Cooperating Intelligent Systems

Utility theory

Chapter 16, AIMA

The utility function $U(S)$

- An agent's preferences between different states S in the world are captured by the *utility function* $U(S)$.

- If $U(S_i) > U(S_j)$ then the agent prefers state S_i over state S_j
- If $U(S_i) = U(S_j)$ then the agent is indifferent between the two states S_i and S_j

Maximize expected utility

A rational agent should choose the action that maximizes the the agent's expected utility (EU):

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

Where $Result_i(A)$ enumerates all the possible resulting states after doing action A given observed environment state (evidence) E.

The basis of utility theory

Notation:

$A \succ B$ A is preferred to B

$A \sim B$ The agent is indifferent between A and B

$A \succsim B$ The agent prefers A to B, or is indifferent between them.

A lottery is described with

$$L = [p_1, C_1; p_2, C_2; \ldots; p_n, C_n]$$

The six axioms of utility theory

Orderability $(A \succ B) \vee (B \succ A) \vee (A \sim B)$ You must make a decision

Transitivity $(A \succ B) \wedge (B \succ C) \Rightarrow (A \succ C)$

Continuity $(A \succ B \succ C) \Rightarrow \exists p\, [p,A;\, 1\text{-}p,C] \sim B$

Substitutability $(A \sim B) \Rightarrow [p,A;\, 1\text{-}p,C] \sim [p,B;\, 1\text{-}p,C]$

Monotonicity $(A \succ B) \Rightarrow (p \geq q \Leftrightarrow [p,A;\, 1\text{-}p,B] \succsim [q,A;\, 1\text{-}q,B])$

Decomposability $[p,A;\, 1\text{-}p,[q,B;\, 1\text{-}q,C]] \sim [p,A;\, (1\text{-}p)q,B;\, (1\text{-}p)(1\text{-}q),C]$

It follows from these axioms that there exists a real-valued function U that operates on states such that

$$U(A) > U(B) \Leftrightarrow A \succ B$$
$$U(A) = U(B) \Leftrightarrow A \sim B$$

The St. Petersburg ”paradox”

You are offered to play the following game (bet): You flip a coin repeatedly until you get your first heads. You will then be paid $2 to the power of every flip you made, including the final one (the price matrix is below).

How much are you willing to pay to participate (participation is not free)?

Toss	**Winning**
H	$2
TH	$4
TTH	$8
TTTH	$16
…	…

The St. Petersburg ”paradox”

What is the expected winning in this betting game?

$$\langle \text{Winning} \rangle_N = \sum_{k=1}^{N} P(k)W(k) = \sum_{k=1}^{N} \frac{2^k}{2^k} = \sum_{k=1}^{N} 1 = N$$

A rational player should be willing to pay any sum of money to participate...
...if $ = Utility

The students in previous years’ classes have offered $4 or less on average...

The St. Petersburg ”paradox”

Bernoulli (1738): The utility of money is not money; it is more like log(Money).

$$\left\langle \text{Utility} \right\rangle_N = \sum_{k=1}^{N} P(k)U(k) = \sum_{k=1}^{N} \frac{\ln(2^k)}{2^k} = \ln(2)\sum_{k=1}^{N} \frac{k}{2^k}$$

$$= 2\ln(2)\left[1 - \left(\frac{1}{2}\right)^N - N\left(\frac{1}{2}\right)^{N+1}\right]$$

