CS 1571 Introduction to AI
Lecture 24

Monte Carlo Inference
Decision making in the presence of uncertaity

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

Monte Carlo approaches

- **MC approximation**:
  – The probability is approximated using sample frequencies
  – **Example:**

$$\tilde{P}(B=T, J=T) = \frac{N_{B=T,J=T}}{N}$$

$N_{B=T,J=T}$: *# samples with $B=T, J=T$*
N: *total # samples*

- **BBN sampling:**

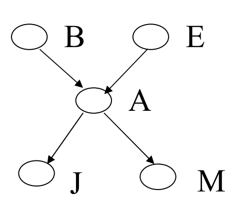

Generate sample in a top down manner, following the links

- **One sample gives one assignment of values to all variables**

BBN sampling example

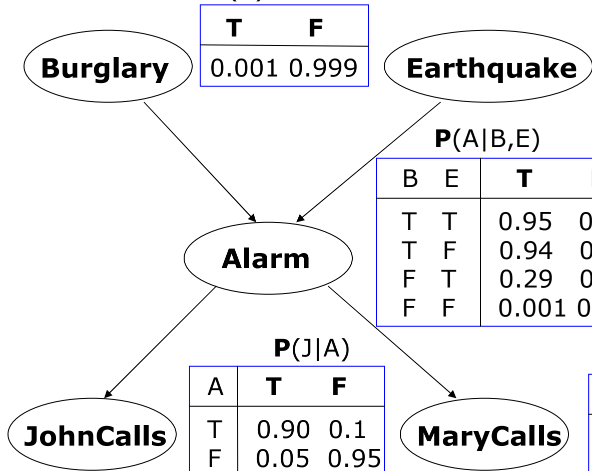

P(B)

T	F
0.001	0.999

P(E)

T	F
0.002	0.998

P(A|B,E)

B	E	T	F
T	T	0.95	0.05
T	F	0.94	0.06
F	T	0.29	0.71
F	F	0.001	0.999

P(J|A)

A	T	F
T	0.90	0.1
F	0.05	0.95

P(M|A)

A	T	F
T	0.7	0.3
F	0.01	0.99

BBN sampling example

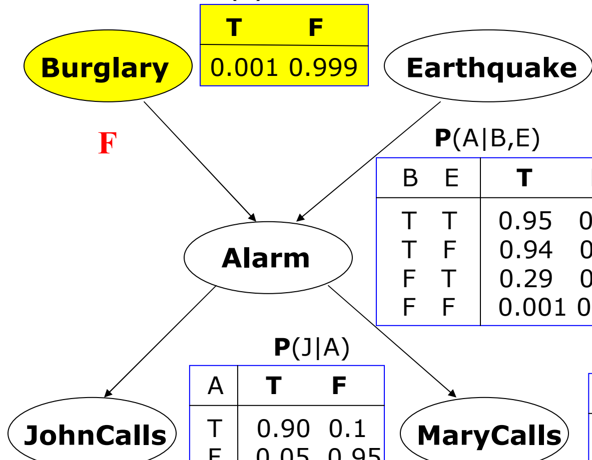

P(B)

T	F
0.001	0.999

P(E)

T	F
0.002	0.998

P(A|B,E)

B	E	T	F
T	T	0.95	0.05
T	F	0.94	0.06
F	T	0.29	0.71
F	F	0.001	0.999

P(J|A)

A	T	F
T	0.90	0.1
F	0.05	0.95

P(M|A)

A	T	F
T	0.7	0.3
F	0.01	0.99

BBN sampling example

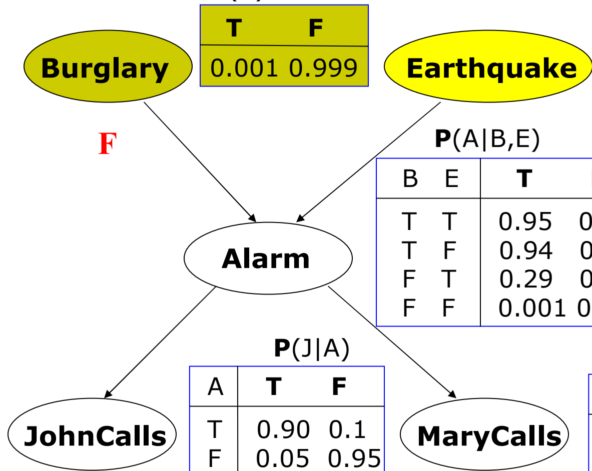

P(B)

T	F
0.001	0.999

P(E)

T	F
0.002	0.998

P(A|B,E)

B	E	T	F
T	T	0.95	0.05
T	F	0.94	0.06
F	T	0.29	0.71
F	F	0.001	0.999

P(J|A)

A	T	F
T	0.90	0.1
F	0.05	0.95

P(M|A)

A	T	F
T	0.7	0.3
F	0.01	0.99

BBN sampling example

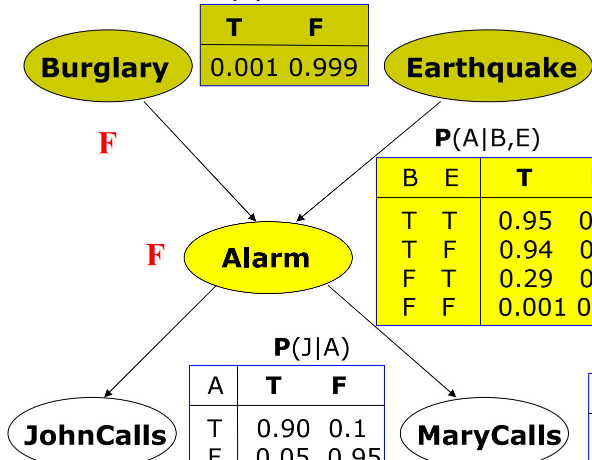

P(B)

T	F
0.001	0.999

P(E)

T	F
0.002	0.998

P(A|B,E)

B	E	T	F
T	T	0.95	0.05
T	F	0.94	0.06
F	T	0.29	0.71
F	F	0.001	0.999

P(J|A)

A	T	F
T	0.90	0.1
F	0.05	0.95

P(M|A)

A	T	F
T	0.7	0.3
F	0.01	0.99

BBN sampling example

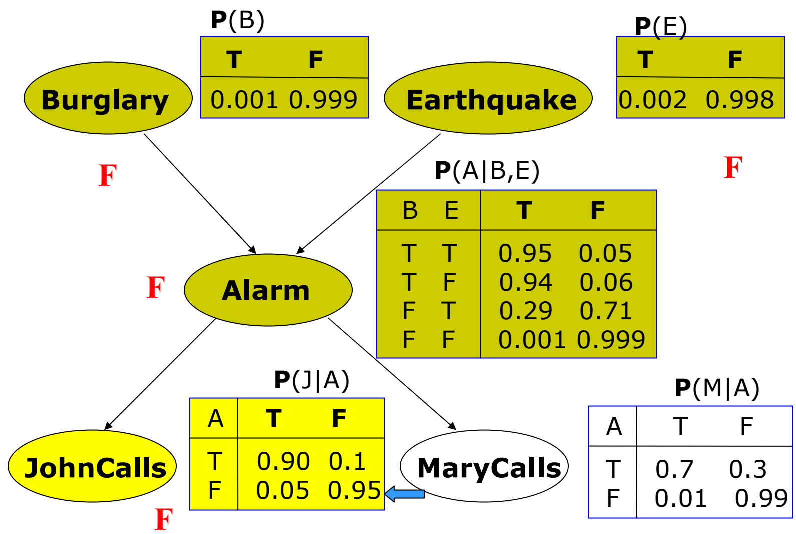

P(B)

T	F
0.001	0.999

P(E)

T	F
0.002	0.998

P(A|B,E)

B	E	T	F
T	T	0.95	0.05
T	F	0.94	0.06
F	T	0.29	0.71
F	F	0.001	0.999

P(J|A)

A	T	F
T	0.90	0.1
F	0.05	0.95

P(M|A)

