Making Simple Decisions: Utility Theory

CMPSCI 383
Nov 3, 2011

Today's topics

- Maximum expected utility principle
- Utility and preferences: rational preferences
- Assessing utility
- Utility of money
  - Risk averse
  - Risk seeking
- What do humans do?
- Multiattribute utility functions
  - Pure dominance
  - Preference structure
- Decision networks

Maximum Expected Utility Principle

- *U(s)*: utility of state *s*
- Expected utility of an action:

Evidence observations

$$EU(a \mid e) = \sum_{s'} P(\text{Result}(a) = s' \mid a, e)\, U(s')$$

$$\text{action} = \arg\max_{a} EU(a \mid e)$$

- MEU principle: formalizes the idea of "doing the right thing"

Why <u>Expected</u> Utility?

- Why not sum of cubes?
- Or minimize worst-case loss?
- Why not just deal with preferences?
- Why should a utility function with the desired properties exist at all?

Preferences

An agent chooses among prizes (A, B, etc.) and lotteries, i.e., situations with uncertain prizes

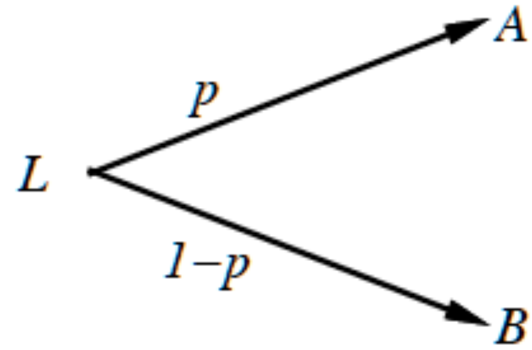

Lottery $L = [p, A; (1-p), B]$

Notation:

$A \succ B$	A preferred to B
$A \sim B$	indifference between A and B
$A \succsim B$	B not preferred to A (prefers A over B or is indifferent between them)

Rational Preferences

Idea: preferences of a rational agent must obey constraints.
Rational preferences $\Rightarrow$
behavior describable as maximization of expected utility

Constraints:

<u>Orderability</u>

$$(A \succ B) \vee (B \succ A) \vee (A \sim B)$$

<u>Transitivity</u>

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

<u>Continuity</u>

$$A \succ B \succ C \Rightarrow \exists p \; [p, A;\; 1-p, C] \sim B$$

<u>Substitutability</u>

$$A \sim B \Rightarrow [p, A;\; 1-p, C] \sim [p, B; 1-p, C]$$

<u>Monotonicity</u>

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

Rational Preferences contd.

Violating the constraints leads to self-evident irrationality

For example: an agent with intransitive preferences can be induced to give away all its money

If $B \succ C$, then an agent who has C would pay (say) 1 cent to get B

If $A \succ B$, then an agent who has B would pay (say) 1 cent to get A

If $C \succ A$, then an agent who has A would pay (say) 1 cent to get C

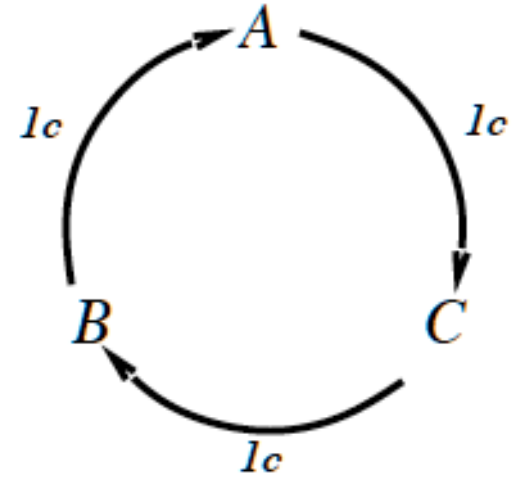

Maximizing Expected Utility

Theorem (Ramsey, 1931; von Neumann and Morgenstern, 1944):
Given preferences satisfying the constraints
there exists a real-valued function U such that

$$U(A) \geq U(B) \quad \Leftrightarrow \quad A \succsim B$$

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

MEU principle:

Choose the action that maximizes expected utility

Note: an agent can be entirely rational (consistent with MEU)
without ever representing or manipulating utilities and probabilities

E.g., a lookup table for perfect tictactoe

Utility function is not unique: $aU(S)+b, \ \ a>0$

Utility assessment

Utilities map states to real numbers. Which numbers?

