CMSC 671
Fall 2010

Thu 11/09/10

Decision Making Under Uncertainty
Chapter 16

Prof. Laura Zavala, laura.zavala@umbc.edu, ITE 373, 410-455-8775

Decision Making Under Uncertainty

- Decision theoretic agents
  - Combine
    - Probability theory
    - Utility theory
  - Make rational decisions based on
    - what it believes
    - what it wants

Decision Making Under Uncertainty

- Many environments have multiple possible outcomes
- Some of these outcomes may be good; others may be bad
- Some may be very likely; others unlikely

- What's a poor agent to do??

Probabilistic outcomes

- The probability of outcome *s'* given evidence observations *e*:
  - $P(Result(a)=s'|a,e)$
- Desirability of a state (agent's preferences):
  - $U(s)$
- Expected utility of an action given the evidence:
  - $EU(a|e) = \sum_{s'} P(Result(a)=s'|a,e)U(s')$
- Maximum expected utility:
  - $action = argmax_a \, EU(a|e)$

Non-Deterministic vs. Probabilistic Uncertainty

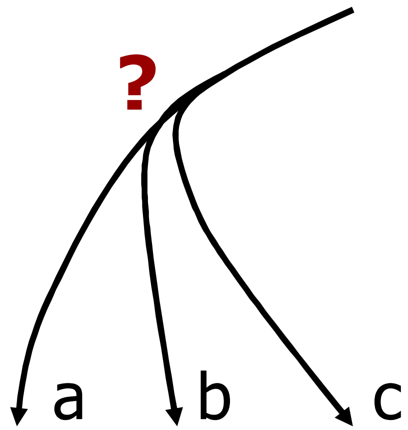

{a,b,c}

→ decision that is best for worst case

Non-deterministic model

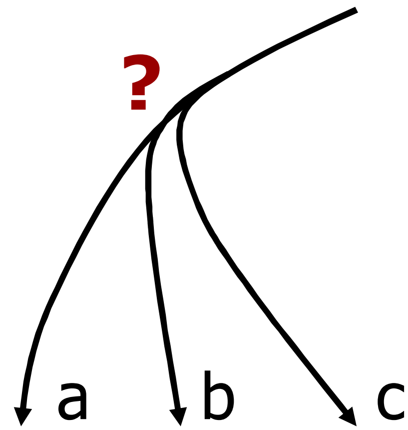

$\{a(p_a), b(p_b), c(p_c)\}$

→ decision that maximizes expected utility value

Probabilistic model

~ Adversarial search

Expected Utility

- Random variable X with n values $x_1,\ldots,x_n$ and distribution $(p_1,\ldots,p_n)$
  E.g.: X is the state reached after doing an action A under uncertainty
- Function U of X
  E.g., U is the utility of a state
- The expected utility of A is