A	T	F
T	0.7	0.3
F	0.01	0.99

BBN sampling example

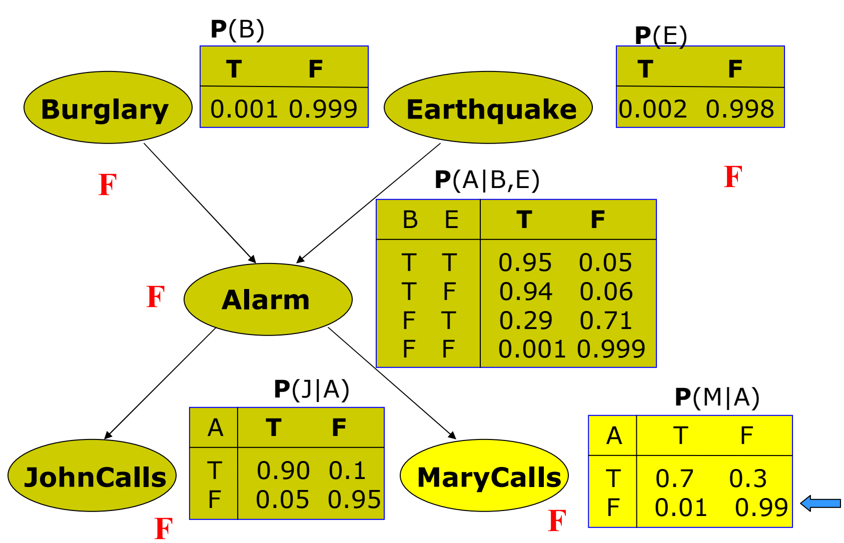

P(B)

T	F
0.001	0.999

P(E)

T	F
0.002	0.998

P(A|B,E)

B	E	T	F
T	T	0.95	0.05
T	F	0.94	0.06
F	T	0.29	0.71
F	F	0.001	0.999

P(J|A)

A	T	F
T	0.90	0.1
F	0.05	0.95

P(M|A)

A	T	F
T	0.7	0.3
F	0.01	0.99

BBN sampling example

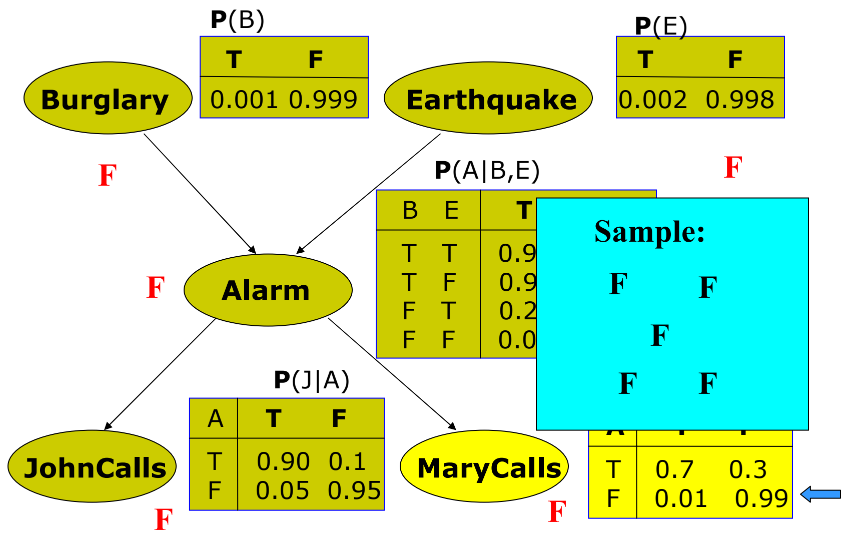

Monte Carlo approaches

- **MC approximation of conditional probabilities**:
  - The probability is approximated using sample frequencies
  - **Example:**

$$\tilde{P}(B = T \mid J = T) = \frac{N_{B=T,J=T}}{N_{J=T}}$$

$N_{B=T,J=T}$: $\#\,samples\ \ with\ \ B = T, J = T$

$N_{J=T}$: $\#\,samples\ \ with\ \ J = T$

- **Rejection sampling:**
  - Generate samples from the full joint by sampling BBN
  - Use only samples that agree with the condition, the remaining samples are rejected
- **Problem:** many samples can be rejected

Likelihood weighting

Idea: generate only samples consistent with an evidence (or conditioning event)

– Benefit: Avoids inefficiencies of rejection sampling

Problem:

- **the distribution generated by enforcing the conditioning variables to set values is biased**
- simple counts are not sufficient to estimate the probabilities

Solution:

- **With every sample keep a weight with which it should count towards the estimate**

$$\tilde{P}(B=T \mid J=T) = \frac{\sum_{samples\ with\ B=T\ and\ J=T} w_{B=T}}{\sum_{samples\ with\ any\ value\ of\ B\ and\ J=T} w_{B=x}}$$

BBN likelihood weighting example

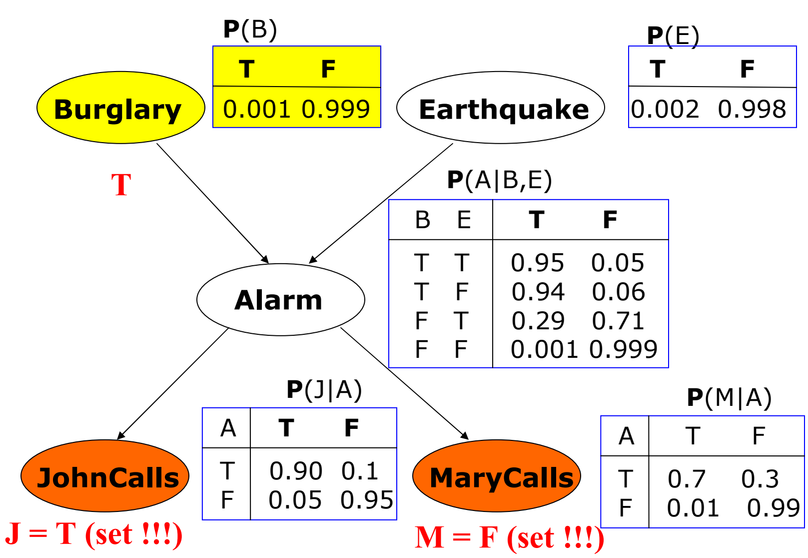

BBN likelihood weighting example

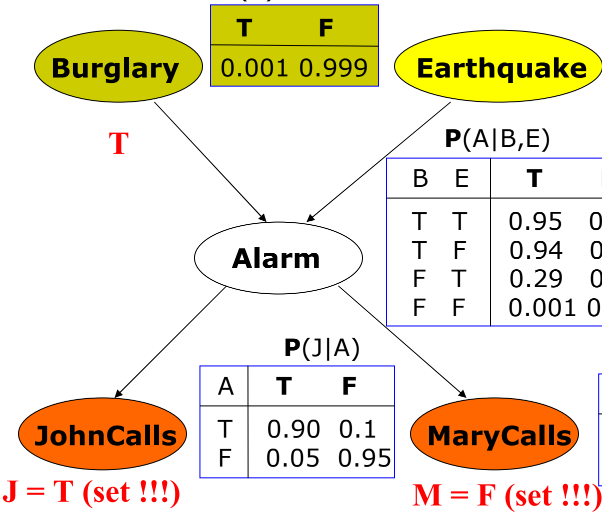

P(B)

T	F
0.001	0.999

P(E)

T	F
0.002	0.998

P(A|B,E)

B	E	T	F
T	T	0.95	0.05
T	F	0.94	0.06
F	T	0.29	0.71
F	F	0.001	0.999

P(J|A)

A	T	F
T	0.90	0.1
F	0.05	0.95

P(M|A)

A	T	F
T	0.7	0.3
F	0.01	0.99

BBN likelihood weighting example

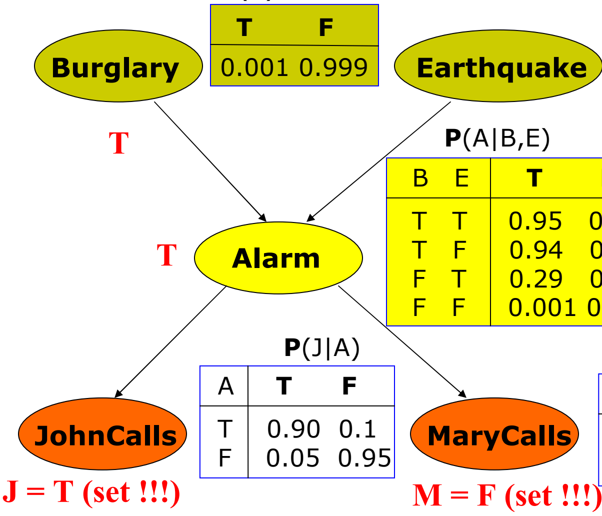

P(B)

