Lecture 20: Uncertainty 5

Victor R. Lesser
CMPSCI 683
Fall 2010

Today's Lecture

- Making Simple One-Shot Decisions
  - Combining Beliefs and Desires Under Uncertainty
  - Basis of Utility Theory

Maximum Expected Utility (MEU)

- The MEU principle says that a rational agent should choose an action that maximizes its expected utility in the current state (E)

$EU(E) = \max_A \Sigma_i P(Result_i(A)|Do(A),E)\ U(Result_i(A))$

- Why isn't the MEU principle all we need in order to build "intelligent agents"?
  - Is it Difficult to Compute P,E or U ?

MEU Computational Difficulties

- Knowing the current state of the world requires perception, learning, knowledge representation and inference.
- Computing P(*) requires a complete causal model of the world.
- Computing U(E) often requires search or planning
  - Calculation of Utility of state may require looking at what utilities could be achieved from that state
- All of the above can be computationally intractable, hence one needs to distinguish between "perfect rationality" and "resource-bounded rationality" or "bounded-optimality".
- Also Need to consider more than one action (one-shot decisions versus sequential decisions).

Still, decision theory offers a good framework

The Foundation of Utility Theory

- Why make decisions based on average or expected utility?
- Why can one assume that utility functions exist?
- Can an agent act rationally by expressing preferences between states without giving them numeric values?
- Can every preference structure be captured by assigning a single number to every state?

Constraints on Rational Preferences

The MEU principle can be derived from a more basic set of assumptions.

- Lotteries: a probability distribution over actual outcomes
  - Key to formalizing preference structures and relating them to MEU
- Different outcomes correspond to different prizes.
  - L = [p;A; 1-p,B].
- Can have any number of outcomes, an outcome of a lottery can be another lottery.
  - L = [p_1;C_1; p_2;C_2 ; …. p_n;C_n].
  - L = [p;A; 1-p [p_1;C_1; p_2;C_2 ; …. p_n;C_n]].
- A lottery with only one outcome written as [1,A] or simply A.

Preference Notation

Let A and B be two possible outcomes:

A > B — Outcome A is preferred to B

A ≡ B — The agent is indifferent between A and B

A ≥ B — The agent prefers A to B or is indifferent between them.

Axioms of Utility Theory

- Orderability (the agent know what it wants)

  $(A > B) \vee (B > A) \vee (A \equiv B)$
- Transitivity

  $(A > B) \wedge (B > C) \Rightarrow (A > C)$
- Continuity

  $A > B > C \Rightarrow \exists p\ [p,A;\ 1\text{-}p,C] \equiv B$
- Substitutability

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

Axioms of Utility Theory cont.

- Monotonicity

$A \geq B \Rightarrow (p \geq q \Leftrightarrow [p,A;\ 1-p,B] \geq [q,A;\ 1-q,B])$

- Decomposability

$[p,A;\ 1-p,[q,B;\ 1-q,C]] \equiv$
$[p,A;\ (1-p)q,B;\ (1-p)(1-q),C]$

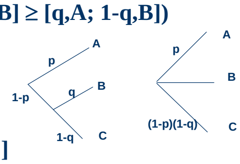

The Utility Principle

Theorem: If an agent's preferences obey the axioms of utility theory, then there exists a real-valued function U that operates on states such that:

$$U(A) > U(B) \Leftrightarrow A > B;$$

and

$$U(A) = U(B) \Leftrightarrow A \equiv B$$

Utility function *follows from* axioms of utility

Maximum Expected Utility Principle

Theorem: The utility of a lottery is the sum of probabilities of each outcome times the utility of that outcome:

$U([p_1,S_1;\ p_2,S_2;\ \ldots;\ p_n,S_n]) = \Sigma_i\ p_i\ U(S_i)$

Expected Monetary Value (EMV)

Example: You can take a $1,000,000 prize or gamble on it by flipping a coin. If you gamble, you will either triple the prize or lose it.

EMV (expected monetary value) of the lottery is $1,500,000, *but does it have higher utility*?

- The utility-theoretic way of thinking about it

Suppose in your current state you have wealth k

EU(*accept*) = ½ $U(S_k)$ + ½ $U(S_{k+3M})$

EU(*decline*) = $U(S_{k+1M})$

Best decision depends on the utils of these 3 states

If $U(S_k)=5$, $U(S_{k+1M})=8$, $U(S_{k+3M})=10$, then?