Standard approach to assessment of human utilities:

compare a given state A to a standard lottery L_p that has

"best possible prize" $u_\top$ with probability p

"worst possible catastrophe" $u_\bot$ with probability $(1-p)$

adjust lottery probability p until $A \sim L_p$

pay $30 ~

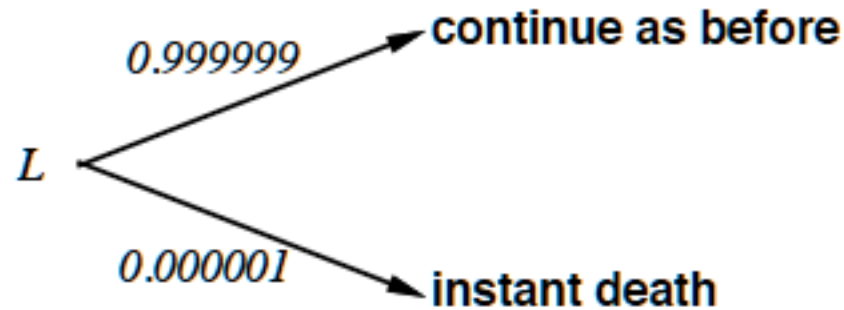

Utility Scale

Normalized utilities: $u_\top = 1.0$, $u_\bot = 0.0$

Micromorts: one-millionth chance of death
useful for Russian roulette, paying to reduce product risks, etc.

QALYs: quality-adjusted life years
useful for medical decisions involving substantial risk

Note: behavior is invariant w.r.t. +ve linear transformation

$$U'(x) = k_1 U(x) + k_2 \quad \text{where } k_1 > 0$$

With deterministic prizes only (no lottery choices), only ordinal utility can be determined, i.e., total order on prizes

Money

- You’re on a game show. You have won a $1M prize. You are offered this gamble:
  - Flip a coin: heads→ nothing; tails $2.5M
- What is the expected monetary value (EMV) of the gamble?
- .5($0) + .5($2.5M) = $1.25M
- What do you do?

- Utility is not directly proportional to EMV
- Risk-averse

Money

- You are $10M in debt.
- You are offered the following gamble:
  - Flip a coin: heads → $10M; tails → –$20M
- What is the EMV?
  - .5($10M) + .5(–$20M) = –$5M
- What do you do?

- Risk-prone or risk-seeking

Money

Money does not behave as a utility function

Given a lottery L with expected monetary value $EMV(L)$,
usually $U(L) < U(EMV(L))$, i.e., people are risk-averse

Typical empirical data, extrapolated with risk-prone behavior:

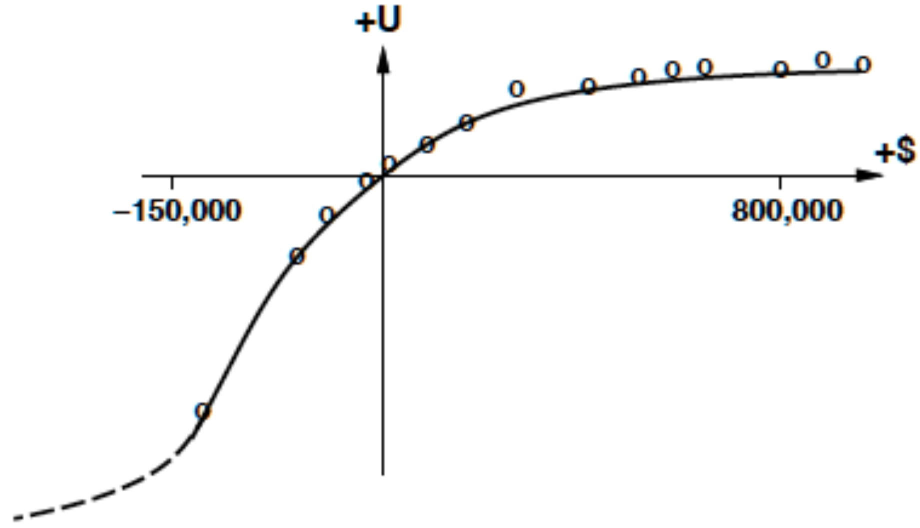

Utility curves

CE: Certainty equivalent: value accepted instead of an uncertain payment

E(U(W)): exp. value of utility of uncertain payment; E(W): exp. value of uncertain payment; U(CE): utility of certainty equivalent; U(E(W)): utility of exp. value of uncertain payment; $U(W_0)$: utility of minimal payment; W_1: maximal payment; RP: risk premium

risk neutral

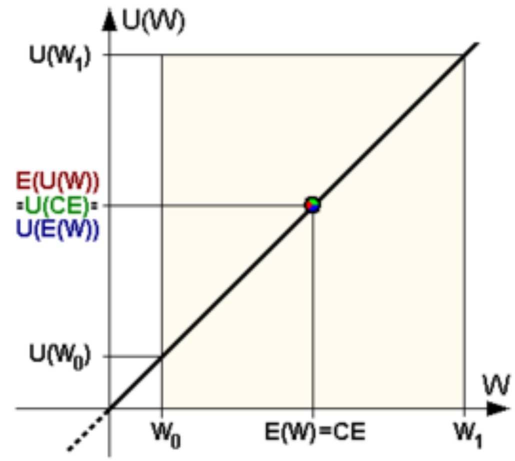

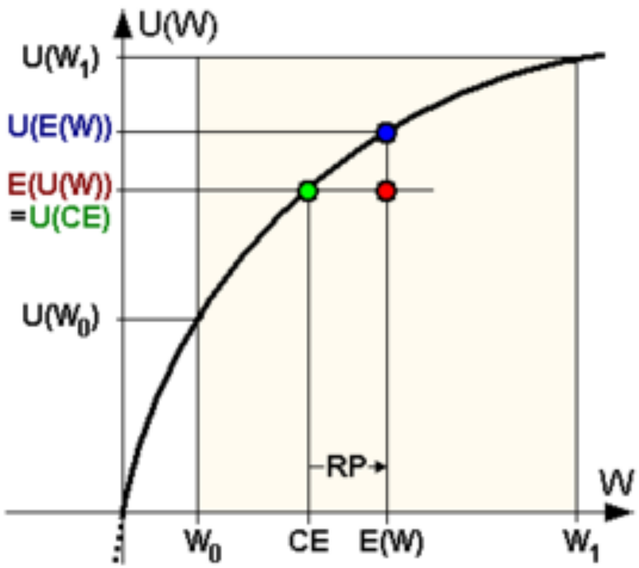

risk averse

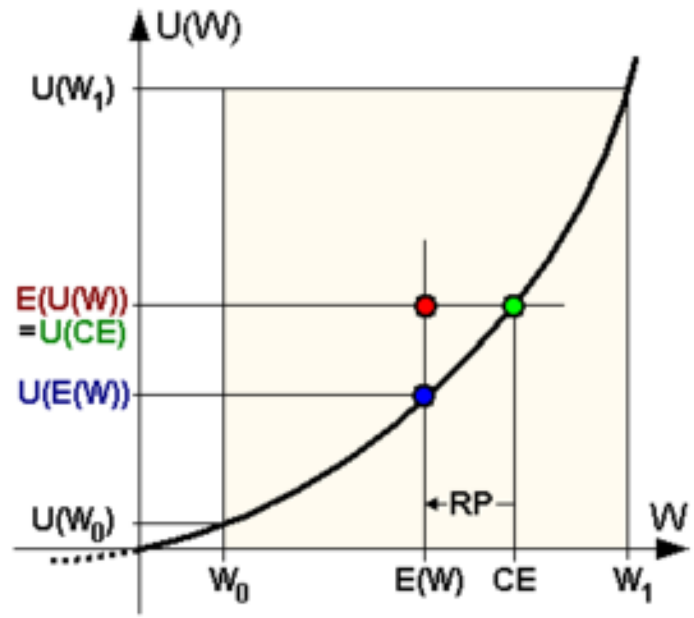

risk seeking

Student Group Utility

For each x, adjust p until half the class votes for lottery (M=10,000)