T	F
0.001	0.999

P(E)

T	F
0.002	0.998

P(A|B,E)

B	E	T	F
T	T	0.95	0.05
T	F	0.94	0.06
F	T	0.29	0.71
F	F	0.001	0.999

P(J|A)

A	T	F
T	0.90	0.1
F	0.05	0.95

P(M|A)

A	T	F
T	0.7	0.3
F	0.01	0.99

BBN likelihood weighting example

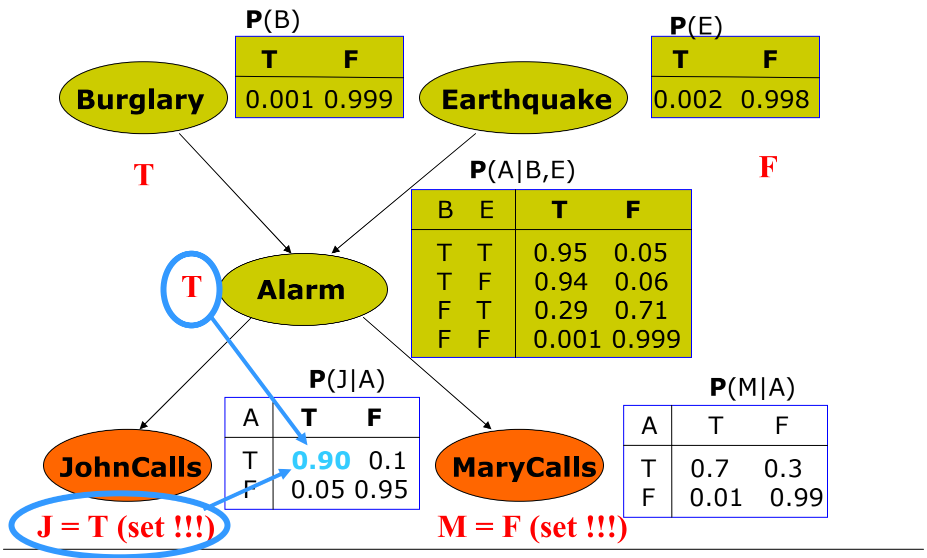

P(B)

T	F
0.001	0.999

P(E)

T	F
0.002	0.998

P(A|B,E)

B	E	T	F
T	T	0.95	0.05
T	F	0.94	0.06
F	T	0.29	0.71
F	F	0.001	0.999

P(J|A)

A	T	F
T	0.90	0.1
F	0.05	0.95

P(M|A)

A	T	F
T	0.7	0.3
F	0.01	0.99

BBN likelihood weighting example

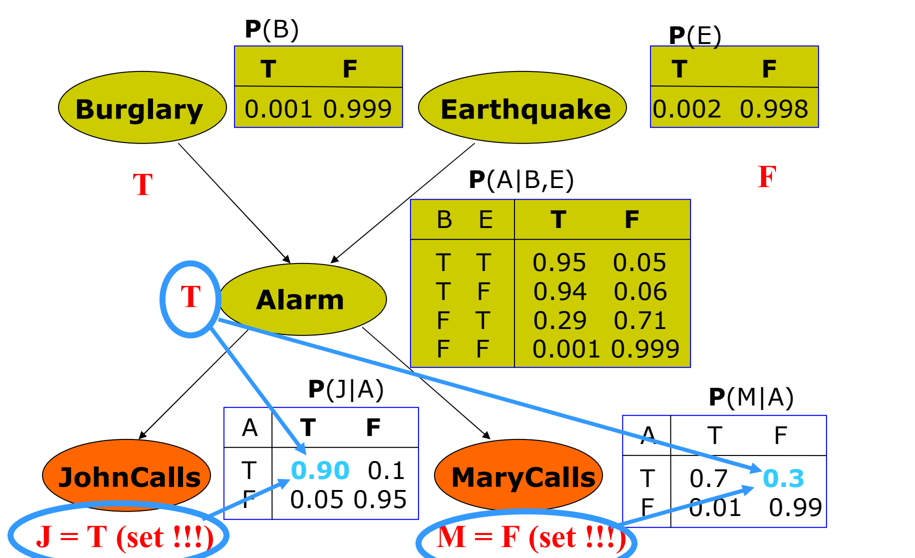

P(B)

T	F
0.001	0.999

P(E)

T	F
0.002	0.998

P(A|B,E)

B	E	T	F
T	T	0.95	0.05
T	F	0.94	0.06
F	T	0.29	0.71
F	F	0.001	0.999

P(J|A)

A	T	F
T	0.90	0.1
F	0.05	0.95

P(M|A)