$$EU[A] = \Sigma_{i=1,\ldots,n}\, p(x_i|A)U(x_i)$$

One State/One Action Example

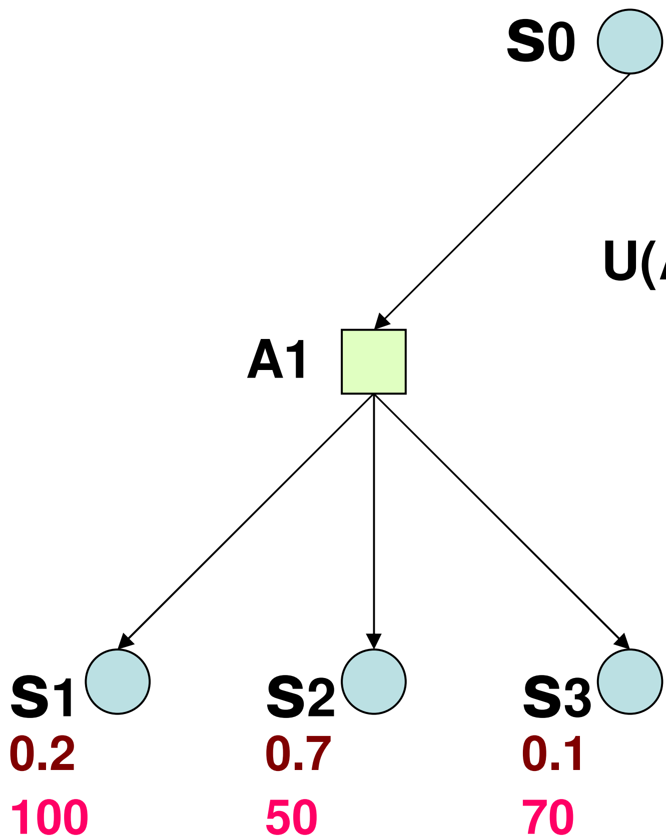

$$U(A1, S0) = 100 \times 0.2 + 50 \times 0.7 + 70 \times 0.1$$
$$= 20 + 35 + 7$$
$$= 62$$