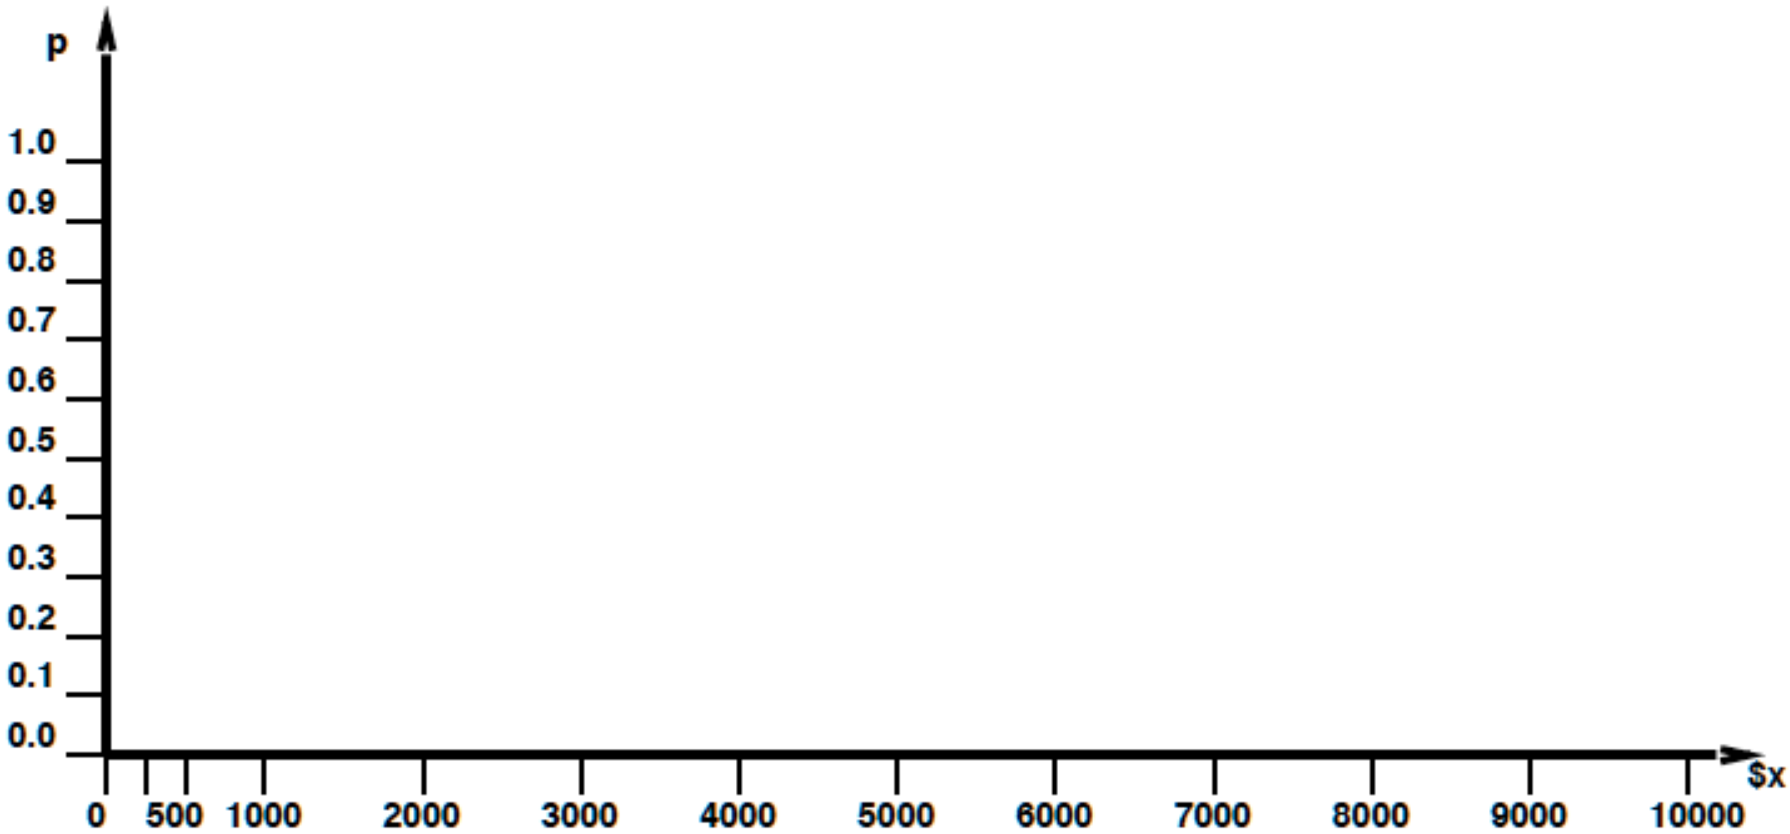

Humans?

- Normative theory: how a rational agent should act
- Descriptive theory: how actual agents—e.g. humans—actually act.

Allais paradox

- Which of these lotteries do you prefer?
  - A: 80% chance of $4000
  - B: 100% chance of $3000
- What about this lottery:
  - C: 20% chance of $4000
  - D: 25% chance of $3000

- Most prefer B over A and D over C
- Is this “rational”?

Allais paradox contd.

$$A \succ B \Rightarrow U(\$3000) > 0.8U(\$4000)$$

$$C \succ D \Rightarrow 0.2U(\$4000) > 0.25U(\$3000) \Rightarrow 0.8U(\$4000) > U(\$3000)$$

No utility function is consistent with these choices!

- Why might this be?
  - Strong attraction to certain gains.
    - Too much computation?
    - Distrust?
    - Too much regret if you lose? (*feeling like an idiot is worse than getting no money..*)

Ellsberg paradox

- You are told an urn contains 1/3 red balls, and 2/3 either black or yellow balls
- Which of these lotteries do you prefer?
  - A: $100 for a red ball
  - B: $100 for a black ball
- What about this lottery:
  - C: $100 for a red or yellow ball
  - D: $100 for a black or yellow ball

- Most prefer A over B and D over C
- Is this “rational”?

Ellsberg paradox contd.

- You are told an urn contains 1/3 red balls, and 2/3 either black or yellow balls
- Which of these lotteries do you prefer?
  - A: $100 for a red ball
  - B: $100 for a black ball
- What about this lottery:
  - C: $100 for a red or yellow ball
  - D: $100 for a black or yellow ball

Most prefer A over B and D over C

- If you think there are more red than black balls, pick A over B and C over D.
- If you think there are fewer red than black balls, pick B over A and D over C.
- Why might people make the usual choice?
  - Ambiguity aversion: 1/3 chance (A) vs. anywhere between 0 and 2/3 chance (B); or 2/3 chance (D) vs. anywhere between 1/3 and 3/3 chance (C)

Framing Effect

- Exact wording can have a big impact
- Which looks better?
  - A medical procedure with a 90% survival rate?
  - A medical procedure with a 10% death rate?

Anchoring Effect

- Anchoring occurs when individuals overly rely on a specific piece of information to govern their thought-process. Once the anchor is set, there is a bias toward adjusting or interpreting other information to reflect the "anchored" information.
- In one of the first studies, Tversky and Kahneman showed that when asked to guess the percentage of African nations which are members of the United Nations, people who were first asked "Was it more or less than 10%?" guessed lower values (25% on average) than those who had been asked if it was more or less than 65% (45% on average).