$$\lim_{N\to\infty} \left\langle \text{Utility} \right\rangle_N = 2\ln(2)$$

”Mr. Beard’s” utility curve

General ”human nature” utility curve

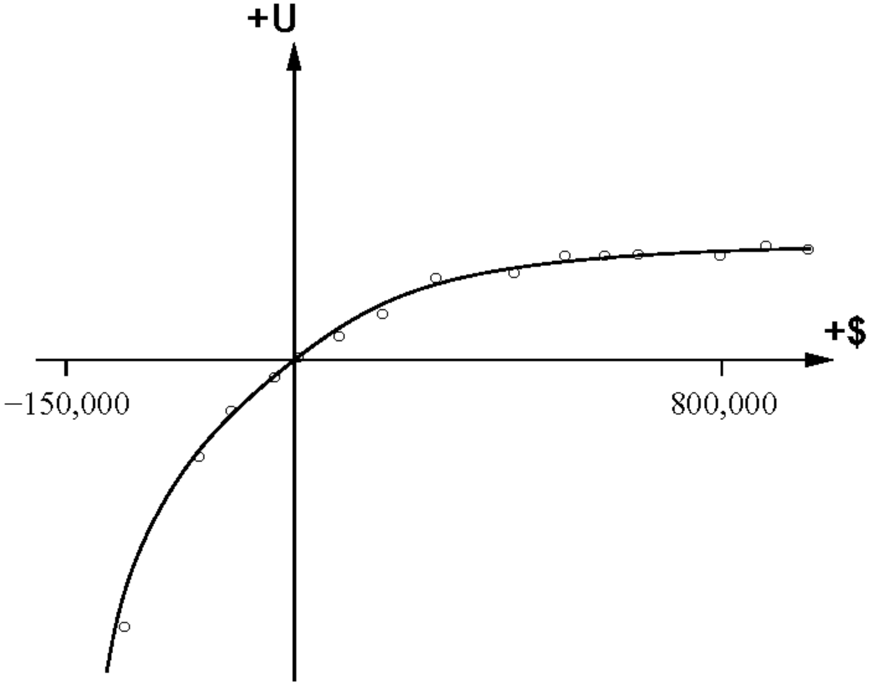

(a)

Risk seeking

Lottery game 1

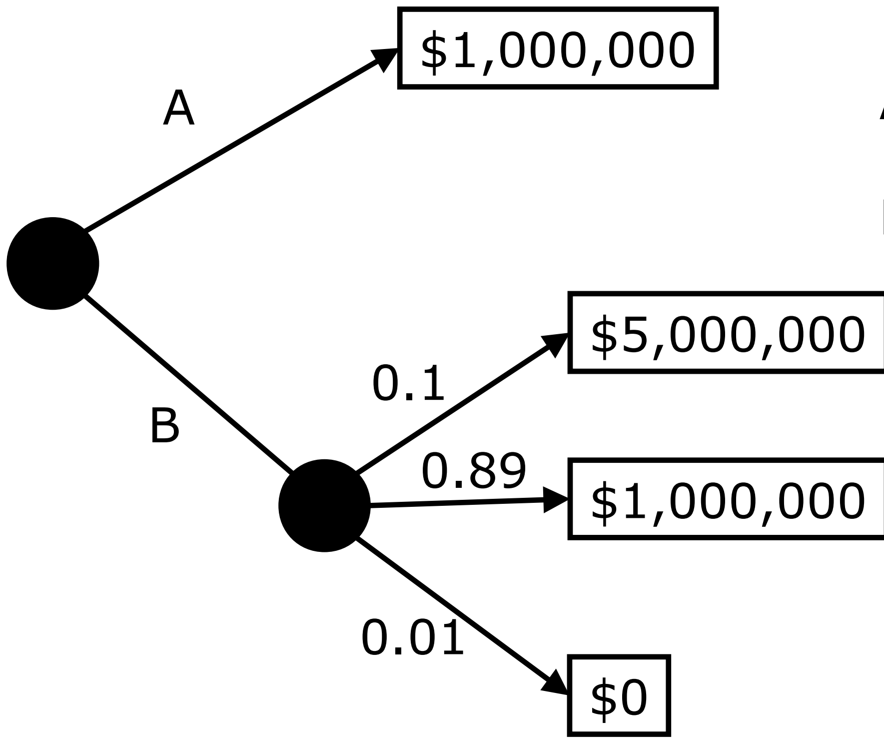

You can choose between alternatives A and B:

A) You get $1,000,000 for sure.

B) You can participate in a lottery where you can win up to $5 mill.