One State/Two Actions Example

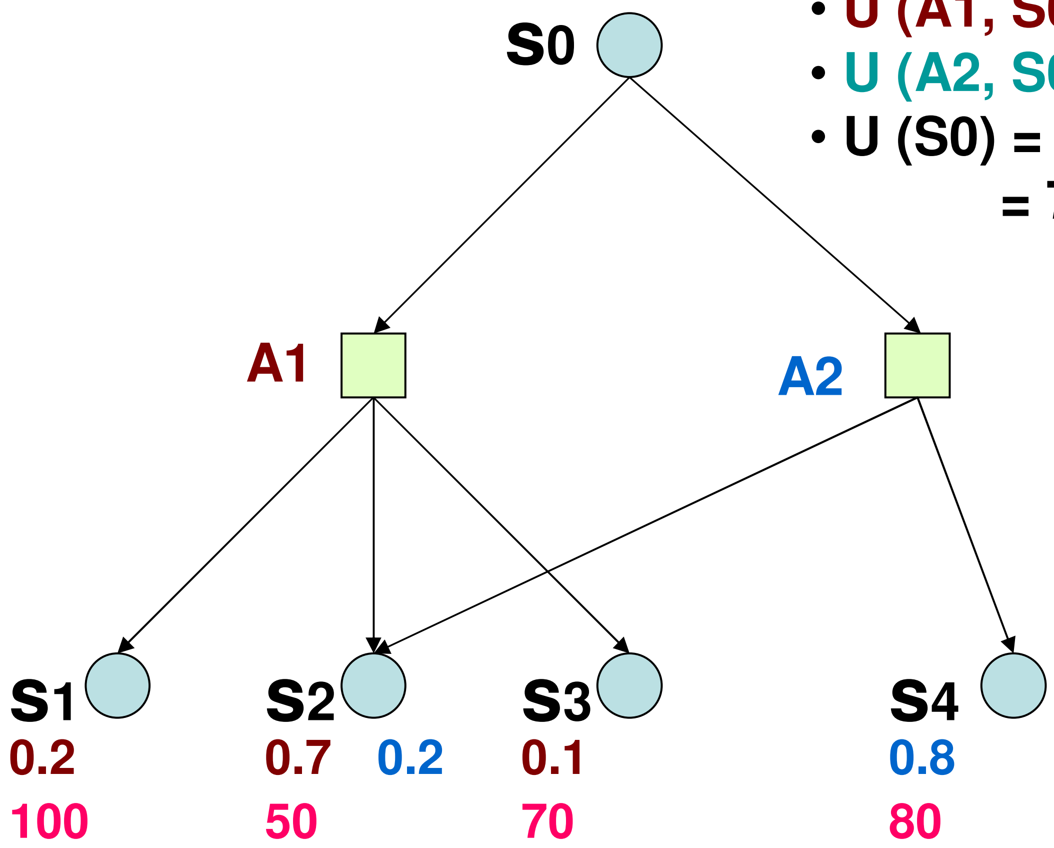

- U (A1, S0) = 62
- U (A2, S0) = 74
- U (S0) = $\max_a\{U(a,S0)\}$
  = 74

Introducing Action Costs

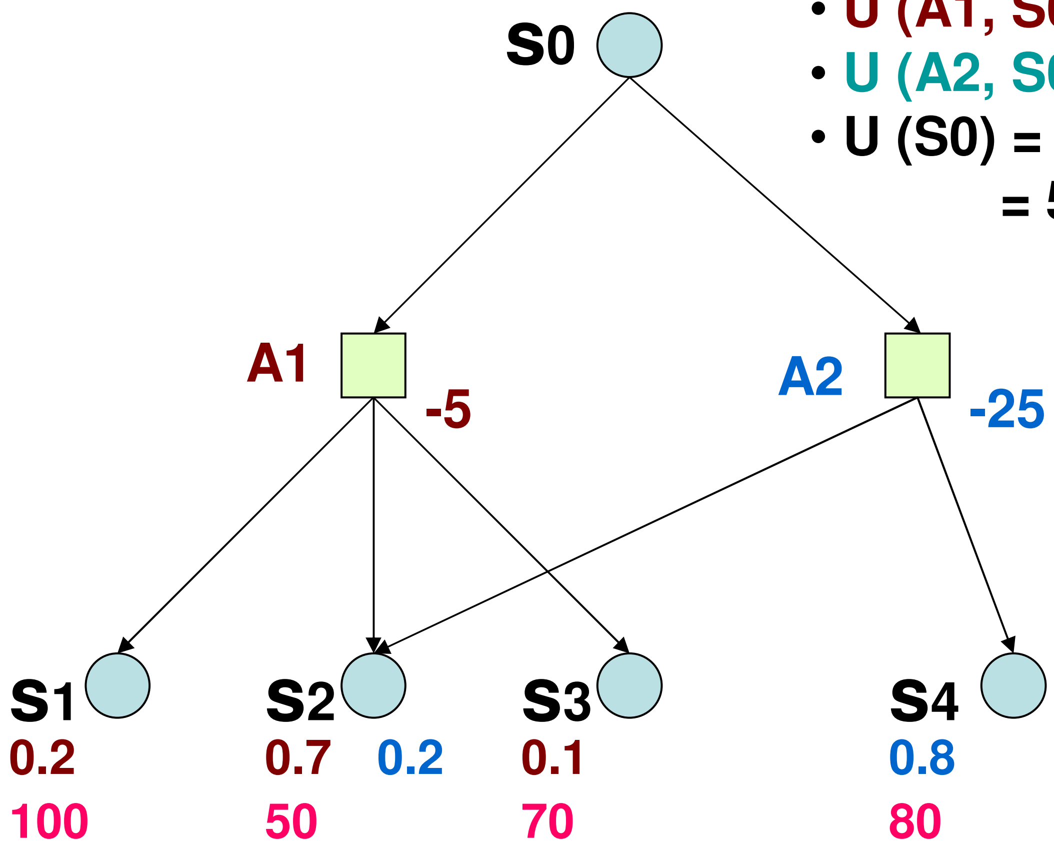

- U (A1, S0) = 62 – 5 = 57
- U (A2, S0) = 74 – 25 = 49
- U (S0) = $\max_a\{U(a, S0)\}$
  = 57

MEU Principle

- A **rational agent** should choose the action that maximizes agent's expected utility
- This is the basis of the field of **decision theory**
- The MEU principle provides a **normative criterion** for rational choice of action

- AI is solved!!!

Not quite…

- Must have a **complete** model of:
  - Actions
  - Utilities
  - States
- Even if you have a complete model, decision making is computationally **intractable**
- In fact, a truly rational agent takes into account the utility of reasoning as well (**bounded rationality**)
- Nevertheless, great progress has been made in this area recently, and we are able to solve much more complex decision-theoretic problems than ever before

Agent preferences (notation)

- A>B
  - The agent prefers A over B
- A~B
  - The agent is indifferent between A and B
- A>~B
  - The agent prefers A over B or is indifferent between them

Axioms of Utility Theory

- Orderability
  - $(A>B) \vee (A<B) \vee (A\sim B)$
- Transitivity
  - $(A>B) \wedge (B>C) \Rightarrow (A>C)$
- Continuity
  - $A>B>C \Rightarrow \exists p\ [p,A;\ 1-p,C] \sim B$
- Substitutability
  - $A\sim B \Rightarrow [p,A;\ 1-p,C]\sim[p,B;\ 1-p,C]$
- Monotonicity
  - $A>B \Rightarrow (p \geq q \Leftrightarrow [p,A;\ 1-p,B] \succsim [q,A;\ 1-q,B])$
  - Also for A~B
- Decomposability
  - $[p,A;\ 1-p,\ [q,B;\ 1-q,\ C]] \sim [p,A;\ (1-p)q,\ B;\ (1-p)(1-q),\ C]$
  - (figure above)