A	T	F
T	0.7	0.3
F	0.01	0.99

BBN likelihood weighting example

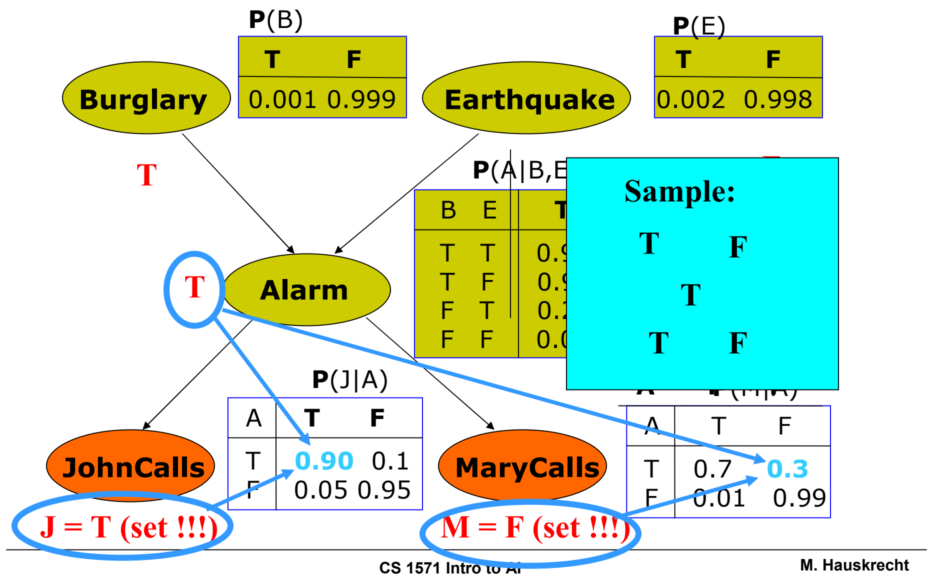

BBN likelihood weighting example

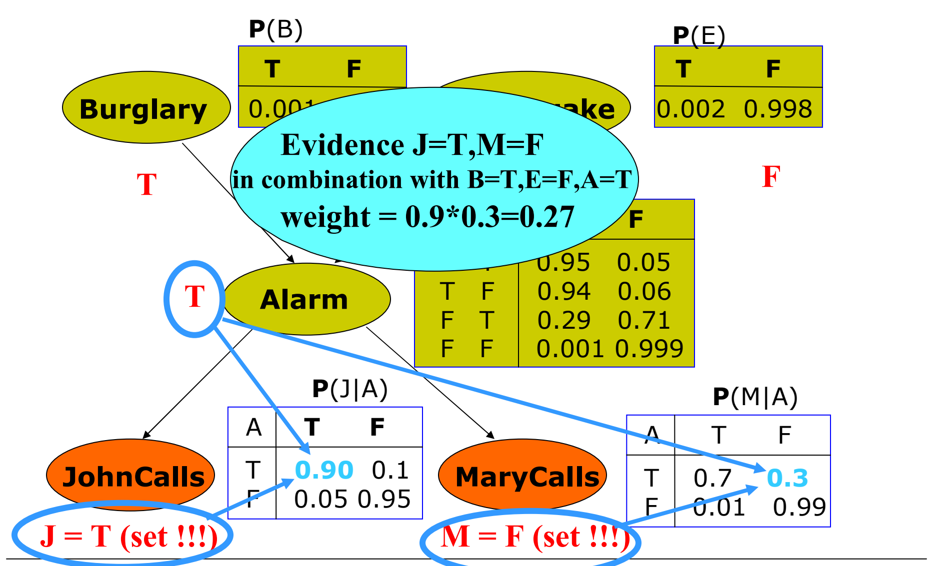

BBN likelihood weighting example

Second sample

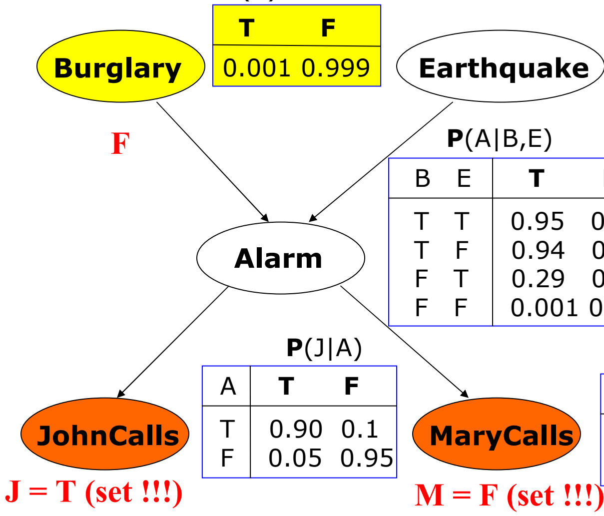

P(B)

T	F
0.001	0.999

P(E)

T	F
0.002	0.998

P(A|B,E)

B	E	T	F
T	T	0.95	0.05
T	F	0.94	0.06
F	T	0.29	0.71
F	F	0.001	0.999

P(J|A)

A	T	F
T	0.90	0.1
F	0.05	0.95

P(M|A)

A	T	F
T	0.7	0.3
F	0.01	0.99

BBN likelihood weighting example

Second sample

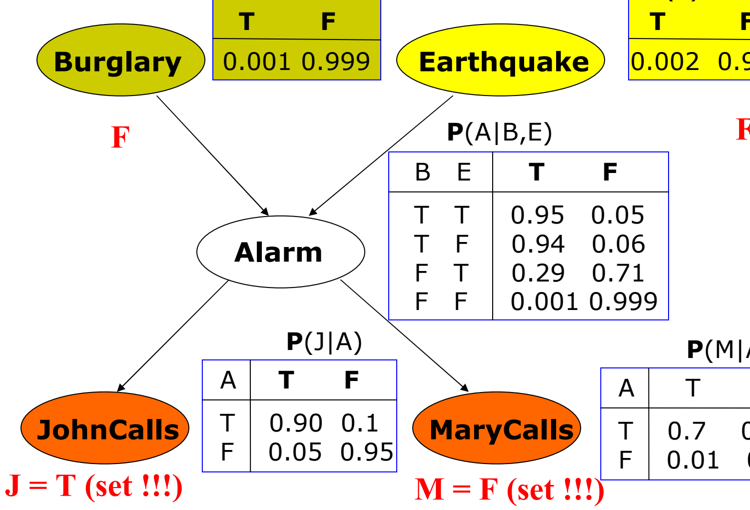

P(B)

T	F
0.001	0.999

P(E)

T	F
0.002	0.998

P(A|B,E)

B	E	T	F
T	T	0.95	0.05
T	F	0.94	0.06
F	T	0.29	0.71
F	F	0.001	0.999

P(J|A)

A	T	F
T	0.90	0.1
F	0.05	0.95

P(M|A)

A	T	F
T	0.7	0.3
F	0.01	0.99

BBN likelihood weighting example

Second sample

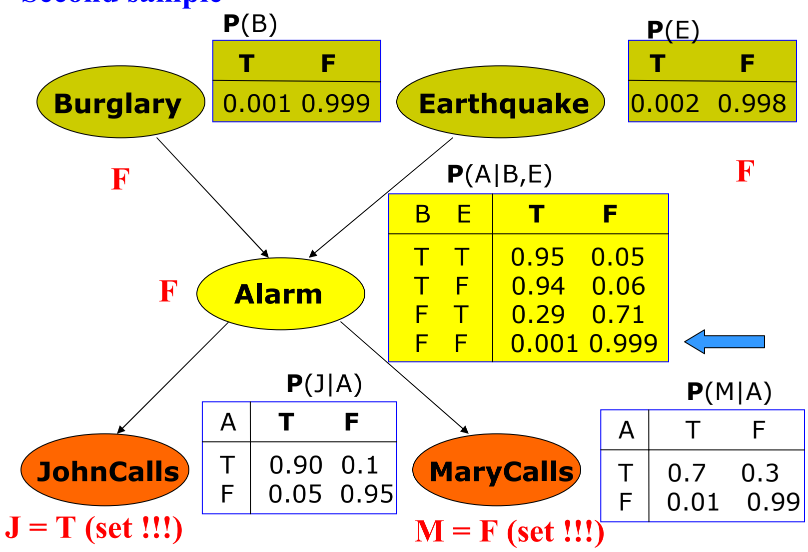

P(B)

T	F
0.001	0.999

P(E)

T	F
0.002	0.998

P(A|B,E)

B	E	T	F
T	T	0.95	0.05
T	F	0.94	0.06
F	T	0.29	0.71
F	F	0.001	0.999

P(J|A)

A	T	F
T	0.90	0.1
F	0.05	0.95

P(M|A)

A	T	F
T	0.7	0.3
F	0.01	0.99

BBN likelihood weighting example

Second sample

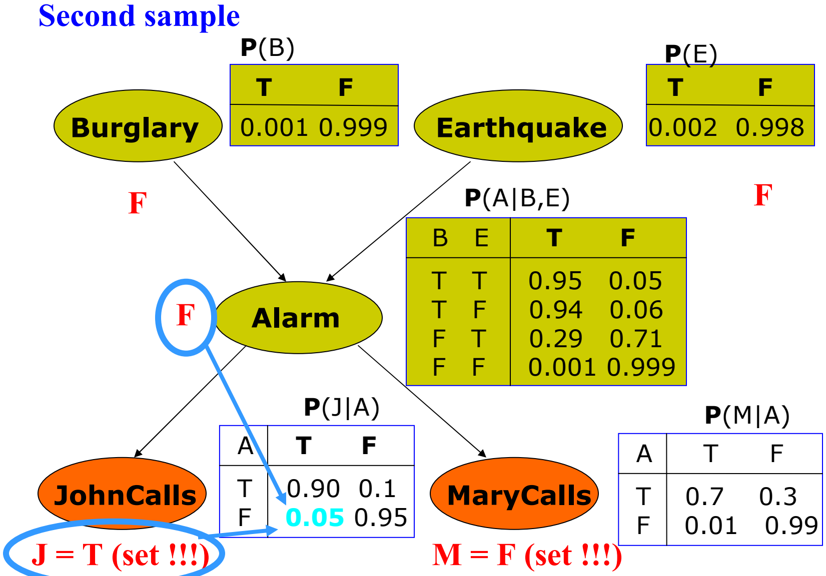

P(B)

T	F
0.001	0.999

P(E)

T	F
0.002	0.998

P(A|B,E)

B	E	T	F
T	T	0.95	0.05
T	F	0.94	0.06
F	T	0.29	0.71
F	F	0.001	0.999

P(J|A)

A	T	F
T	0.90	0.1
F	0.05	0.95

P(M|A)

A	T	F
T	0.7	0.3
F	0.01	0.99

BBN likelihood weighting example

Second sample

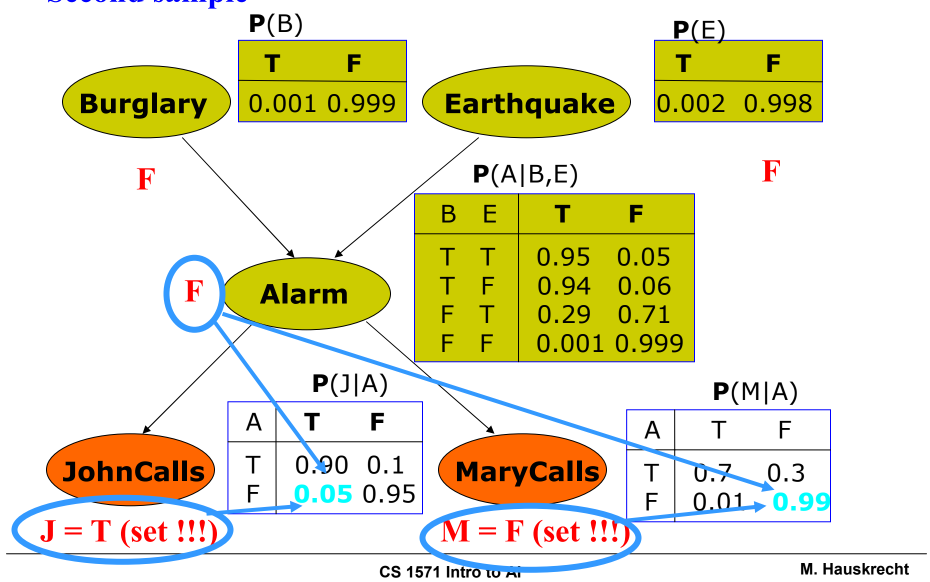

P(B)

