

430.457

Introduction to Intelligent Systems

Prof. Songhwai Oh
ECE, SNU

Ch. 15 Making Simple Decisions

UTILITY THEORY

Principle of Maximum Expected Utility

- **Utility function** $U(s)$: assigns a single number to express the desirability of a state.
- **Expected utility** $EU(a|\mathbf{e})$

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

- **Principle of maximum expected utility**: a rational agent should choose the action that maximizes the agent's expected utility – *do the right thing*.

$$action = \arg \max_a EU(a|\mathbf{e})$$

I.e., if an agent acts so as to maximize a utility function that correctly reflects the performance measure, then the agent will achieve the highest possible performance score (averaged over all the possible environments).

Notations

Preference relations:

- $A \succ B$, the agent prefers A over B .
- $A \sim B$, the agent is indifferent between A and B .
- $A \succeq B$, the agent prefers A over B or is indifferent between them.

Lottery: $L = [p_1, S_1; p_2, S_2; \dots; p_n, S_n]$


- $S_1, \dots, S_n$, possible outcomes
- $p_1, \dots, p_n$, probabilities of outcomes
- S_i can be either an atomic state or another lottery.

Axioms of Utility Theory

1. **Orderability:** Exactly one of $(A \succ B)$, $(B \succ A)$, or $(A \sim B)$ holds.
2. **Transitivity:** $(A \succ B) \wedge (B \succ C) \Rightarrow (A \succ C)$.
3. **Continuity:** $A \succ B \succ C \Rightarrow \exists p \ [p, A; 1 - p, C] \sim B$.
4. **Substitutability:**
 - $A \sim B \Rightarrow [p, A; 1 - p, C] \sim [p, B; 1 - p, C]$ for any p .
 - $A \succ B \Rightarrow [p, A; 1 - p, C] \succ [p, B; 1 - p, C]$ for any p .
5. **Monotonicity:** $A \succ B \Rightarrow (p > q \Leftrightarrow [p, A; 1 - p, B] \succ [q, A; 1 - q, B])$.
6. **Decomposability:**

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

For nontransitive preferences $A \succ B \succ C \succ A$, irrational behavior is possible.



Existence of Utility Function

- **Existence of Utility Function:** If an agent's preferences obey the axioms of utility, then there exists a function U 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 .

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

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

- **Expected Utility of a Lottery:** The utility of a lottery is the sum of the probability of each outcome times the utility of that outcome.

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

(von Neumann and Morgenstern, 1944)

- Hence, an agent can act rationally, i.e., consistent with its preferences, by choosing an action that maximizes expected utility.
- Utility is a function that maps from lotteries to real numbers.
- A utility function exists but it is not unique. If $U(S)$ is a utility function, then so is $U'(S) = aU(S) + b$.

Utility Assessment and Scales

- $u_{\top}$, the largest possible value of U .
- $u_{\perp}$, the smallest possible value of U .
- **Normalized utility** if $u_{\perp} = 0$ and $u_{\top} = 1$.
- **Standard lottery** $[p, u_{\top}; (1 - p), u_{\perp}]$
- To assess the utility of outcome S
 - Choose p such that the agent is indifferent between S and the standard lottery.
 - Then the utility of S is p (for normalized utilities).

Utility of Money


- Suppose you have triumphed over the other competitors in a television game show. The host now offers you a choice: either you can take the \$1,000,000 prize or you can gamble it on the flip of a coin. If the coin comes up heads, you end up with nothing, but if it comes up tails, you get \$2,500,000. What would do?

Utility of Money

- Suppose you have triumphed over the other competitors in a television game show. The host now offers you a choice: either you can take the \$1,000,000 prize or you can gamble it on the flip of a coin. If the coin comes up heads, you end up with nothing, but if it comes up tails, you get \$2,500,000. What would do?
- **Expected monetary value** (EMV) = $\frac{1}{2}(\$0) + \frac{1}{2}(\$2,500,000) = \$1,250,000$.
- Let S_n be the state of possessing total wealth $\$n$ and let $\$k$ be your current wealth. Then the expected utilities of two actions are:


$$\begin{aligned}EU(Accept) &= \frac{1}{2}U(S_k) + \frac{1}{2}U(S_{k+2,500,000}) \\EU(Decline) &= U(S_{k+1,000,000})\end{aligned}$$

-
- Utility is not directly proportional to monetary value.
 - Grayson (1960) found that the utility of money was almost exactly proportional to the logarithm of the amount.