Evolutionary Psychology

- Briefly: Utility theory does not account for the complexity of the problems that animals faced through evolutionary history.

Behavioral Economics

- Uses social, cognitive and emotional factors in understanding the economic decisions of people and institutions performing economic functions.

Neuroeconomics

- Seeks to explain human decision making by studying the brain. It combines research methods from neuroscience, experimental and behavioral economics, and cognitive and social psychology.

Multiattribute Utility Functions

How can we handle utility functions of many variables $X_1 \ldots X_n$?
E.g., what is $U(Deaths, Noise, Cost)$?

- Assume each variable has a numerical value
- Assume higher value corresponds to higher utility

Strict Dominance

Typically define attributes such that U is monotonic in each

Strict dominance: choice B strictly dominates choice A iff

$$\forall i \;\; X_i(B) \geq X_i(A) \quad \text{(and hence } U(B) \geq U(A)\text{)}$$

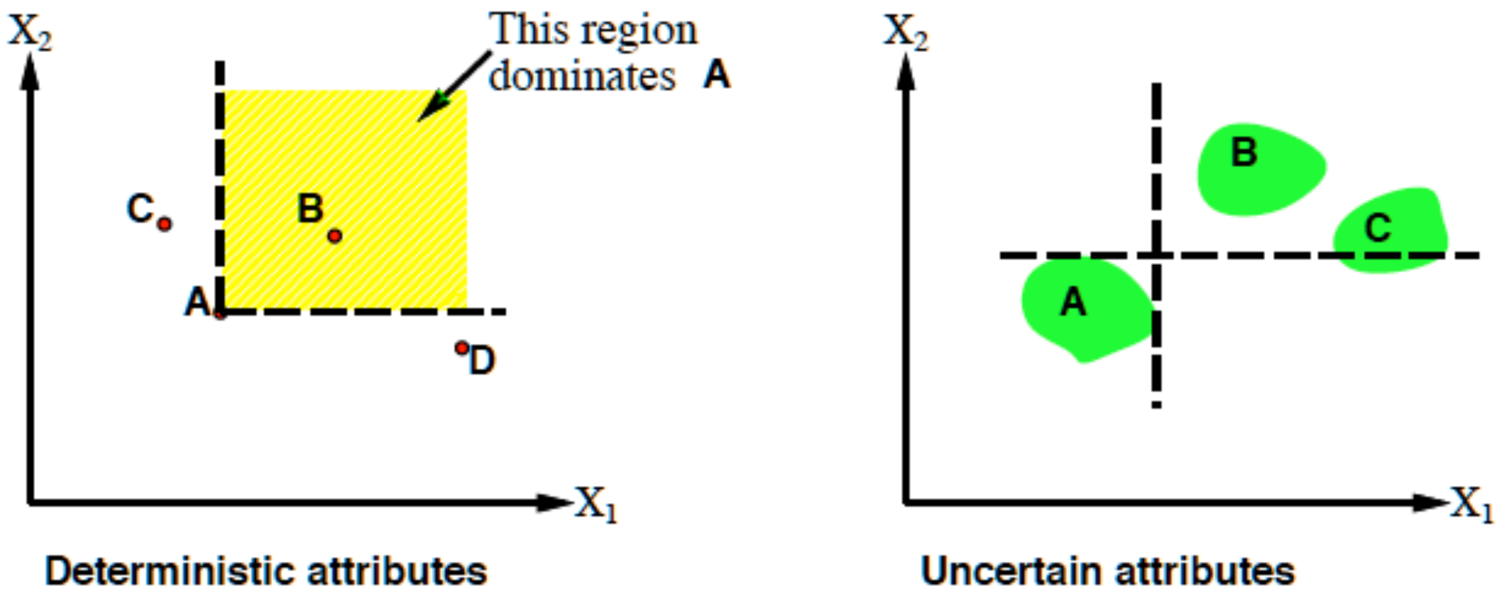

Strict dominance seldom holds in practice

Stochastic Dominance

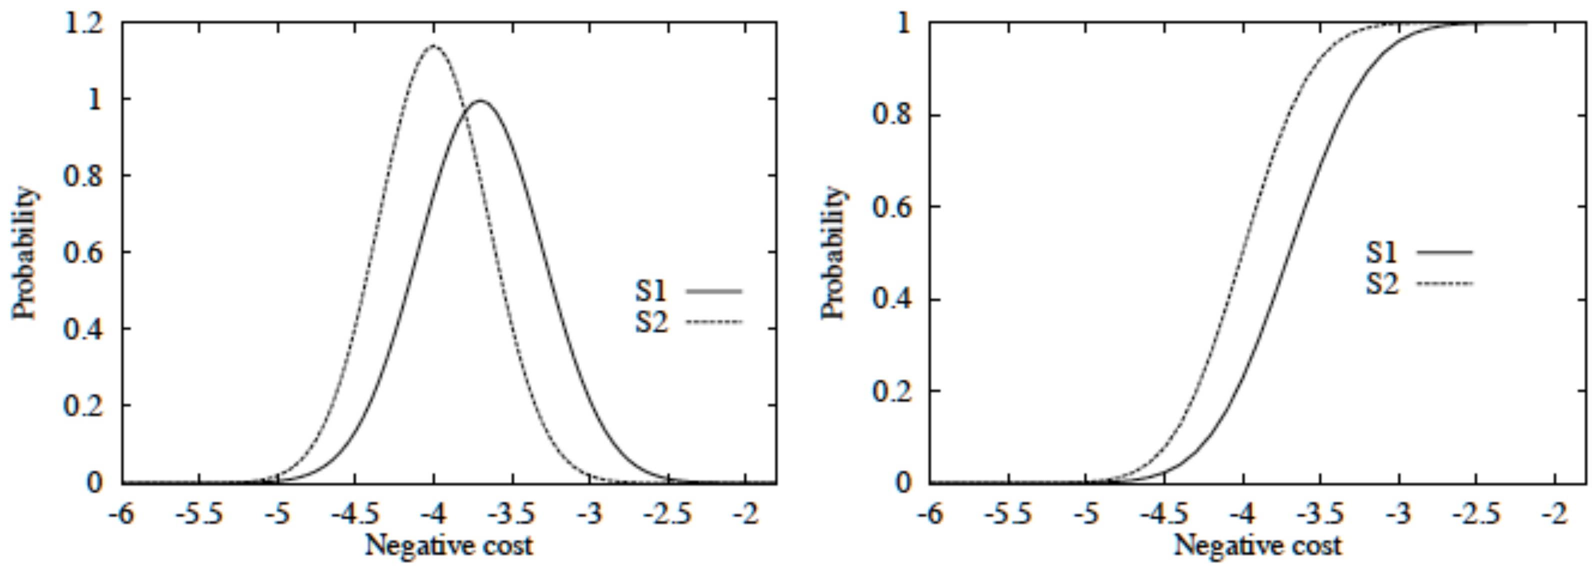

Distribution p_1 stochastically dominates distribution p_2 iff

