

CS 1571 Introduction to AI Lecture 22

Decision making in the presence of uncertainty

Milos Hauskrecht

milos@cs.pitt.edu

5329 Sennott Square

CS 1571 Intro to AI

Decision-making in the presence of uncertainty

- Computing the probability of some event may not be our ultimate goal
- Instead we are often interested in **making decisions about our future actions so that we satisfy goals**
- **Example: medicine**
 - Diagnosis is typically only the first step
 - The ultimate goal is to manage the patient in the best possible way. Typically many options available:
 - Surgery, medication, collect the new info (lab test)
 - There is an **uncertainty in the outcomes** of these procedures: patient can be improve, get worse or even die as a result of different management choices.

CS 1571 Intro to AI

Decision-making in the presence of uncertainty

Main issues:

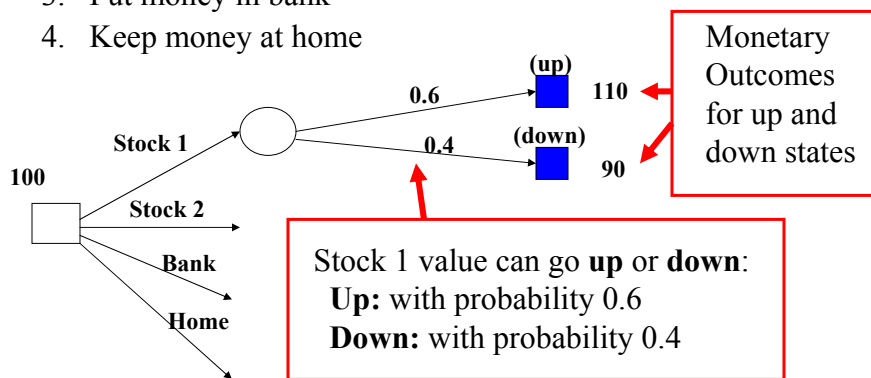
- How to model the decision process with uncertain outcomes in the computer ?
- How to make decisions about actions in the presence of uncertainty?

The field of **decision-making** studies ways of making decisions in the presence of uncertainty.

Decision making example.


Assume we want to invest \$100 for 6 months

- We have 4 choices:
 1. Invest in Stock 1
 2. Invest in Stock 2
 3. Put money in bank
 4. Keep money at home



Decision making example.


Investing of \$100 for 6 months



CS 1571 Intro to AI

Decision making example.

We need to make a choice whether to invest in Stock 1 or 2, put money into bank or keep them at home. But how?




CS 1571 Intro to AI

Decision making example.

Assume the simplified problem with the Bank and Home choices only.

The result is guaranteed – the outcome is deterministic




What is the rational choice assuming our goal is to make money?

Decision making. Deterministic outcome.

Assume the simplified problem with the Bank and Home choices only.

These choices are deterministic.



Our goal is to make money. What is the rational choice?


Answer: Put money into the bank. The choice is always strictly better in terms of the outcome

But what to do if we have uncertain outcomes?

Decision making. Stochastic outcome

- How to quantify the goodness of the stochastic outcome?

We want to compare it to deterministic and other stochastic outcomes.




?

Decision making. Stochastic outcome

- How to quantify the goodness of the stochastic outcome?

We want to compare it to deterministic and other stochastic outcomes.



Idea: Use expected value of the outcome


Expected value

- Let X be a random variable representing the monetary outcome with a discrete set of values Ω_X .
- Expected value** of X is:

$$E(X) = \sum_{x \in \Omega_X} xP(X = x)$$

Intuition: Expected value summarizes all stochastic outcomes into a single quantity.

- Example:**



- What is the expected value of the outcome of Stock 1 option?