Optimizer's Curse

- Often, the true expected utility $EU(a|e)$ is not available and we have to work with estimates $\widehat{EU}(a|e)$.
- Decision problem with k choices, each of which has true estimated utility of 0
- Error in each utility estimate has zero mean and variance of 1
- Since we select the action with the highest utility estimate, we favor the overly optimistic estimates (source of the bias)
- **Optimizer's curse:** tendency for the estimated expected utility of the best choice to be too high (e.g., utility estimate with high variance)
- In order to avoid the curse, get good probability models $P(\widehat{EU}|EU)$ and $P(EU)$.



What do you prefer?

- A: 80% chance of \$4000
- B: 100% chance of \$3000

What do you prefer?

- C: 20% chance of \$4000
- D: 25% chance of \$3000

Allais Paradox

What do you prefer?

- A: 80% chance of \$4000
- B: 100% chance of \$3000


What do you prefer?

- C: 20% chance of \$4000
- D: 25% chance of \$3000

- Most people consistently prefer B over A (taking the sure thing), and C over D (taking the higher EMV).
- $B \succ A$ implies $U(\$3000) > 0.8U(\$4000)$.
- $C \succ D$ implies $0.8U(\$4000) > U(\$3000)$.
- It is known as the Allais paradox (Allais, 1953).
- **Certainty effect:** people are strongly attracted to gains that are certain
- Reasons: Don't want to compute. Distrust. Regret...

Human Irrationality

- Certainty effect
- Ambiguity aversion
 - Most people elect the known probability rather than the unknown unknowns.
- Framing effect
 - “90% survival rate” vs. “10% death rate”
- Anchoring effect
 - It seems 100,000 KW wine is cheap from the menu with wines listed as 500,000 KW.
- Evolutionary psychology
- If you are interested in topics like these, read **Stumbling on Happiness** by Daniel Gilbert (2007).




Multiple Attributes


- Often outcomes are characterized by many attributes.
- Attributes $\mathbf{X} = X_1, \dots, X_n$
- Example: New airport construction requires following considerations:
 - the disruption caused by construction
 - the cost of land
 - the distance from centers of population
 - the noise of flight operations
 - safety issues arising from local topography and weather conditions, etc.
- **Multiattribute utility function** is required.

Dominance


Strict dominance of S_1 over S_2



Stochastic dominance of S_1 over S_2



(a)



(b)

Stochastic Dominance

- If two actions A_1 and A_2 lead to probability distributions $p_1(x)$ and $p_2(x)$ on attribute X , then A_1 stochastically dominates A_2 on X if

$$\forall x \quad \int_{-\infty}^x p_1(y)dy \leq \int_{-\infty}^x p_2(y)dy.$$

- If A_1 stochastically dominates A_2 , then for any monotonically nondecreasing utility function $U(x)$, the expected utility of A_1 is at least as high as the expected utility of A_2 .
- Hence, if an action is stochastically dominated by another action on all attributes, then it can be discarded.

Multiattribute Utility Function

- **Preference independence:** Two attributes X_1 and X_2 are preferentially independent of a third attribute X_3 if the preference between outcomes $\langle x_1, x_2, x_3 \rangle$ and $\langle x'_1, x'_2, x_3 \rangle$ does not depend on the particular value x_3 for attribute X_3 .
- **Mutual preferential independence (MPI):** Each attribute does not affect the way in which one trades off the other attributes against each other.
- If attributes $X_1, \dots, X_n$ are mutually preferentially independent, then the agents preference behavior can be described as maximizing the function

$$V(x_1, \dots, x_n) = \sum_i V_i(x_i),$$

where each V_i is a value function referring only to the attribute X_i .

- **Utility independence:** A set of attributes $\mathbf{X}$ is utility independent of a set of attributes $\mathbf{Y}$ if preferences between lotteries on the attributes in $\mathbf{X}$ are independent of the particular values of the attributes in $\mathbf{Y}$.
- **Mutual utility independence (MUI):** if each of its subsets is utility-independent of the remaining attributes.
- MUI implies that the agents behavior can be described using a **multiplicative utility function** (Keeney, 1974).