Artificial Intelligence

Utility Theory
for Risk-Sensitive Decision Making

this is not in the Nilsson

Russell and Norvig - Chapter 16.1 - 16.3

Utility Theory (1)

probability 1.0: win 4,500,000 dollars

probability 0.5: win 10,000,000 dollars
probability 0.5: win 0 dollars

Utility Theory (2)

prob	reward	exp reward	utility	exp utility
1.0	4,500,000	4,500,000	0.5	0.6
0.5 0.5	10,000,000 0	5,000,000	0.7 0.0	0.4

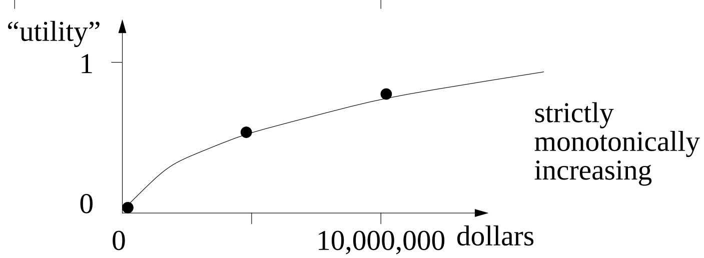

utility function: money -> utility
its existence follows from simple axioms

Utility Theory (3)

[p, A; 1-p, B]	lottery; A occurs with probability p and B occurs with probability 1-p
A > B	the agent prefers A over B
A = B	the agent is indifferent between A and B

axioms

Orderability
(A > B) OR (B > A) OR (A = B)

Transitivity
(A > B) AND (B > C) IMPLIES (A > C)

Continuity
A > B > C IMPLIES THERE EXISTS p SUCH THAT [p, A; 1-p, C] = B

Substitutability
A = B IMPLIES [p, A; 1-p, C] = [p, B; 1-p, C]

Monotonicity
A > B IMPLIES (p >= q EQUIVALENT [p, A; 1-p, B] >= [q, A; 1-q, B])

Decomposability
[p, A; 1-p, [q, B; 1-q, C]] = [p, A; (1-p)q, B; (1-p)(1-q), C]

it follows from these axioms that

Utility Principle
There exists a real-valued function U that operates on states such that
U(A) > U(B) if and only if A is preferred to B, and
U(A) = U(B) if and only if the agent is indifferent between A and B.

Maximum Expected Utility Principle
The utility of a lottery is the sum of the probabilities of each outcome times the utility of that outcome.
With this extension, the utility principle applies to lotteries as well.

Utility Theory (4)

utility function 1:
$u(x)$

lottery 1:
prize x_i with probability p_i

utility function 2:
$u'(x) = m\, u(x) + n$ with $m > 0$
(positively linear transformation)

lottery 2:
prize y_j with probability q_j

utility function 1:
compare $\Sigma_i\, p_i\, u(x_i)$ and $\Sigma_j\, q_j\, u(y_j)$

utility function 2:
compare $\Sigma_i\, p_i\, u'(x_i)$ and $\Sigma_j\, q_j\, u'(y_j)$
compare $\Sigma_i\, p_i\, (m\, u(x_i) + n)$ and $\Sigma_j\, q_j\, (m\, u(y_j) + n)$
compare $m\, (\Sigma_i\, p_i\, u(x_i)) + n$ and $m\, (\Sigma_j\, q_j\, u(y_j)) + n$
compare $\Sigma_i\, p_i\, u(x_i)$ and $\Sigma_j\, q_j\, u(y_j)$

utility functions are defined only up to positively linear transformations

Utility Theory (5)

people have different utility functions
how to elicit them from decision makers?

what is the value of p that makes you indifferent between:

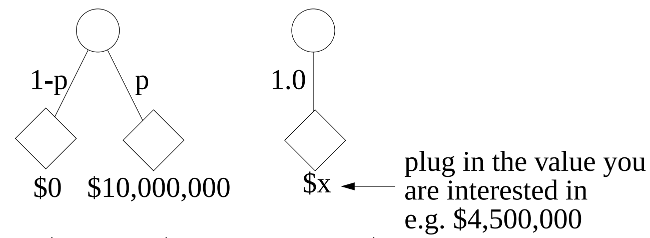

$(1\text{-}p)\ u(\$0) + p\ u(\$10{,}000{,}000) = u(\$x)$
$(1\text{-}p)\ 0 + p\ 1 = u(\$x)$
$p = u(x)$

Utility Theory (6)

normative theory, not descriptive theory

A: probability 0.80: win $4,000; probability 0.20: win $0
B: probability 1.00: win $3,000

C: probability 0.20: win $4,000; probability 0.80: win $0
D: probability 0.25: win $3,000; probability 0.75: win $0

Choosing B over A, and C over D is inconsistent.
Set u($0) = 0 and solve!