent's utility function
- Simplified form:
  chance nodes describing the outcome state omitted;
  action-utility tables: expected utility of actions

Slide 13

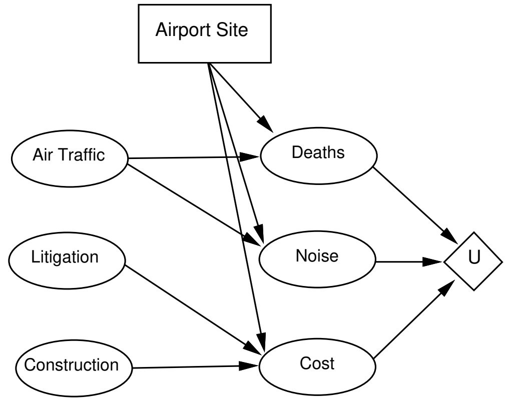

Slide 14

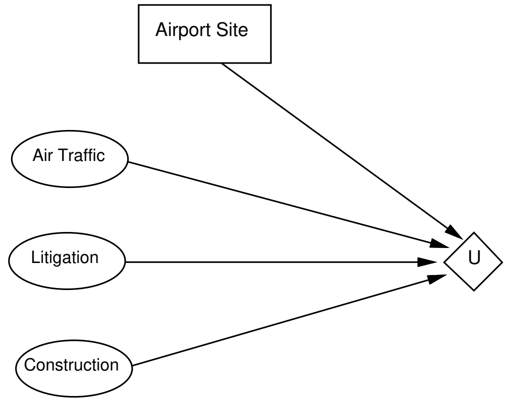

Slide 15

Evaluating decision networks

1. Set the evidence variables for the current state.
2. For each possible value of the decision node:
   (a) Set the decision node to that value.
   (b) Calculate the posterior probabilities for the parent nodes of the utility node.
   (c) Calculate the resulting utility for the action.
3. Return the action with the highest utility.

Slide 16

Value of Information

- The value of a piece of information:
  the difference in expected value between best actions before and after information is obtained.
- Simple example: n blocks for sale.
  Exactly one block contains oil worth C dollars.
  Price of a block C/n dollars.
  How much to pay to know whether block 3 contains oil?
  Expected value of the information:
  $$\frac{1}{n}(C - \frac{C}{n}) + \frac{n-1}{n}(\frac{C}{n-1} - \frac{C}{n}) - (\frac{C}{n} - \frac{C}{n}) = \frac{C}{n}$$