T	F
0.001	0.999

P(E)

T	F
0.002	0.998

P(A|B,E)

B	E	T	F
T	T	0.95	0.05
T	F	0.94	0.06
F	T	0.29	0.71
F	F	0.001	0.999

P(J|A)

A	T	F
T	0.90	0.1
F	0.05	0.95

P(M|A)

A	T	F
T	0.7	0.3
F	0.01	0.99

BBN likelihood weighting example

Second sample

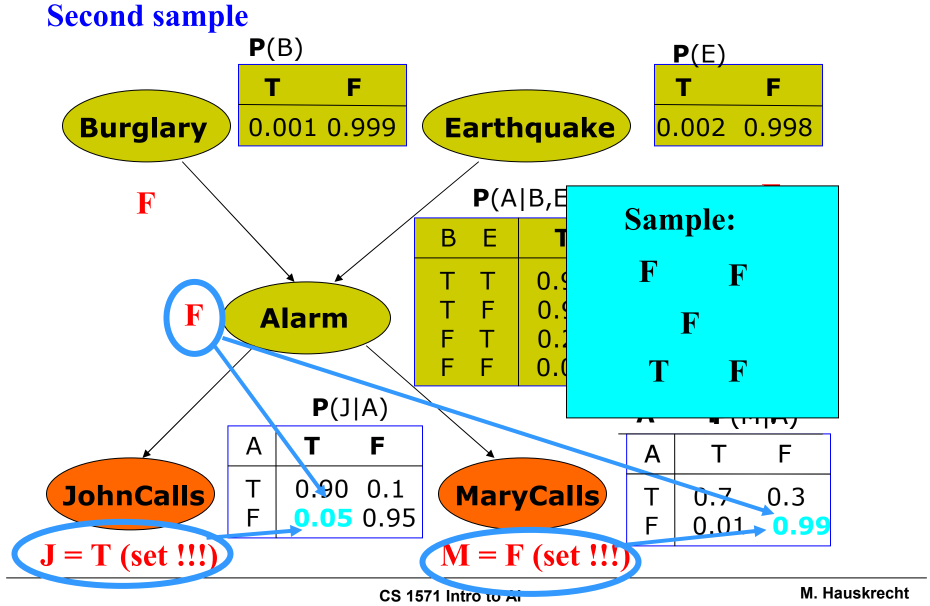

P(B)

T	F
0.001	0.999

P(E)

T	F
0.002	0.998

P(A|B,E

B	E	T
T	T	0.9
T	F	0.9
F	T	0.2
F	F	0.0

P(J|A)

A	T	F
T	0.90	0.1
F	0.05	0.95

A	T	F
T	0.7	0.3
F	0.01	0.99

BBN likelihood weighting example

Second sample

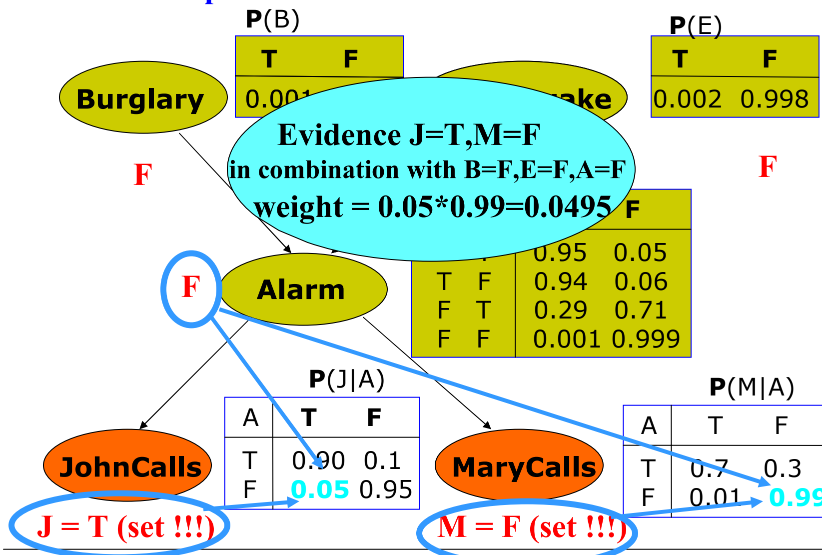

Likelihood weighting

- Assume we have generated the following M samples:

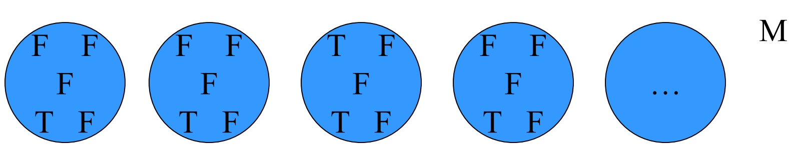

How to make the samples consistent?

Weight each sample by probability with which it agrees with the conditioning evidence P(e).

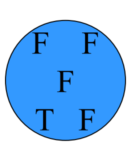

← **Weight 0.0495**

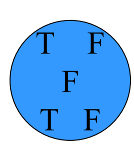

← **Weight 0.27**

Decision-making in the presence of uncertainty

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.

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

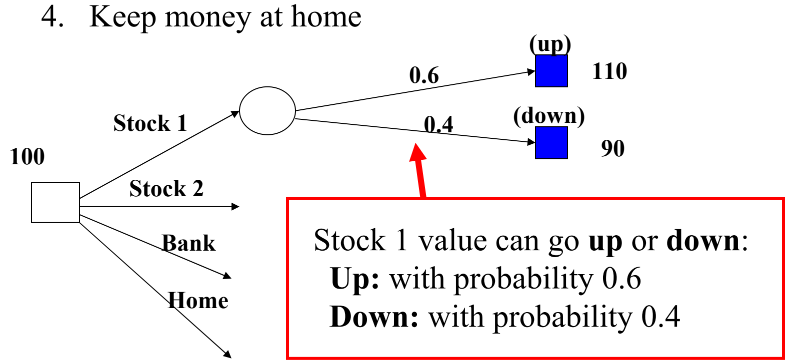

Decision making example.

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

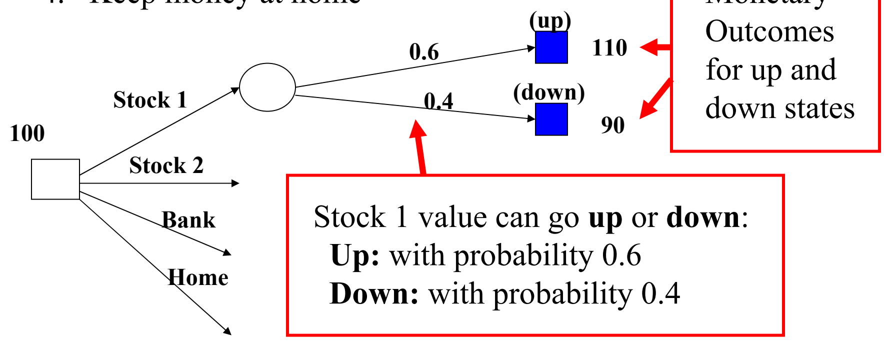

Decision making example.