Expected value

- Let X be a random variable representing the monetary outcome with a discrete set of values Ω_X .
- Expected value** of X is:

$$E(X) = \sum_{x \in \Omega_X} xP(X = x)$$

- Expected value** summarizes all stochastic outcomes into a single quantity


- Example:**



Expected value for the outcome of the Stock 1 option is:
 $0.6 \times 110 + 0.4 \times 90 = 66 + 36 = 102$

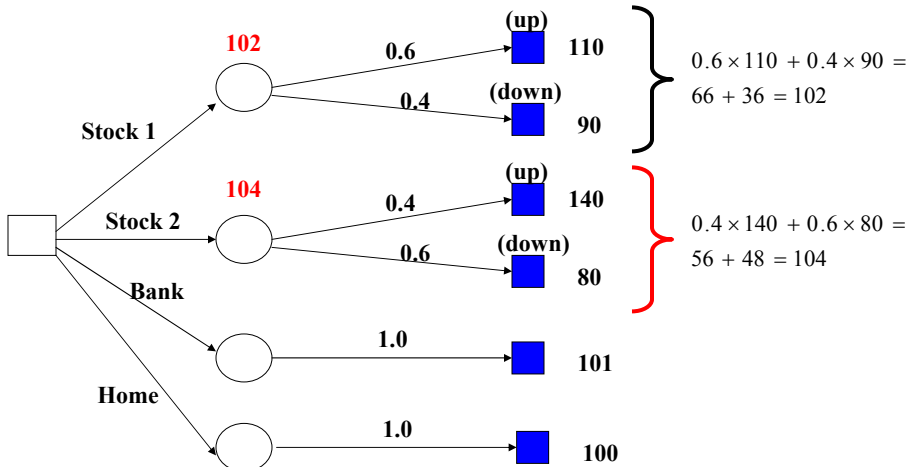
Expected values

Investing \$100 for 6 months



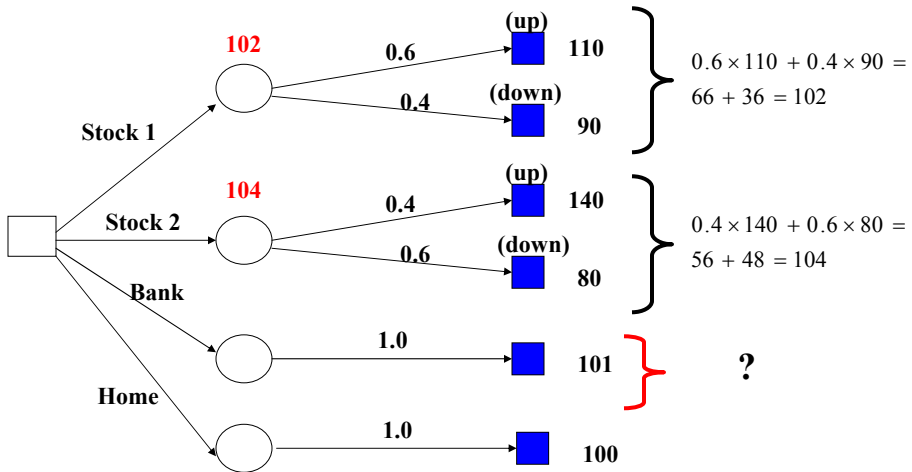
Expected values

Investing \$100 for 6 months