Lottery game 2

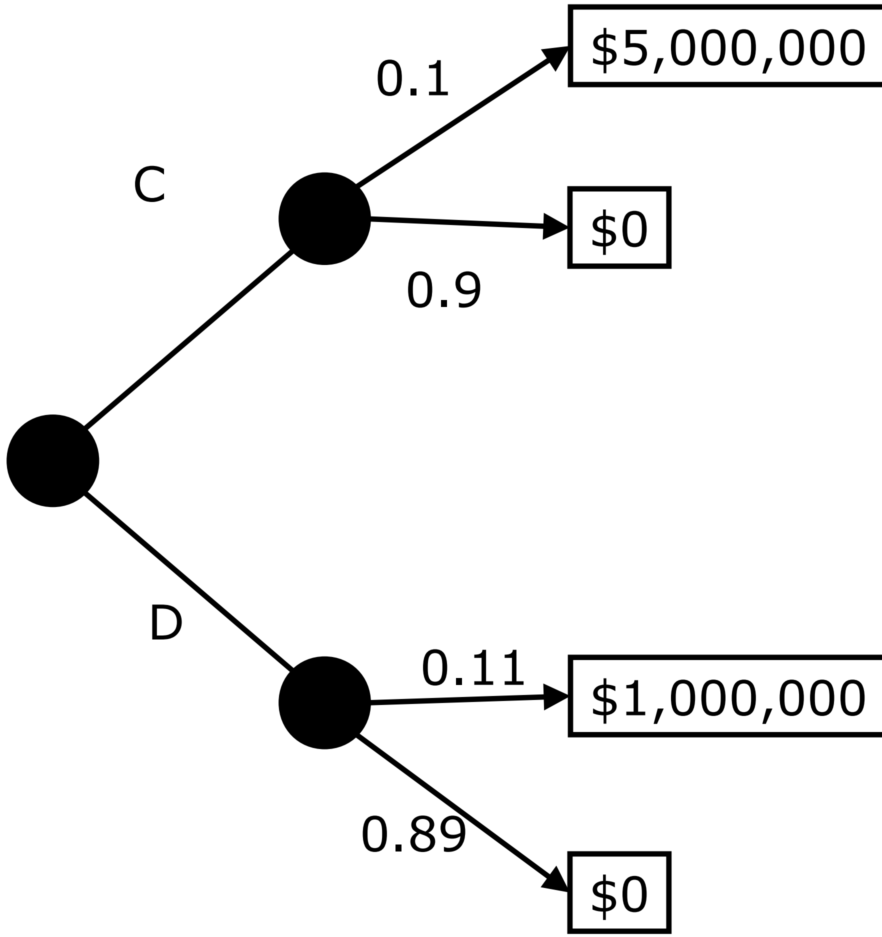

You can choose between alternatives C and D:

C) A lottery where you can win $5 mill.

D) A lottery where you can win $1 mill.

Lottery preferences

- People should select A and D, or B and C. Otherwise they are not being consistent...

$$U(A) - U(B) = U(\$1\text{M}) - 0.1U(\$5\text{M}) - 0.89U(\$1\text{M}) - 0.01U(\$0) =$$
$$0.11U(\$1\text{M}) - 0.1U(\$5\text{M}) - 0.01U(\$0)$$

$$U(D) - U(C) = 0.11U(\$1\text{M}) + 0.89U(\$0) - 0.1U(\$5\text{M}) - 0.9U(\$0) =$$
$$0.11U(\$1\text{M}) - 0.1U(\$5\text{M}) - 0.01U(\$0)$$

Allais paradox. Utility function does not capture a human's fear of looking like a complete idiot.

In last years classes, fewer than 50% have been consistent...

Form of *U(S)*

If the value of one attribute does not influence one's opinion about the preference for another attribute, then we have mutual preferential independence and can write:

$$V(X_1, X_2) = V_1(X_1) + V_2(X_2)$$

Where *V(X)* is a value function (expressing the [monetary] value)

Example: The party problem

We are about to give a wedding party. It will be held during summer-time.

Should we be outdoors or indoors?

The party is such that we can’t change our minds on the day of the party (different locations for indoors and outdoors).

What is the rational decision?