$$\forall\, t \quad \int_{-\infty}^{t} p_1(x)dx \leq \int_{-\infty}^{t} p_2(t)dt$$

If U is monotonic in x, then A_1 with outcome distribution p_1 stochastically dominates A_2 with outcome distribution p_2:

$$\int_{-\infty}^{\infty} p_1(x)U(x)dx \geq \int_{-\infty}^{\infty} p_2(x)U(x)dx$$

Multiattribute case: stochastic dominance on all attributes $\Rightarrow$ optimal

Stochastic Dominance contd.

Stochastic dominance can often be determined without exact distributions using qualitative reasoning

E.g., construction cost increases with distance from city
S_1 is closer to the city than S_2
$\Rightarrow$ S_1 stochastically dominates S_2 on cost

E.g., injury increases with collision speed

Preference Structure: deterministic

X_1 and X_2 preferentially independent of X_3 iff
preference between $\langle x_1, x_2, x_3 \rangle$ and $\langle x'_1, x'_2, x_3 \rangle$
does not depend on x_3

E.g., $\langle Noise, Cost, Safety \rangle$:
⟨20,000 suffer, \$4.6 billion, 0.06 deaths/mpm⟩ vs.
⟨70,000 suffer, \$4.2 billion, 0.06 deaths/mpm⟩

Theorem (Leontief, 1947): if every pair of attributes is P.I. of its complement, then every subset of attributes is P.I of its complement: mutual P.I..