Expected values

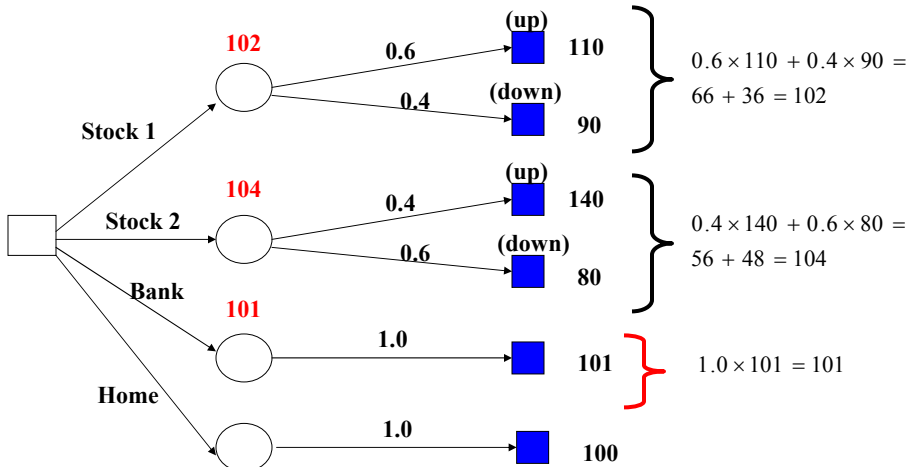
Investing \$100 for 6 months



CS 1571 Intro to AI

Expected values

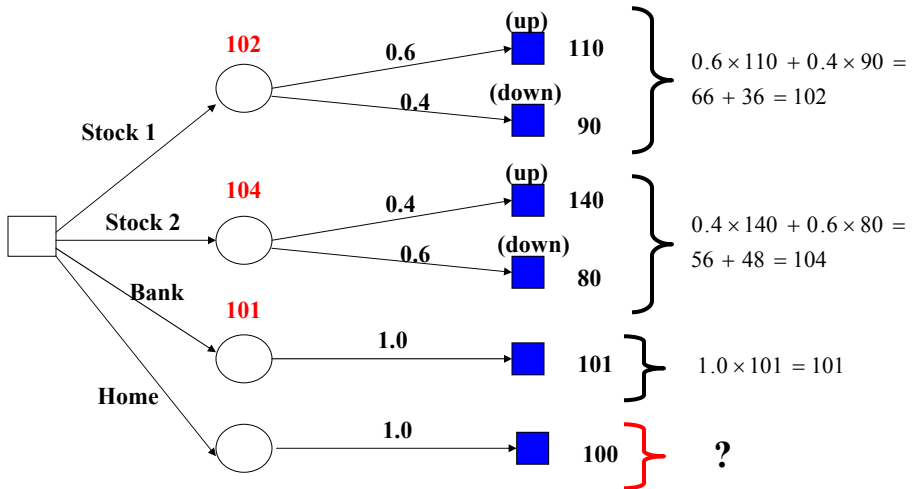
Investing \$100 for 6 months



CS 1571 Intro to AI

Expected values

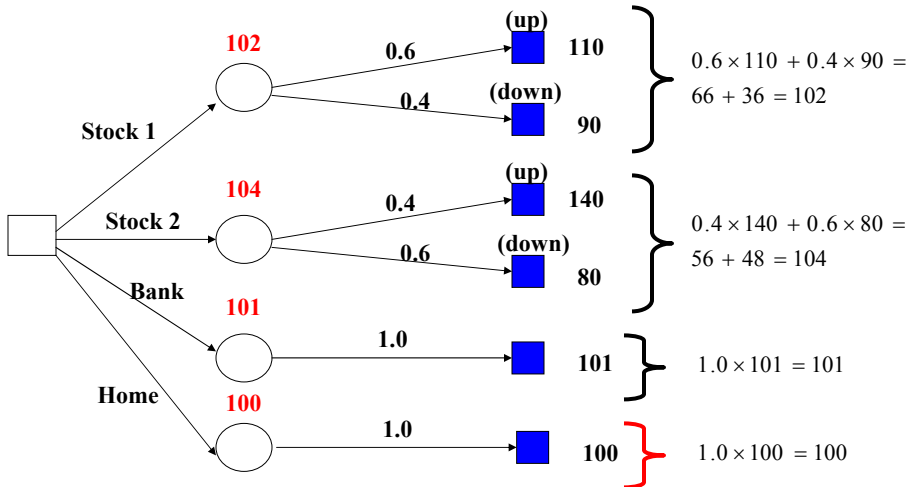
Investing \$100 for 6 months



CS 1571 Intro to AI

Expected values

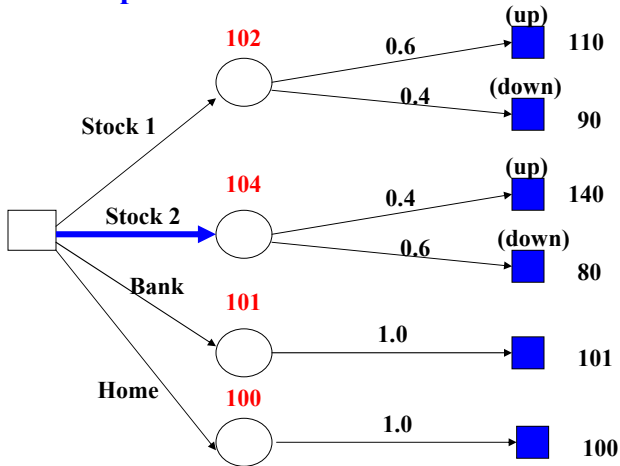
Investing \$100 for 6 months



CS 1571 Intro to AI

Selection based on expected values


The optimal action is the option that maximizes the expected outcome:



CS 1571 Intro to AI

Relation to the game search

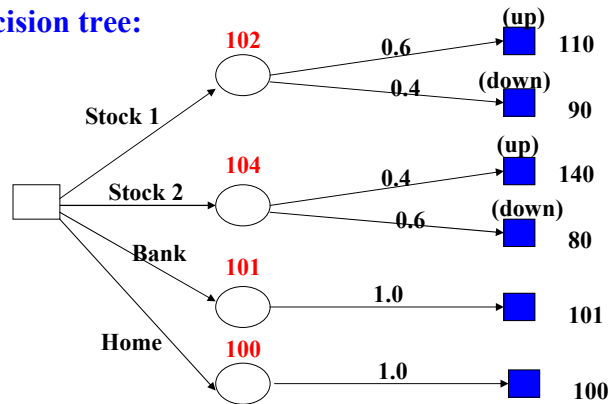
- **Game search: minimax algorithm**
 - considers the rational opponent and its best move
- **Decision making: maximizes the expectation**
 - play against the nature - stochastic non-malicious "opponent"



CS 1571 Intro to AI

(Stochastic) Decision tree

- Decision tree:



- decision node
- chance node
- outcome (value) node

CS 1571 Intro to AI

Sequential (multi-step) problems

The decision tree can be build to capture multi-step decision problems:

- Choose an action
- Observe the stochastic outcome
- And repeat

How to make decisions for multi-step problems?

- Start from the leaves of the decision tree (outcome nodes)
- Compute expectations at chance nodes
- Maximize at the decision nodes


Algorithm is sometimes called **expectimax**

CS 1571 Intro to AI

Multi-step problem example

Assume:

- Two investment periods
- Two actions: stock and bank

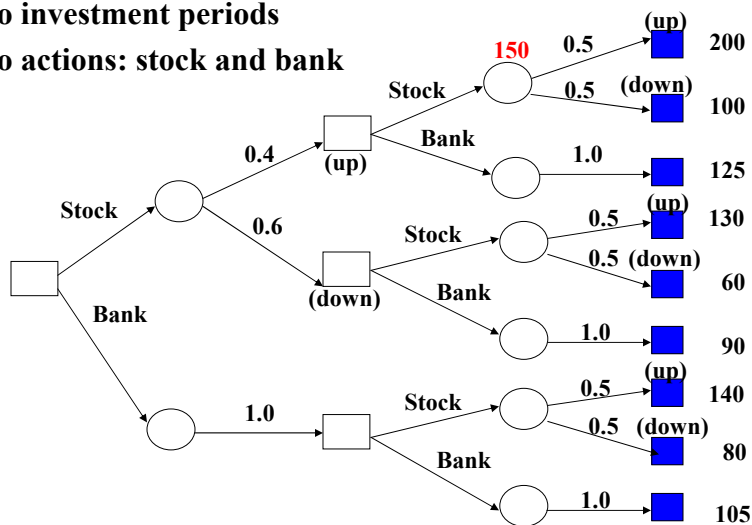


CS 1571 Intro to AI

Multi-step problem example

Assume:

- Two investment periods
- Two actions: stock and bank

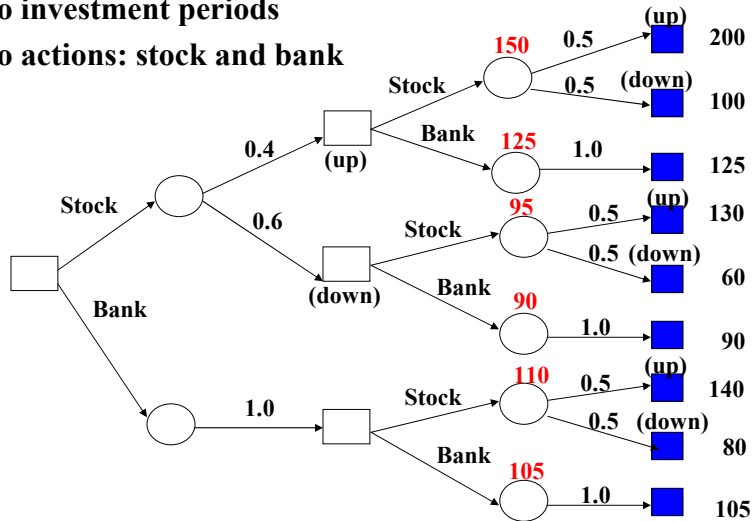


CS 1571 Intro to AI

Multi-step problem example

Assume:

- Two investment periods
- Two actions: stock and bank




CS 1571 Intro to AI

Multi-step problem example

Assume:

- Two investment periods
- Two actions: stock and bank




CS 1571 Intro to AI

Multi-step problem example

Assume:

- Two investment periods
- Two actions: stock and bank




CS 1571 Intro to AI

Multi-step problem example

Assume:

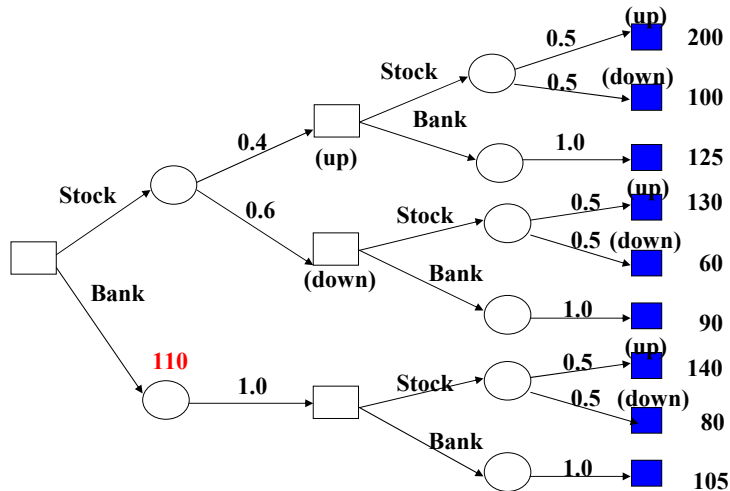
- Two investment periods
- Two actions: stock and bank



CS 1571 Intro to AI

Multi-step problems. Conditioning.

- Notice that the probability of up and down in the 2nd step is independent of the 1st step (=0.5)



CS 1571 Intro to AI

Conditioning in the decision tree

Probabilities:


- Later outcomes can be conditioned on the earlier stochastic outcomes and actions

Example: stock movement probabilities. Assume:

$$P(1^{\text{st}}=\text{up})=0.4$$

$$P(2^{\text{nd}}=\text{up}|1^{\text{st}}=\text{up})=0.4$$

$$P(2^{\text{nd}}=\text{up}|1^{\text{st}}=\text{down})=0.5$$



CS 1571 Intro to AI

Multi-step problems. Conditioning.