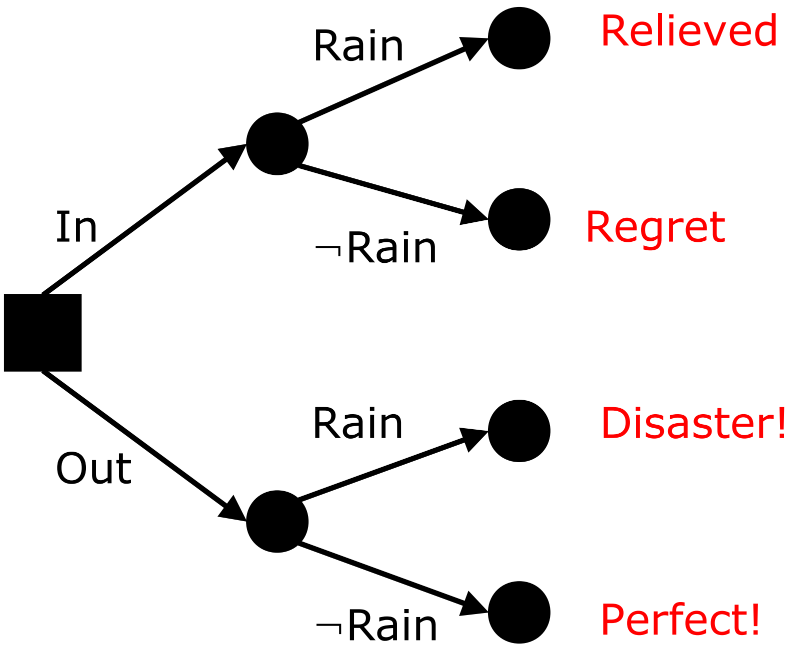

Example adapted from Breese & Kooler 1997

Example: The party problem

The value function: Assign a numerical (monetary) value to each outcome.

(We avoid the question on how this is done for the time being)

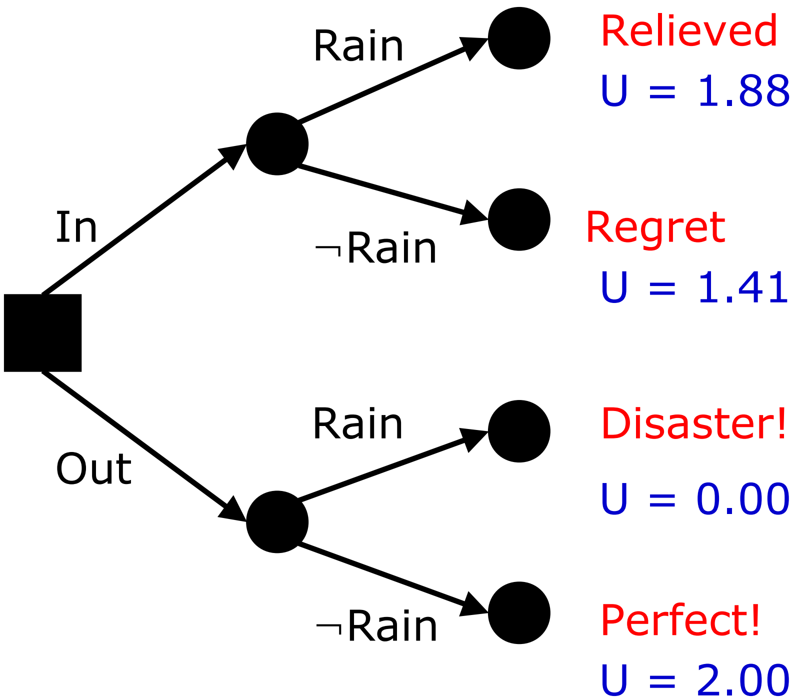

Location	**Weather**	**Value**
Indoors	Sun	$25
Indoors	Rain	$75
Outdoors	Sun	$100
Outdoors	Rain	$0

Let U(S) = log[V(S)+1]

Example adapted from Breese & Kooler 1997

Example: The party problem

Get weather statistics for your location in the summer (June).