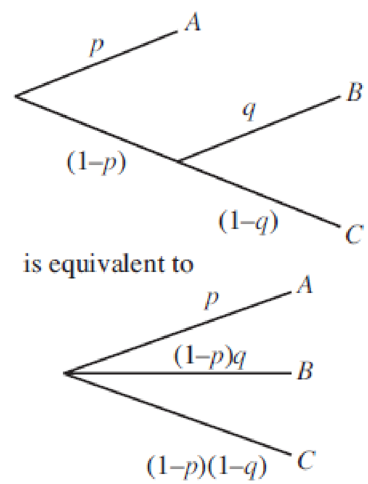

Value Function

- Provides a ranking of alternatives, but not a meaningful metric scale
- Also known as an "ordinal utility function"

- Sometimes, only relative judgments (value functions) are necessary
- At other times, absolute judgments (utility functions) are required

Money Versus Utility

- Money <> Utility
  - More money is better, but not always in a linear relationship to the amount of money
- Expected Monetary Value
  - TV game show: you have won $1,000,000
  - Take it or Gamble (flip coin to get $2,500,000)
  - EMV = .5($0) + .5(2,500,000) = 1,250,000
  - 1,250,000 > 1,000,000: rational agent should take it right?
  - Assign utilities (might be different for average people and billionares)

Money Versus Utility

- Money <> Utility
  - More money is better, but not always in a linear relationship to the amount of money
- Expected Monetary Value
- Risk-averse: $U(L) < U(S_{EMV(L)})$
- Risk-seeking: $U(L) > U(S_{EMV(L)})$
- Risk-neutral: $U(L) = U(S_{EMV(L)})$

Multiattribute Utility Theory

- A given state may have multiple utilities
  - ...because of multiple evaluation criteria
  - ...because of multiple agents (interested parties) with different utility functions
- We will talk about this more later in the semester, when we discuss multi-agent systems and game theory

Decision Networks

- Extend BNs to handle **actions** and **utilities**
- Also called *influence diagrams*
- Use BN inference methods to solve
- Perform *Value of Information* calculations

Decision Networks cont.

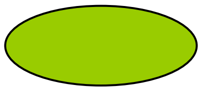

- Chance nodes: random variables, as in BNs

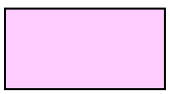

- Decision nodes: actions that a decision maker can take

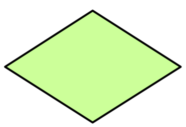

- Utility/value nodes: the utility of an outcome state

R&N example

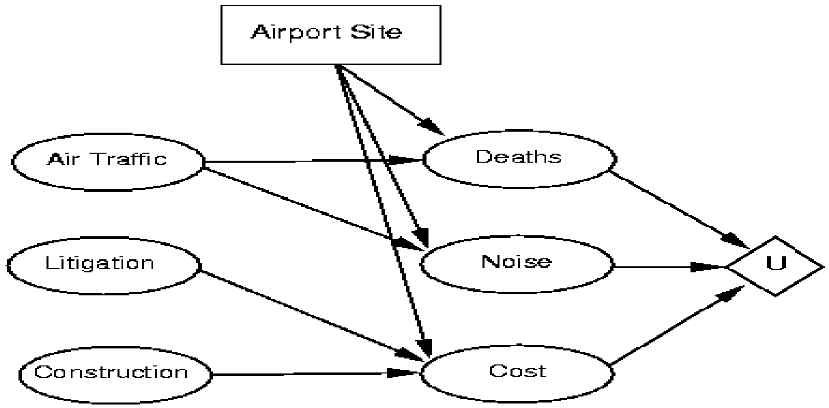

- Siting a new airport requires consideration of the disruption caused by construction; the cost of land; the distance from centers of population; the noise of flight operations; safety issues arising from local topography; weather conditions; and so on.

