CS 2710 Foundations of AI
Lecture 20-a

Utility theory

Milos Hauskrecht
milos@cs.pitt.edu
5329 Sennott Square

Selection based on expected values

- **Until now:** The optimal action choice was the option that maximized the expected monetary value.
- **But is the expected monetary value always the quantity we want to optimize?**

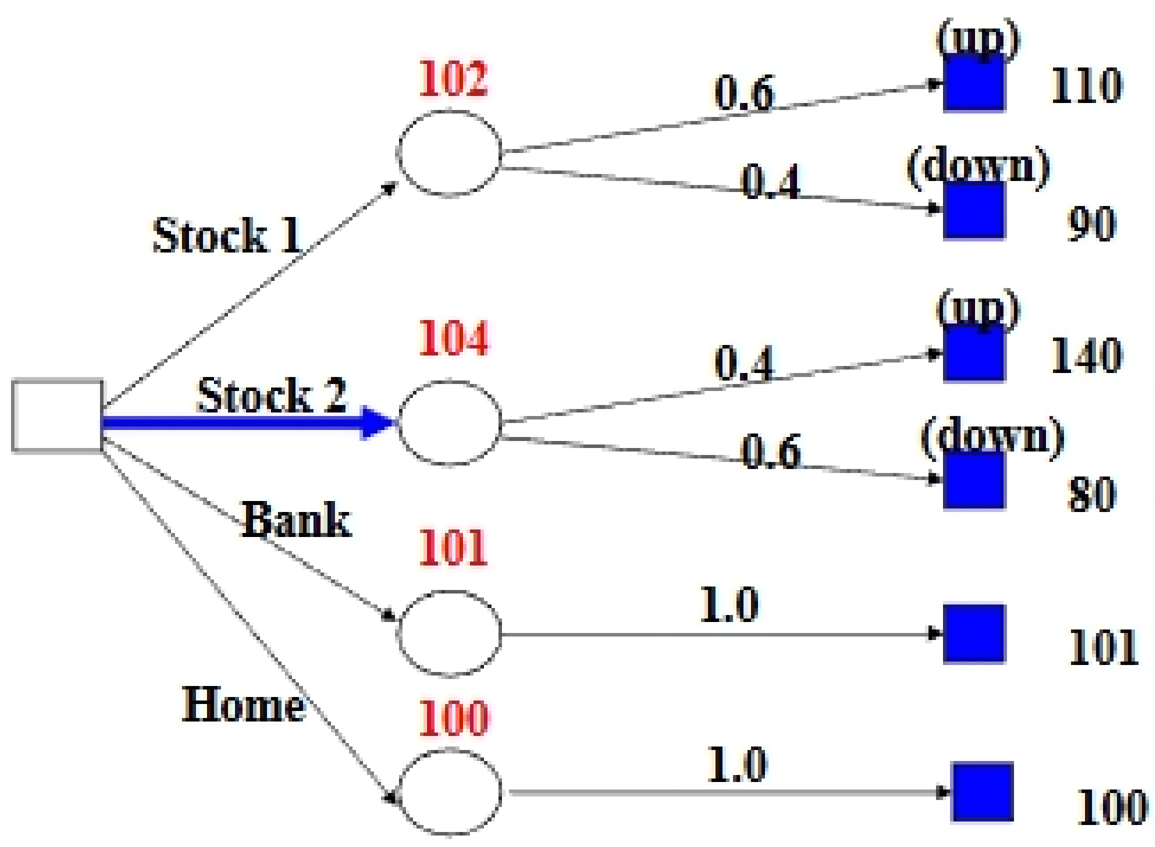

Selection based on expected values

- **Is the expected monetary value always the quantity we want to optimize?**
- **Answer:** Yes, but only if we are risk-neutral.

- But what if **we do not like the risk (we are risk-averse)?**
- In that case we may want to get the premium for undertaking the risk (of loosing the money)
- **Example:**
  - we may prefer to get \$101 for sure against \$102 in expectation but with the risk of loosing the money
- **Problem:** How to model decisions and account for the risk?
- **Solution:** use **utility function, and utility theory**

Utility function

- **Utility function (denoted U)**
  - Quantifies how we "value" outcomes, i.e., it reflects our preferences
  - Can be also applied to "value" outcomes other than money and gains (e.g. utility of a patient being healthy, or ill)
- **Decision making:**
  - uses expected utilities (denoted EU)

$$EU(X) = \sum_{x \in \Omega_x} P(X = x) U(X = x)$$

$U(X = x)$ the utility of outcome x

Important !!!

- Under some conditions on preferences **we can always design the utility function that fits our preferences**

Utility theory

- Defines axioms on preferences that involve uncertainty and ways to manipulate them.
- Uncertainty is modeled through **lotteries**
  - **Lottery:**

    $$[p : A; (1-p) : C]$$

    - Outcome A with probability p
    - Outcome C with probability (1-p)
- The following six constraints are known as the axioms of utility theory. The axioms are the most obvious semantic constraints on preferences with lotteries.
- **Notation:**

  $\succ$ - preferable

  $\sim$ - indifferent (equally preferable)

Axioms of the utility theory

- **Orderability:** Given any two states, the a rational agent prefers one of them, else the two as equally preferable.

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

- **Transitivity:** Given any three states, if an agent prefers A to B and prefers B to C, agent must prefer A to C.

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

- **Continuity:** If some state B is between A and C in preference, then there is a p for which the rational agent will be indifferent between state B and the lottery in which A comes with probability p, C with probability (1-p).

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