Location	P(Rain)
Stockholm, Sweden	18/30
Bergen, Norway	19/30
San Fransisco, USA	1/30
Seattle, USA	9/30
Paris, France	14/30
Haifa, Israel	0/30

Rain probabilities from Weatherbase www.weatherbase.com/

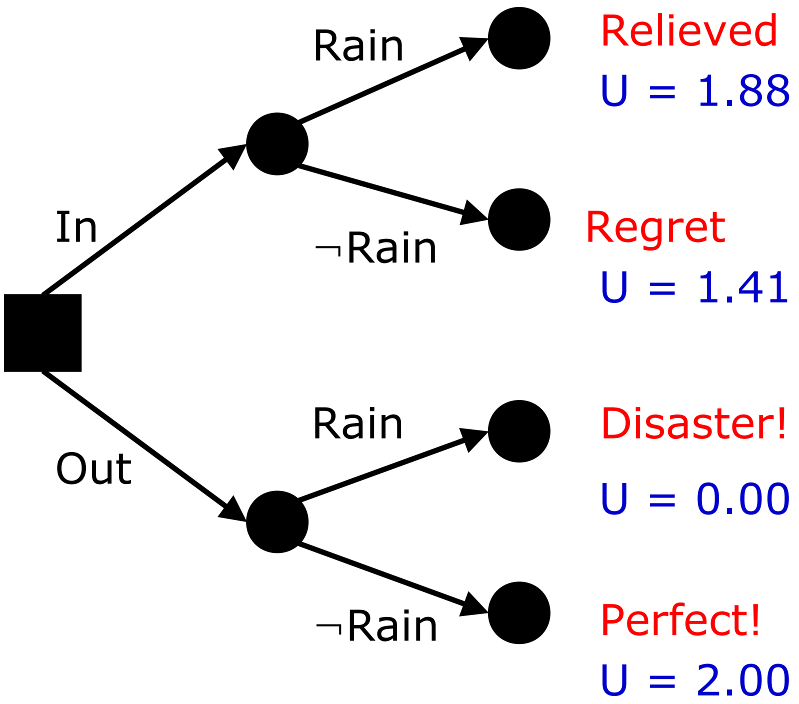

Example adapted from Breese & Kooler 1997

Example: The party problem

Example:
Stockholm, Sweden

$$EU(in) = \frac{18}{30} \times 1.88 + \frac{12}{30} \times 1.41 = 1.7$$

$$EU(out) = \frac{18}{30} \times 0.00 + \frac{12}{30} \times 2.00 = 0.8$$

Be indoors!

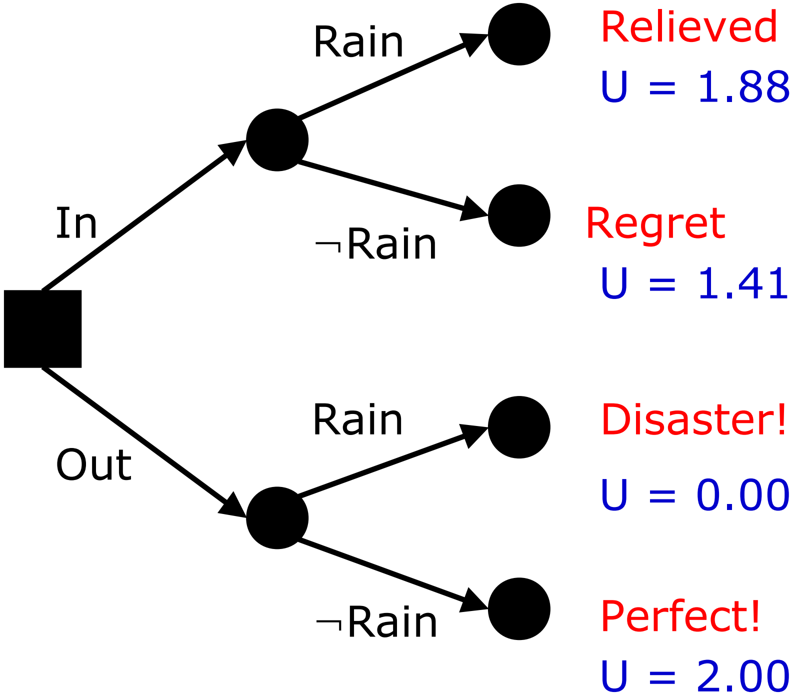

Example adapted from Breese & Kooler 1997

Example: The party problem

Example:
San Fransisco, California

$$EU(in) = \frac{1}{30} \times 1.88 + \frac{29}{30} \times 1.41 = 1.4$$

$$EU(out) = \frac{1}{30} \times 0.00 + \frac{29}{30} \times 2.00 = 1.9$$

Be outdoors!