Umbrella Network

take/don't take

P(rain) = 0.4

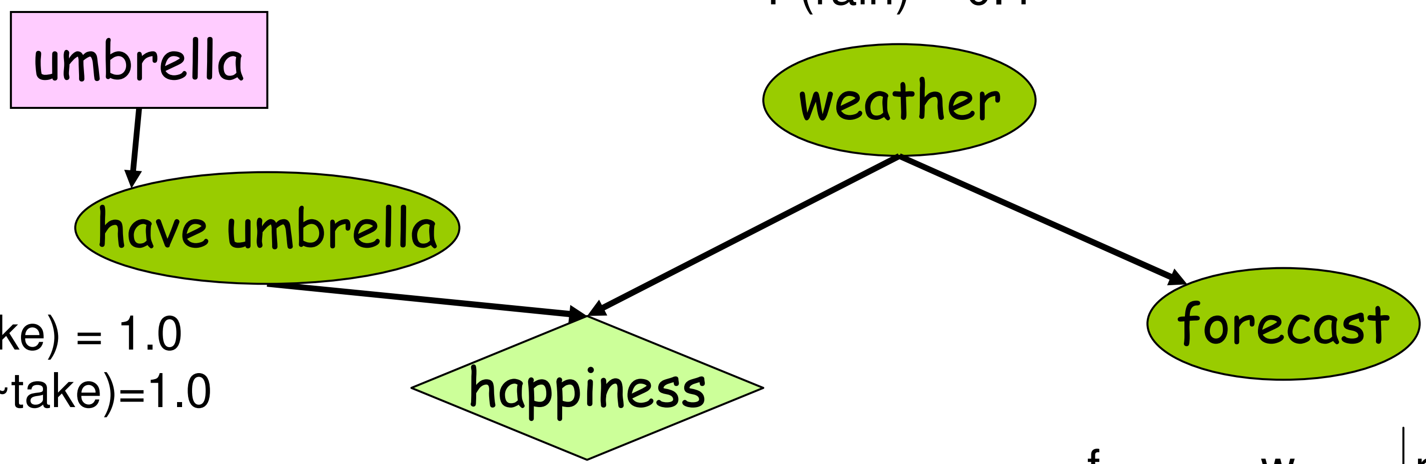

P(have|take) = 1.0
P(~have|~take)=1.0

U(have,rain) = -25
U(have,~rain) = 0
U(~have, rain) = -100
U(~have, ~rain) = 100

f	w	p(f|w)
sunny	rain	0.3
rainy	rain	0.7
sunny	no rain	0.8
rainy	no rain	0.2

Evaluating Decision Networks

- Set the evidence variables for current state
- For each possible value of the decision node:
  - Set decision node to that value
  - Calculate the posterior probability of the parent nodes of the utility node, using BN inference
  - Calculate the resulting utility for each action
- Return the action with the highest utility

Decision Making: Umbrella Network

Should I take my umbrella??