Expected Monetary Value (EMV)

Bernoulli's 1738 St. Petersburg Paradox: Toss a coin until it comes up heads. If it happens after n times, you receive 2^n dollars.

EMV(St. P.) = $\Sigma_i\ 1/(2^i)\ 2^i$ = inf.

How much should you pay to participate in this game?

Risk-Averter's Curve

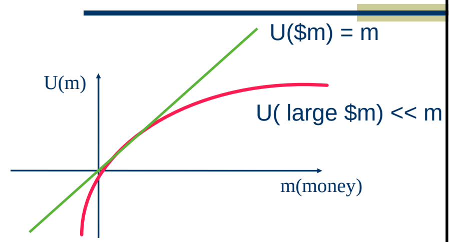

Decreasing marginal utility for money. Will buy affordable insurance. Will only take gambles with substantial positive expected monetary payoff.

Risk-Seeker's Curve

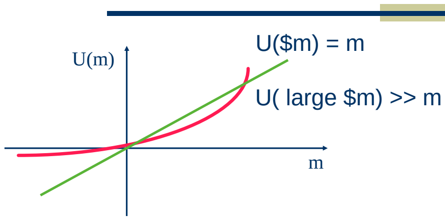

Increasing marginal utility for money. Will not buy insurance. Will sometimes participate in unfavorable gamble having negative expected monetary payoff if there is a chance for high payoff if successful.

Utility Curves

- Risk-neutral agents (linear curve).
- Regardless of the attitude towards risk, the utility function can always be approximated by a straight line over a small range of monetary outcome.
- The certainty equivalent of a lottery.
  - Example: Most people will accept about $400 in lieu of a gamble that gives $1000 half the time and $0 the other half.

Human Judgment under Uncertainty

- Is decision theory compatible with human judgment under uncertainty?
- Does it outperform human judgment in micro/macro worlds?
- Are people "experts" in reasoning under uncertainty? How well do they perform? What kind of heuristics do they use?

Is Human Judgment Rational?

- Choose between lotteries A and B, and then between C and D:

A: 80% chance of $4000 C: 20% chance of $4000
B: 100% chance of $3000 D: 25% chance of $3000

- The majority of the subjects choose B over A and C over D. But if U($m) = m, we get:
  0.8 U($4000) < U($3000) and
  0.2 U($4000) > 0.25 U($3000)
  ...contradicts the axioms.

[.8,4000,.2,0]<[1,3000,.0,0], [.25,3000,.75,0]>[.2,4000,.8,0]

Issue of utility function does not factor in probability of outcome

Utility Scales and Utility Assessment

- Utility functions are not unique (for a given preference structure): U'(S) = a + b U(S)
- Normalized utility:
  - U– = 0 = Utility(worst possible catastrophe)
  - U+ = 1 = Utility(best possible prize)
- Can find the utility of a state S by adjusting the probability p of a standard lottery: [p,U–; 1-p,U+] that makes the agent indifferent between S and the lottery.

Multi-Attribute Utility Functions

- Why multi-attribute?
  - Example: evaluating a new job offer (salary, commute time, quality of life, etc.)
  - U(a,b,c,…)= f[f1(a),f2(b)…..] where f is a simple function such as addition
    - f=+, In case of mutual preference independence which occurs when it is always preferable to increase the value of an attribute given all other attributes are fixed
- Dominance (strict dominance vs. stochastic dominance).
  - For every point
  - Probablistic view

Strict Dominance

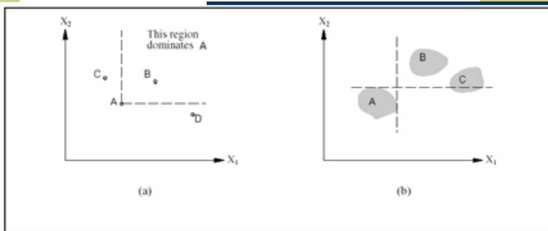

Figure 16.3 Strict dominance. (a) Deterministic: Option A is strictly dominated by B but not by C or D. (b) Uncertain: A is strictly dominated by B but not by C.