The change from outdoors to indoors occurs at P(Rain) > 7/30

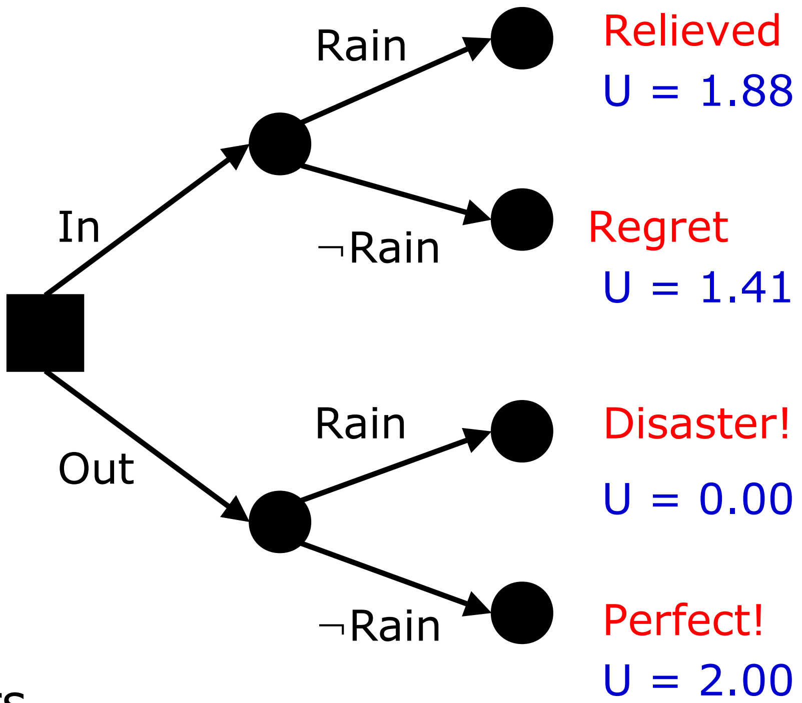

Example adapted from Breese & Kooler 1997

Decision network for the party problem

Decision represented by a rectangle

Chance (random variable) represented by an oval.

Utility function represented by a diamond.

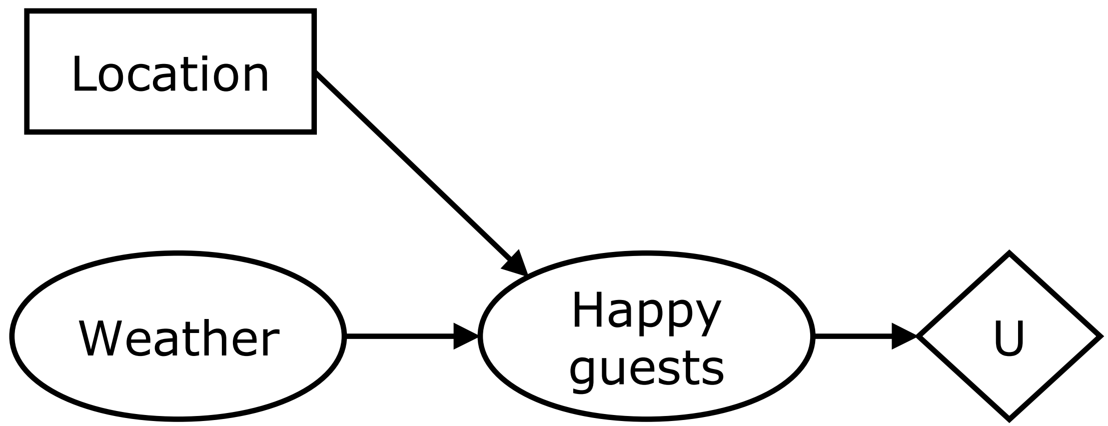

The value of information

- The value of a given piece of information is the difference in expected utility value between best actions before and after information is obtained.

- Information has value to the extent that it is likely to cause a change of plan and to the extent that the new plan will be significantly better than the old plan.