Investing of $100 for 6 months

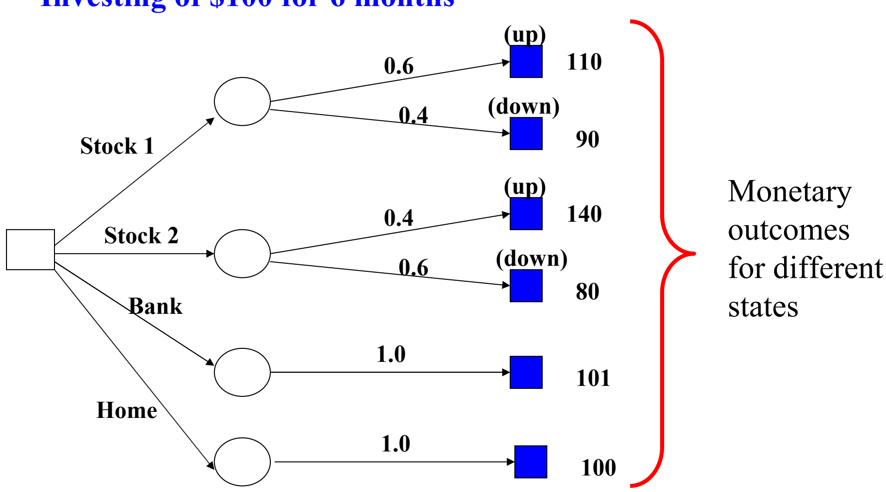

Decision making example.

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?

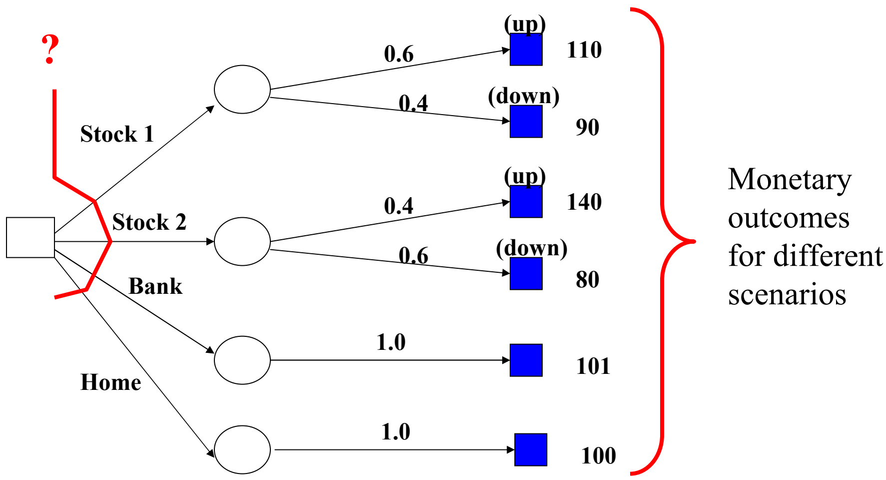

Decision making example.

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

The result is guaranteed – the outcome is deterministic

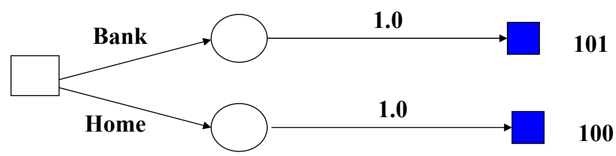

What is the rational choice assuming our goal is to make money?

Decision making. Deterministic outcome.

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

These choices are deterministic.

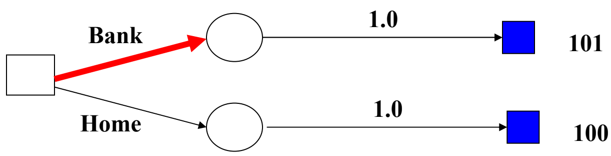

Our goal is to make money. What is the rational choice?

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.

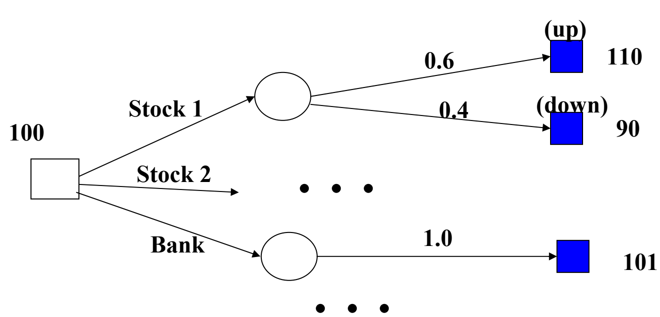

?

Decision making. Stochastic outcome

- **How to quantify the goodness of the stochastic outcome?**
  We want to compare it to deterministic and other stochastic outcomes.

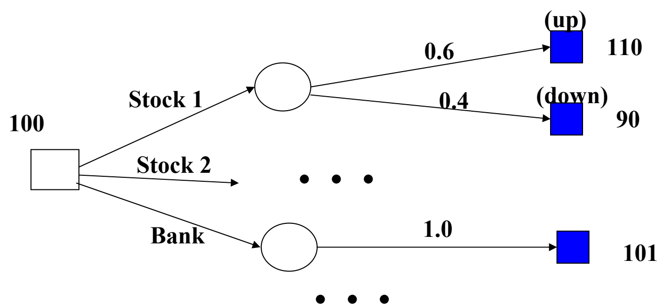

Idea: Use the expected value of the outcome

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:

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

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

- **Example:**

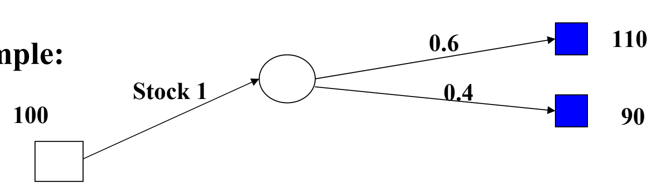

- What is the expected value of the outcome of Stock 1 option?

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:

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

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

- **Example:**

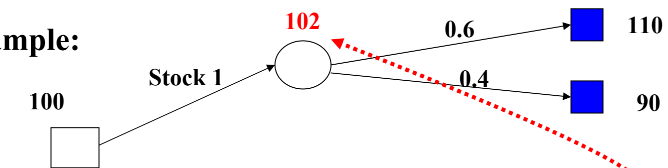

Expected value for the outcome of the Stock 1 option is:

$$0.6 \times 110 + 0.4 \times 90 = 66 + 36 = \boxed{102}$$

Expected values

Investing $100 for 6 months

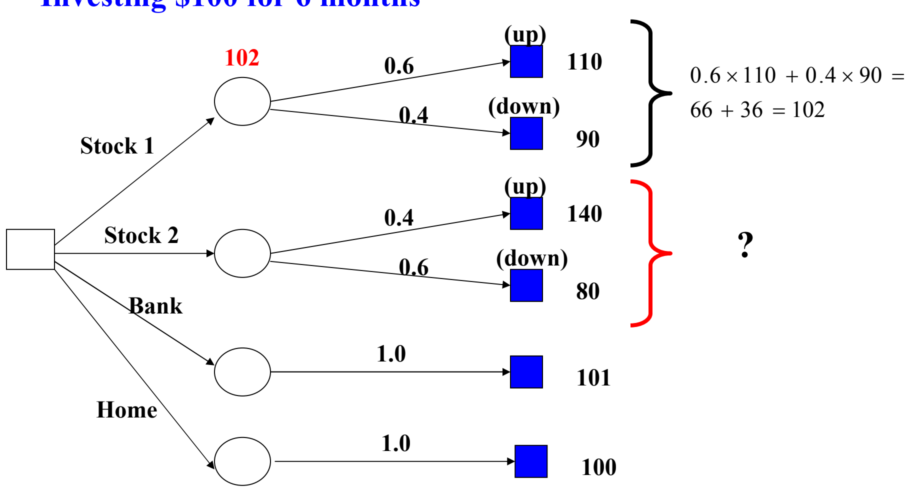

Expected values

Investing $100 for 6 months

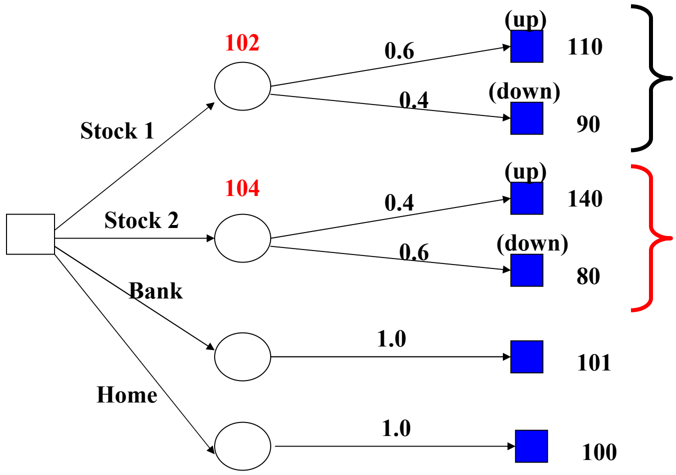

$0.6 \times 110 + 0.4 \times 90 =$
$66 + 36 = 102$

$0.4 \times 140 + 0.6 \times 80 =$
$56 + 48 = 104$

Expected values

Investing $100 for 6 months

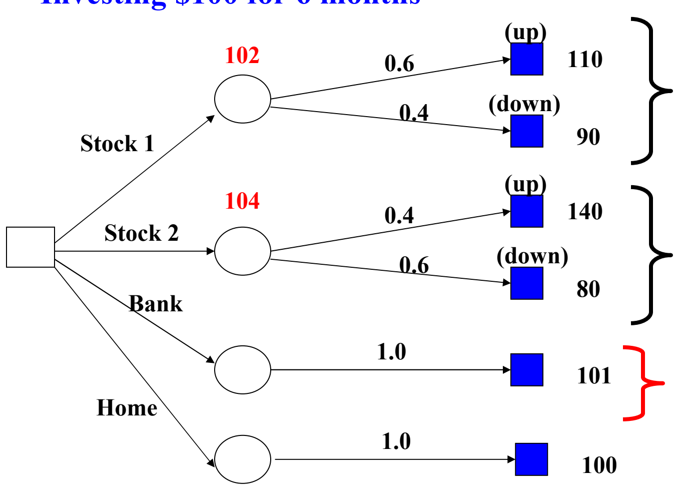

$0.6 \times 110 + 0.4 \times 90 =$
$66 + 36 = 102$

$0.4 \times 140 + 0.6 \times 80 =$
$56 + 48 = 104$

?

Expected values

Investing \$100 for 6 months

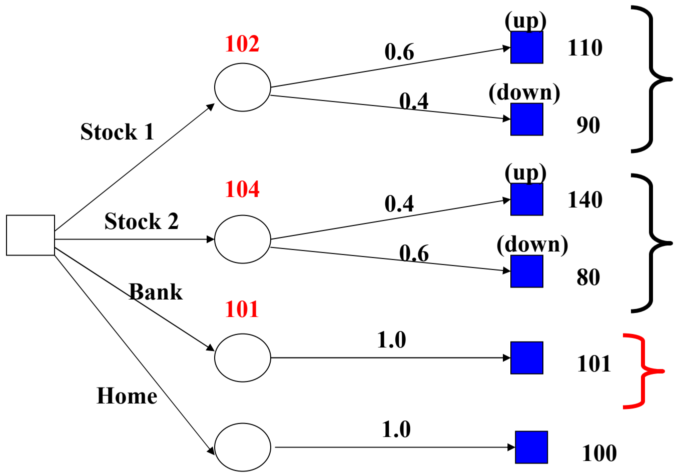

$0.6 \times 110 + 0.4 \times 90 =$
$66 + 36 = 102$

$0.4 \times 140 + 0.6 \times 80 =$
$56 + 48 = 104$

$1.0 \times 101 = 101$

Expected values

Investing \$100 for 6 months

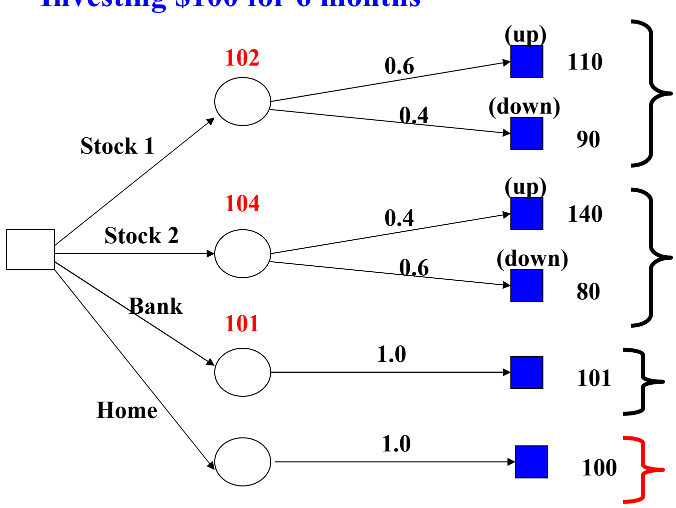

$0.6 \times 110 + 0.4 \times 90 =$
$66 + 36 = 102$

$0.4 \times 140 + 0.6 \times 80 =$
$56 + 48 = 104$

$1.0 \times 101 = 101$

?

Expected values

Investing $100 for 6 months

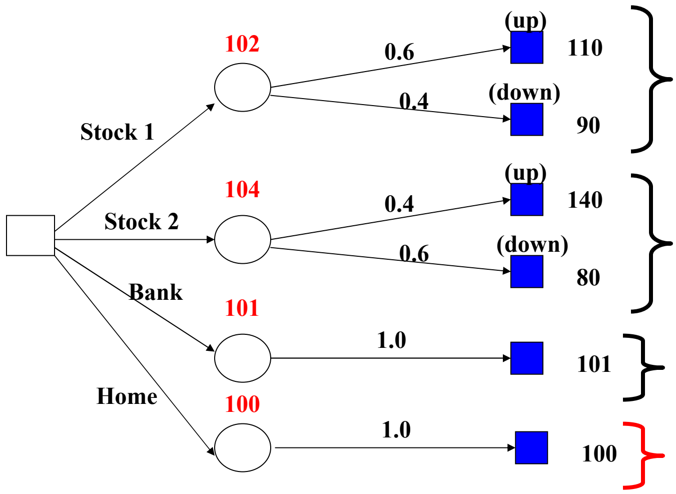

$0.6 \times 110 + 0.4 \times 90 =$
$66 + 36 = 102$

$0.4 \times 140 + 0.6 \times 80 =$
$56 + 48 = 104$

$1.0 \times 101 = 101$

$1.0 \times 100 = 100$

Selection based on expected values

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

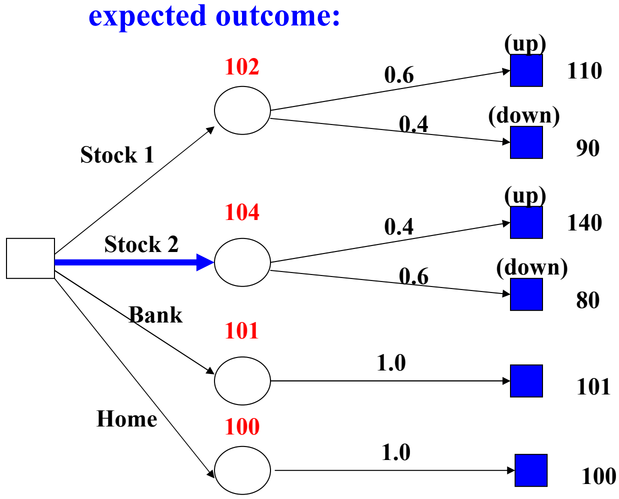

Relation to the game search

- 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"

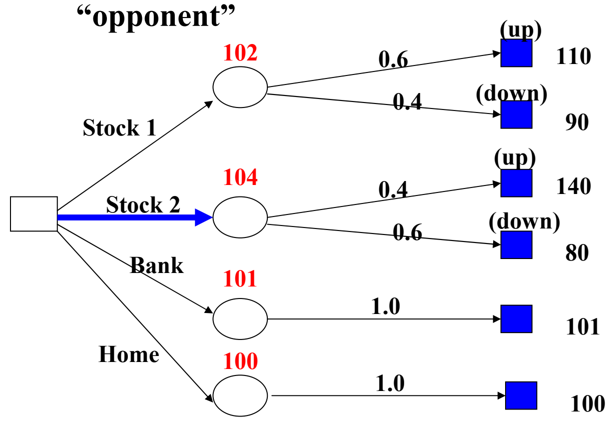

(Stochastic) Decision tree

- Decision tree:

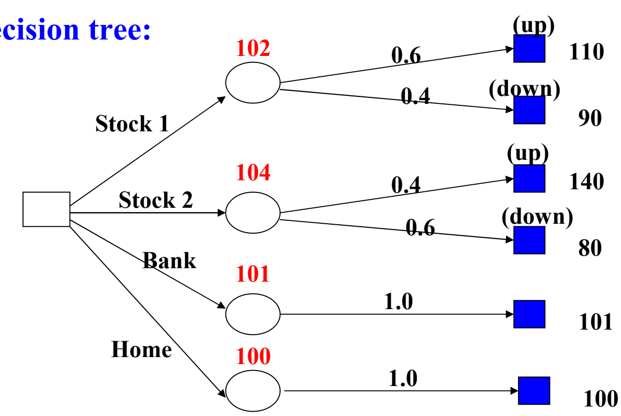

- decision node
- chance node
- outcome (value) node

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**

Multi-step problem example

Assume:

- **Two investment periods**
- **Two actions: stock and bank**

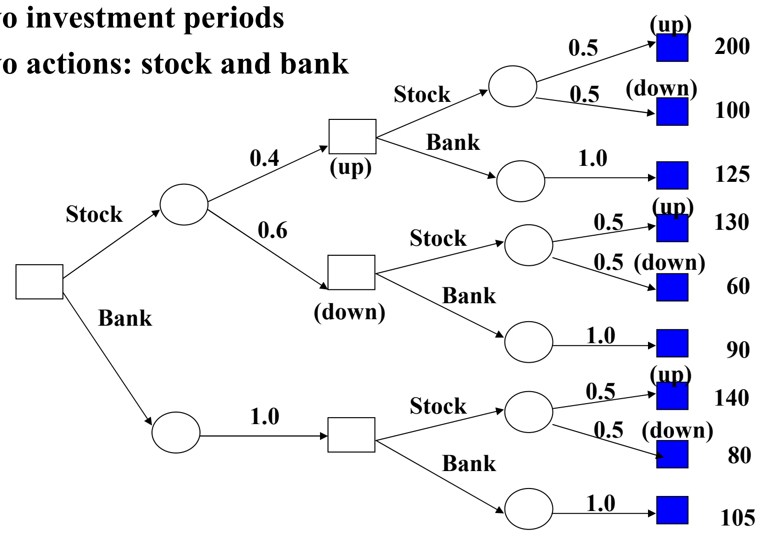

Multi-step problem example

Assume:

- Two investment periods
- Two actions: stock and bank

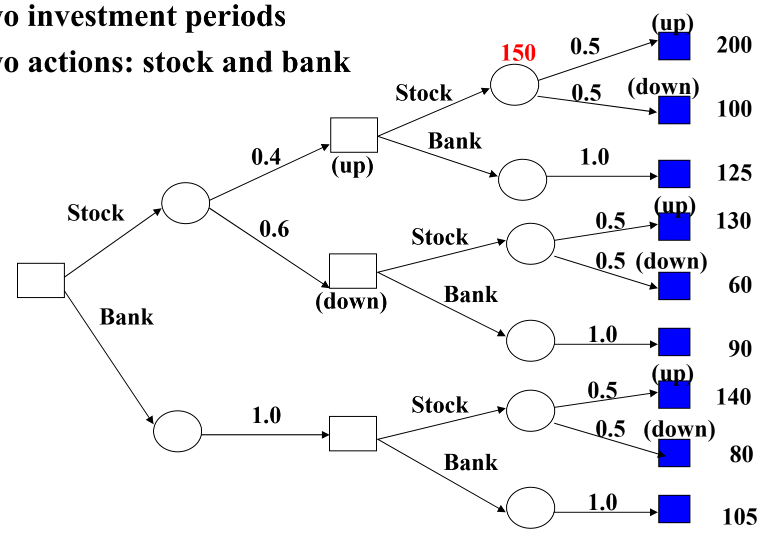

Multi-step problem example

Assume:

- Two investment periods
- Two actions: stock and bank

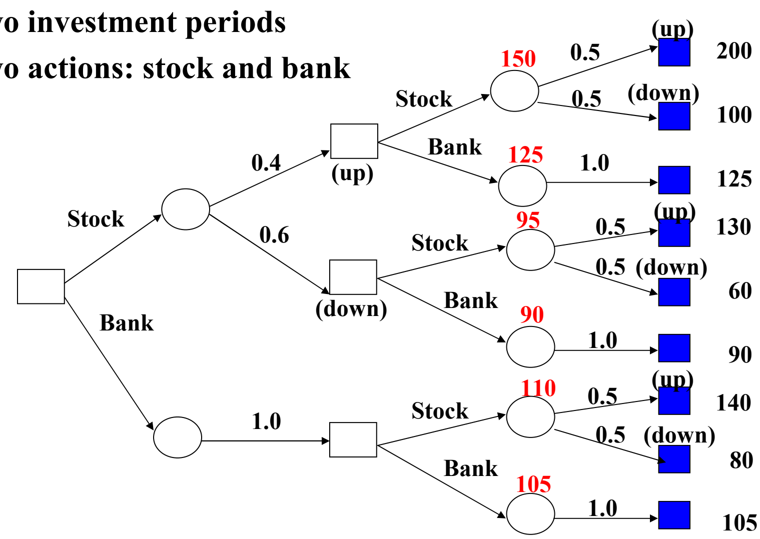

Multi-step problem example

Assume:

- Two investment periods
- Two actions: stock and bank

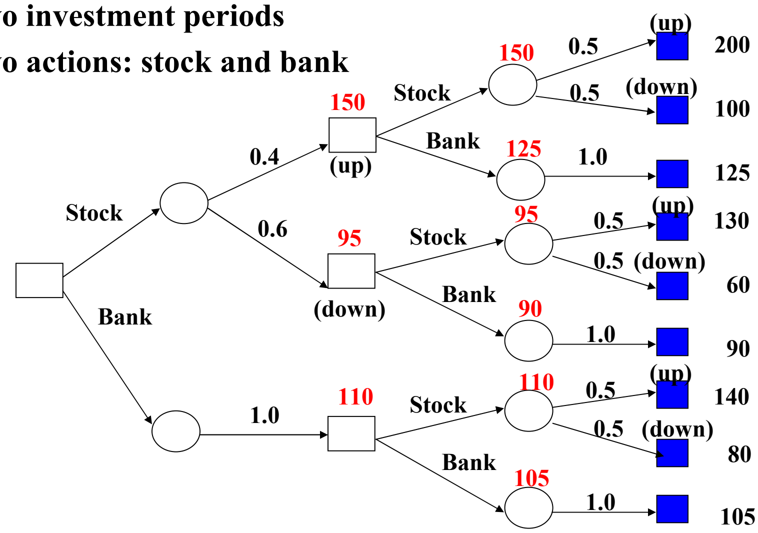

Multi-step problem example

Assume:

- Two investment periods
- Two actions: stock and bank

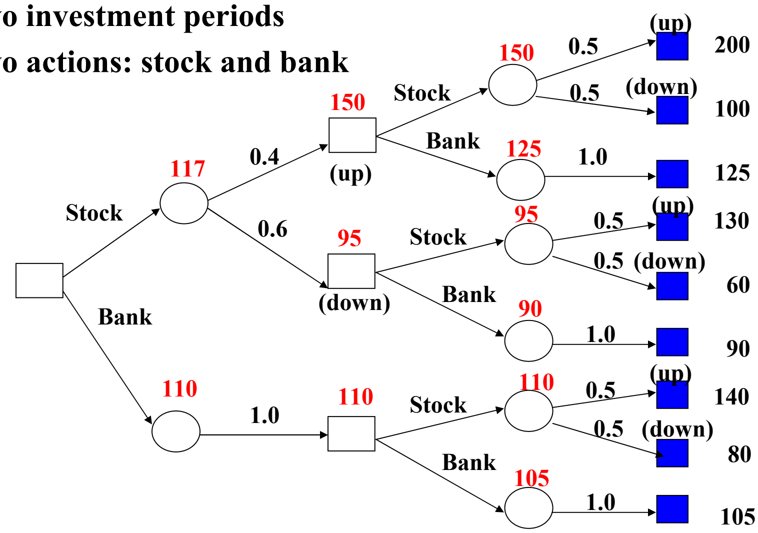

Multi-step problem example

Assume:

- **Two investment periods**
- **Two actions: stock and bank**