Strict dominance occurs if an option is of lower value on all attributes than some other option

Stochastic Dominance

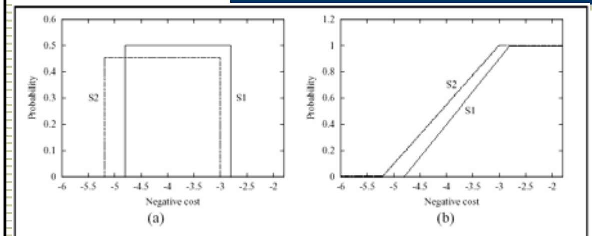

Figure 16.4 Stochastic dominance. (a) S_1 stochastically dominates S_2 on cost. (b) Cumulative distributions for the negative cost of S_1 and S_2.

The value of information

- Example 1: You consider buying a program to manage your finances that costs $100. There is a prior probability of 0.7 that the program is suitable in which case it will have a positive effect on your work worth $500. There is a probability of 0.3 that the program is not suitable in which case it will have no effect.

- What is the value of knowing whether the program is suitable before buying it?

Example 1 Answer

- Expected utility given information
  - [0.7*(500-100)+0.3(0)]
    - Why not 1.0 *(500-100)
- Expected utility not given information
  - [0.7(500-100)+0.3(0-100)]
- Value of Information
  - [0.7*(500-100)+0.3(0)] - [0.7(500-100)+0.3(0-100)] = 280 - 250 = $30

The Value of Information – Example 2

Example 2: Suppose an oil company is hoping to buy one of n blocks of ocean drilling rights.

- Exactly one block contains oil worth C dollars.
- The price of each block is C/n dollars.
- If the company is risk-neutral, it will be indifferent between buying a block or not.-- WHY?
- A seismologist offers the company a survey indicating whether block #3 contains oil.
- How much should the company be willing to pay for the information?

The Value of Information cont*.

- What can the company do with the information?
- Case 1: block #3 contains oil (p=1/n).
  Company will buy it and make a profit of:
  C - C/n = (n-1) C/n dollars.
- Case 2: block #3 contains no oil (p=(n-1)/n).
  Company will buy different block and make:
  C/(n-1) - C/n = C/(n (n-1)) dollars.
- Now, the overall expected profit is C/n.
  (1/n)((n-1) C/n) + ((n-1)/n)(C/(n (n-1))
- What is the value of information?

Value of Perfect Information

- The general case: We assume that exact evidence can be obtained about the value of some random variable E_j.
- The agent's current knowledge is E.
- The value of the current best action α is defined by:

$$EU(\alpha|E) = \max_A \Sigma_i P(Result_i(A)|Do(A),E)\ U(Result_i(A))$$

VPI cont.

- With the information, the value of the new best action will be: $EU(\alpha_{Ej}|E,E_j) = \max_A \Sigma_i P(Result_i(A) \mid Do(A),E,E_j)\ U(Result_i(A))$
- But E_j is a random variable whose value is currently unknown, so we must average over all possible values e_{jk} using our current belief:

$$VPI_E(E_j) = (\Sigma_k P(E_j=e_{jk} \mid E)\ EU(\alpha_{ejk} \mid E, E_j = e_{jk})) - EU(\alpha \mid E)$$

Properties of the Value of Information*

- In general:
  $VPI_E(E_j,E_k) \neq VPI_E(E_j) + VPI_E(E_k)$
- But the order is not important:
  $VPI_E(E_j,E_k) =$
  $VPI_E(E_j) + VPI_{E,E_j}(E_k) =$
  $VPI_E(E_k) + VPI_{E,E_k}(E_j)$
- What about the value of imperfect information?

Value of Information

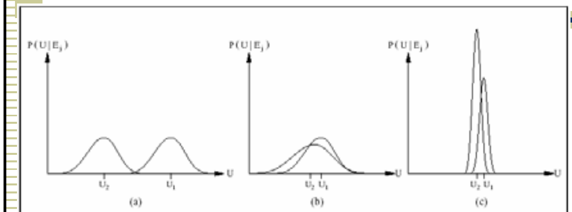

Figure 16.7 Three generic cases for the value of information. In (a), A_1 will almost certainly remain superior to A_2, so the information is not needed. In (b), the choice is unclear and the information is crucial. In (c), the choice is unclear but because it makes little difference, the information is less valuable.

Utility Distributions for Actions A_1 and A_2 over the range of the random variable E_j. The question is whether knowing more about E_j helps?

Next Lecture

- Decision Trees and Networks