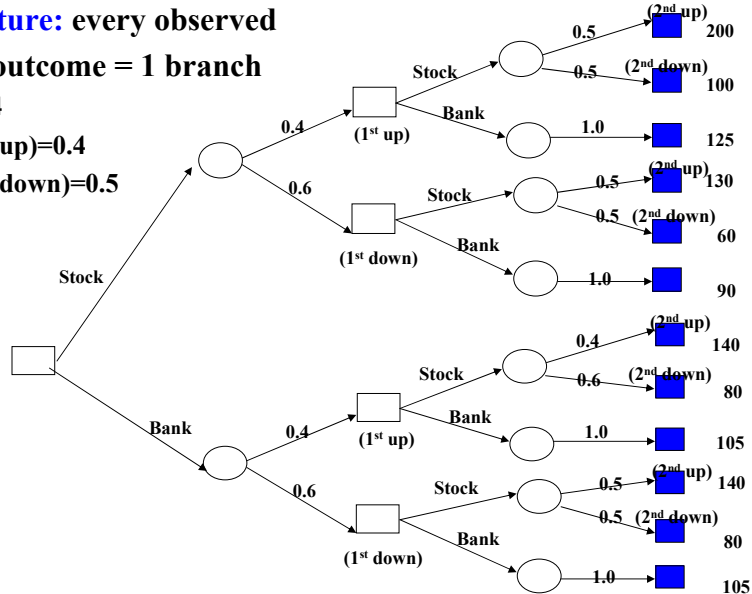
Tree Structure: every observed

stochastic outcome = 1 branch

$$P(1^{\text{st}}=\text{up})=0.4$$

$$P(2^{\text{nd}}=\text{up}|1^{\text{st}}=\text{up})=0.4$$

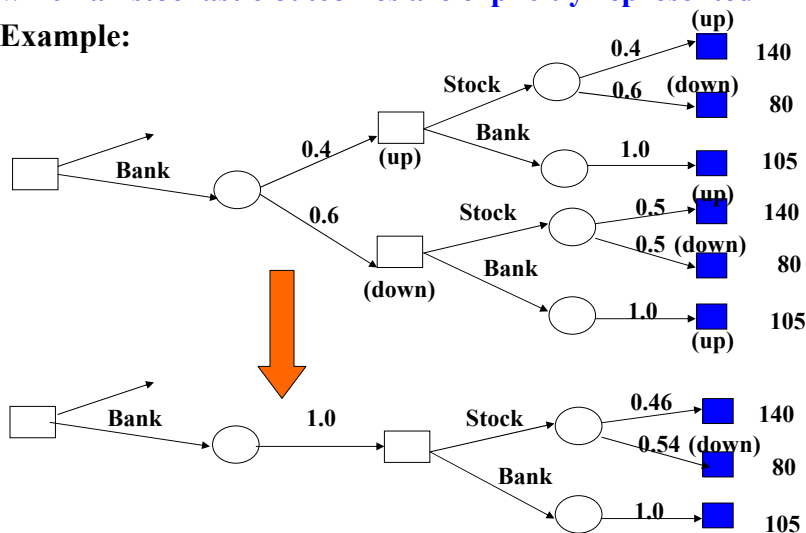
$$P(2^{\text{nd}}=\text{up}|1^{\text{st}}=\text{down})=0.5$$



CS 1571 Intro to AI

Multi-step problems. Conditioning.

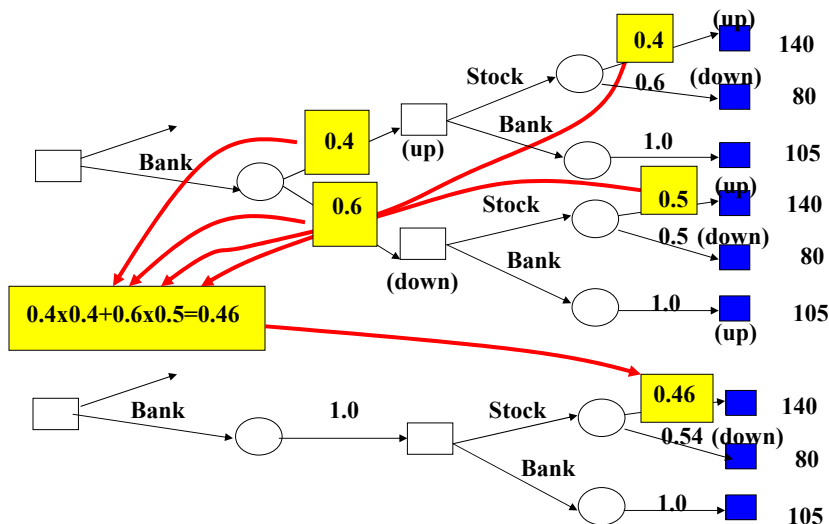
- If conditioning is present – we can always build the tree in which all stochastic outcomes are explicitly represented
- Example:



CS 1571 Intro to AI

Multi-step problems. Conditioning.