Theorem (Debreu, 1960): mutual P.I. $\Rightarrow$ $\exists$ additive value function:

$$V(S) = \Sigma_i V_i(X_i(S))$$

Hence assess n single-attribute functions; often a good approximation

Preference Structure: stochastic

Here is just one main result to illustrate

Need to consider preferences over lotteries:

X, Y: sets of attributes

$\mathbf{X}$ is utility-independent of $\mathbf{Y}$ iff

preferences over lotteries in $\mathbf{X}$ do not depend on $\mathbf{y}$

Mutual U.I.: each subset is U.I of its complement

$\Rightarrow$ $\exists$ multiplicative utility function:

For case of just 3 attributes

$$U = k_1U_1 + k_2U_2 + k_3U_3 \\ + k_1k_2U_1U_2 + k_2k_3U_2U_3 + k_3k_1U_3U_1 \\ + k_1k_2k_3U_1U_2U_3$$

Routine procedures and software packages for generating preference tests to identify various canonical families of utility functions

Decision Networks

Add action nodes and utility nodes to belief networks
to enable rational decision making

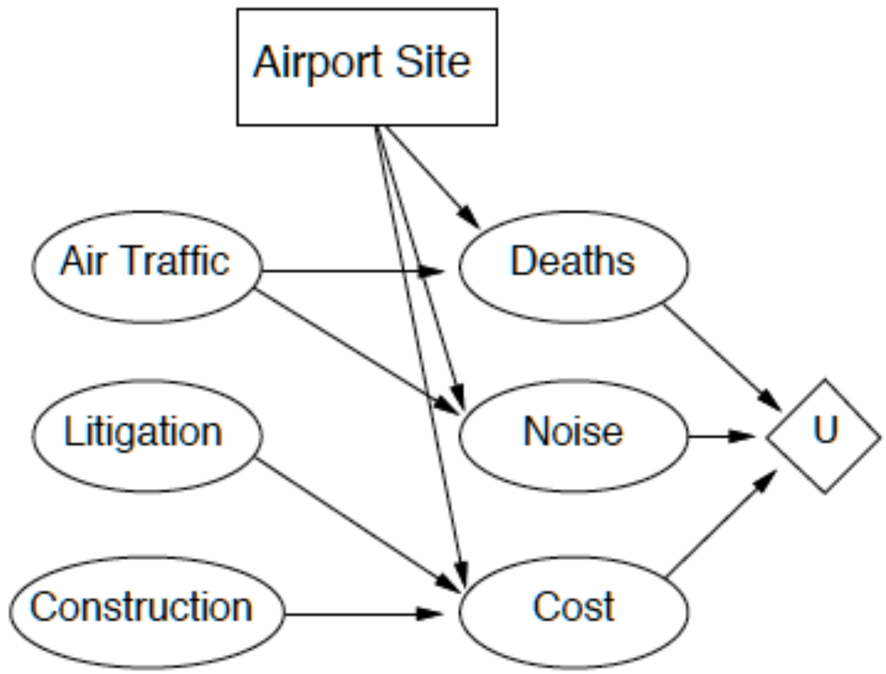

Algorithm:

For each value of action node
    compute expected value of utility node given action, evidence
Return MEU action

Summary

- Maximum expected utility principle
- Utility and preferences: rational preferences
- Assessing utility
- Utility of money
  - Risk averse
  - Risk seeking
- What do humans do?
- Multiattribute utility functions
  - Pure dominance
  - Preference structure
- Decision networks

Next Class

- Making decisions over time
- Secs. 17.1 – 17.3
- Phil will be lecturing….