take/don't take

P(rain) = 0.4

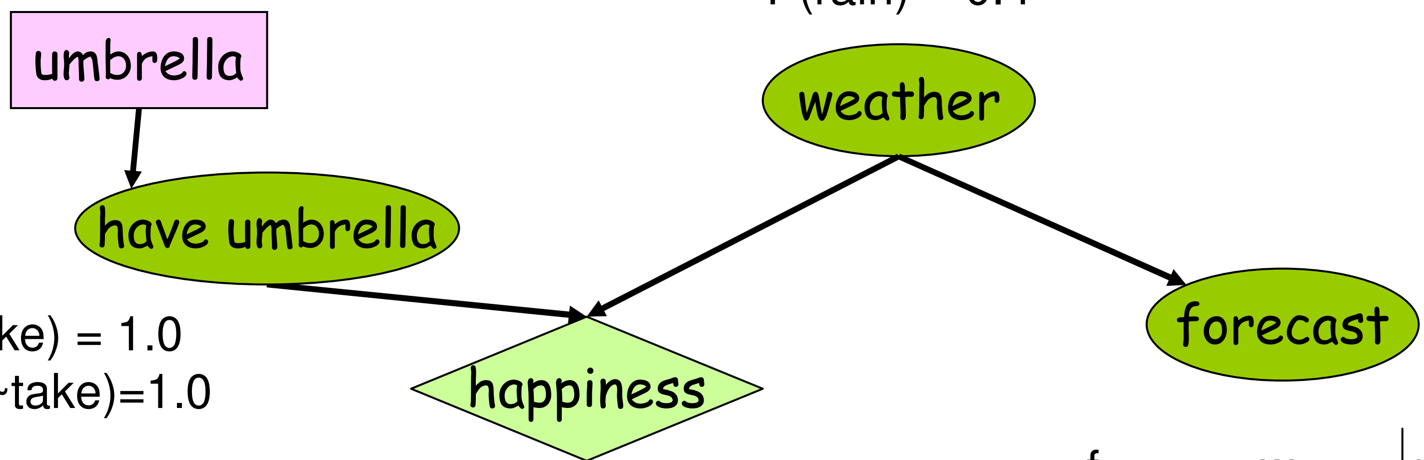

P(have|take) = 1.0
P(~have|~take)=1.0

U(have,rain) = -25
U(have,~rain) = 0
U(~have, rain) = -100
U(~have, ~rain) = 100

f	w	p(f|w)
sunny	rain	0.3
rainy	rain	0.7
sunny	no rain	0.8
rainy	no rain	0.2

Value of Information (VOI)

- Suppose an agent's current knowledge is e. The value of the current best action α is (slide 4):

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

- The value of the new best action after new evidence e' is obtained (E'=e'):

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

- The value of information for E' is therefore:

$$VOI(E') = \sum_k P(e_k \mid e) EU(\alpha_{e_k} \mid e_k, e) - EU(\alpha \mid e)$$

Value of Information: Umbrella Network

What is the value of knowing the weather forecast?

take/don't take

P(rain) = 0.4

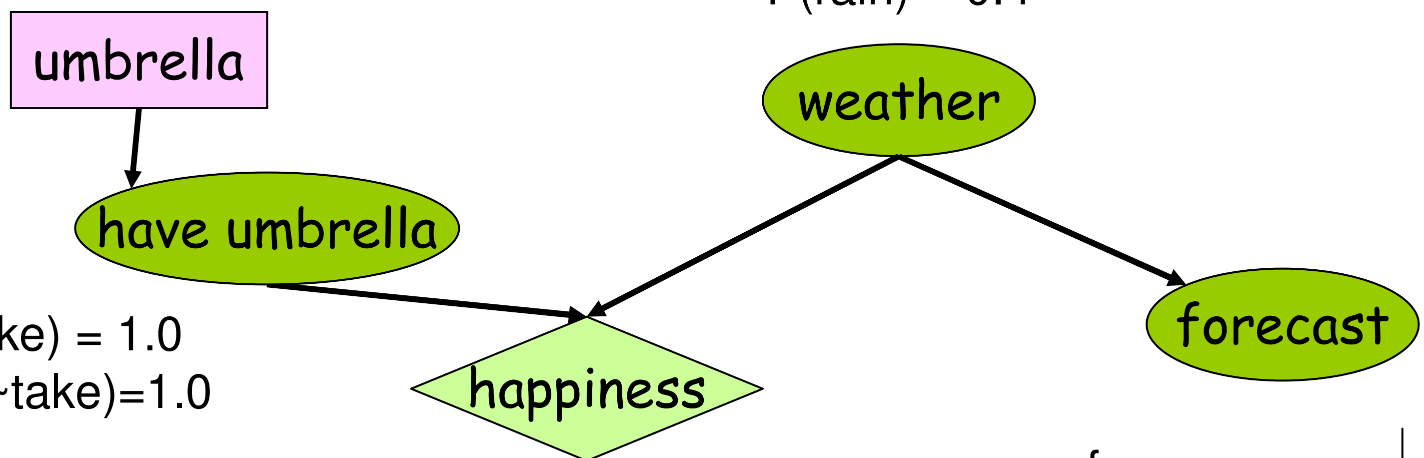

P(have|take) = 1.0
P(~have|~take)=1.0

U(have,rain) = -25
U(have,~rain) = 0
U(~have, rain) = -100
U(~have, ~rain) = 100

f	w	p(f|w)
sunny	rain	0.3
rainy	rain	0.7
sunny	no rain	0.8
rainy	no rain	0.2