- If conditioning is present – we can always build the tree in which all stochastic outcomes are explicitly represented

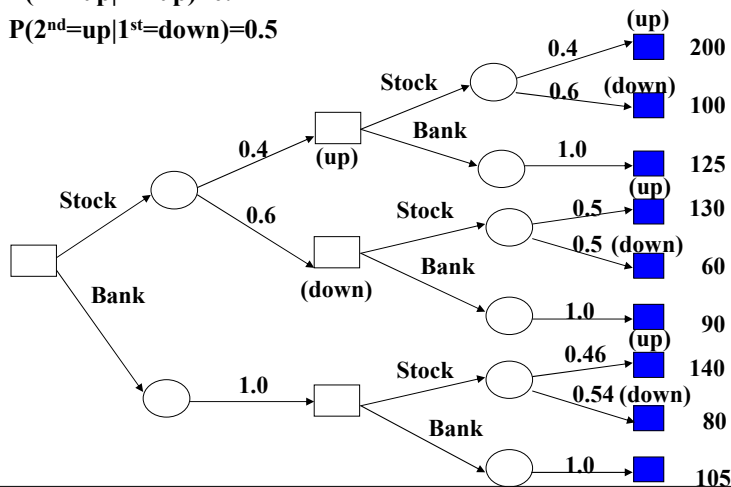


CS 1571 Intro to AI

Multi-step problems. Conditioning.

2nd step outcomes to be dependent on the 1st period outcomes:

- $P(1^{st}=up)=0.4$
- $P(2^{nd}=up|1^{st}=up)=0.4$
- $P(2^{nd}=up|1^{st}=down)=0.5$



CS 1571 Intro to AI

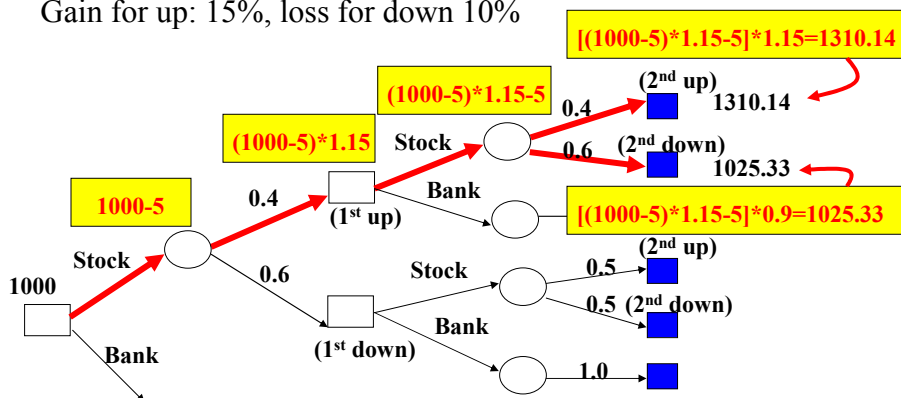
Trajectory payoffs

- Outcome values at leaf nodes (e.g. monetary values)
 - Rewards and costs for the path trajectory

Example: stock fees and gains. **Assume:**

Fee per period: \$5 paid at the beginning

Gain for up: 15%, loss for down 10%



CS 1571 Intro to AI

Constructing a decision tree

- The decision tree is rarely given to you directly.
 - Part of the problem is to construct the tree.

Example: stocks, bonds, bank for k periods

Stock:

- Probability of stocks going up in the first period: 0.3
- Probability of stocks going up in subsequent periods:
 - $P(\text{kth step}=\text{Up} | (\text{k}-1)\text{th step}=\text{Up})=0.4$
 - $P(\text{kth step}=\text{Up} | (\text{k}-1)\text{th step}=\text{Down})=0.5$
- Return if stock goes up: 15 % if down: 10%
- Fixed fee per investment period: \$5

Bonds:

- Probability of value up: 0.5, down: 0.5
- Return if bond value is going up: 7%, if down: 3%
- Fee per investment period: \$2

Bank:

- Guaranteed return of 3% per period, no fee

CS 1571 Intro to AI