


CS 188: Artificial Intelligence

Decision Networks and VPI




Instructor: Saagar Sanghavi—University of California, Berkeley

[Slides Credit: Dan Klein, Pieter Abbeel, Anca Dragan, and many others.]

Recap: Utilities and Rationality

- Utilities and Rationality
- Rational Preferences

Orderability: $(A > B) \vee (B > A) \vee (A \sim B)$
 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: $(A > B) \Rightarrow$
 $(p \geq q) \Leftrightarrow [p, A; 1-p, B] \geq [q, A; 1-q, B]$




MEU Principle:

Given Rational Preferences, Exists $U(X)$ s.t.

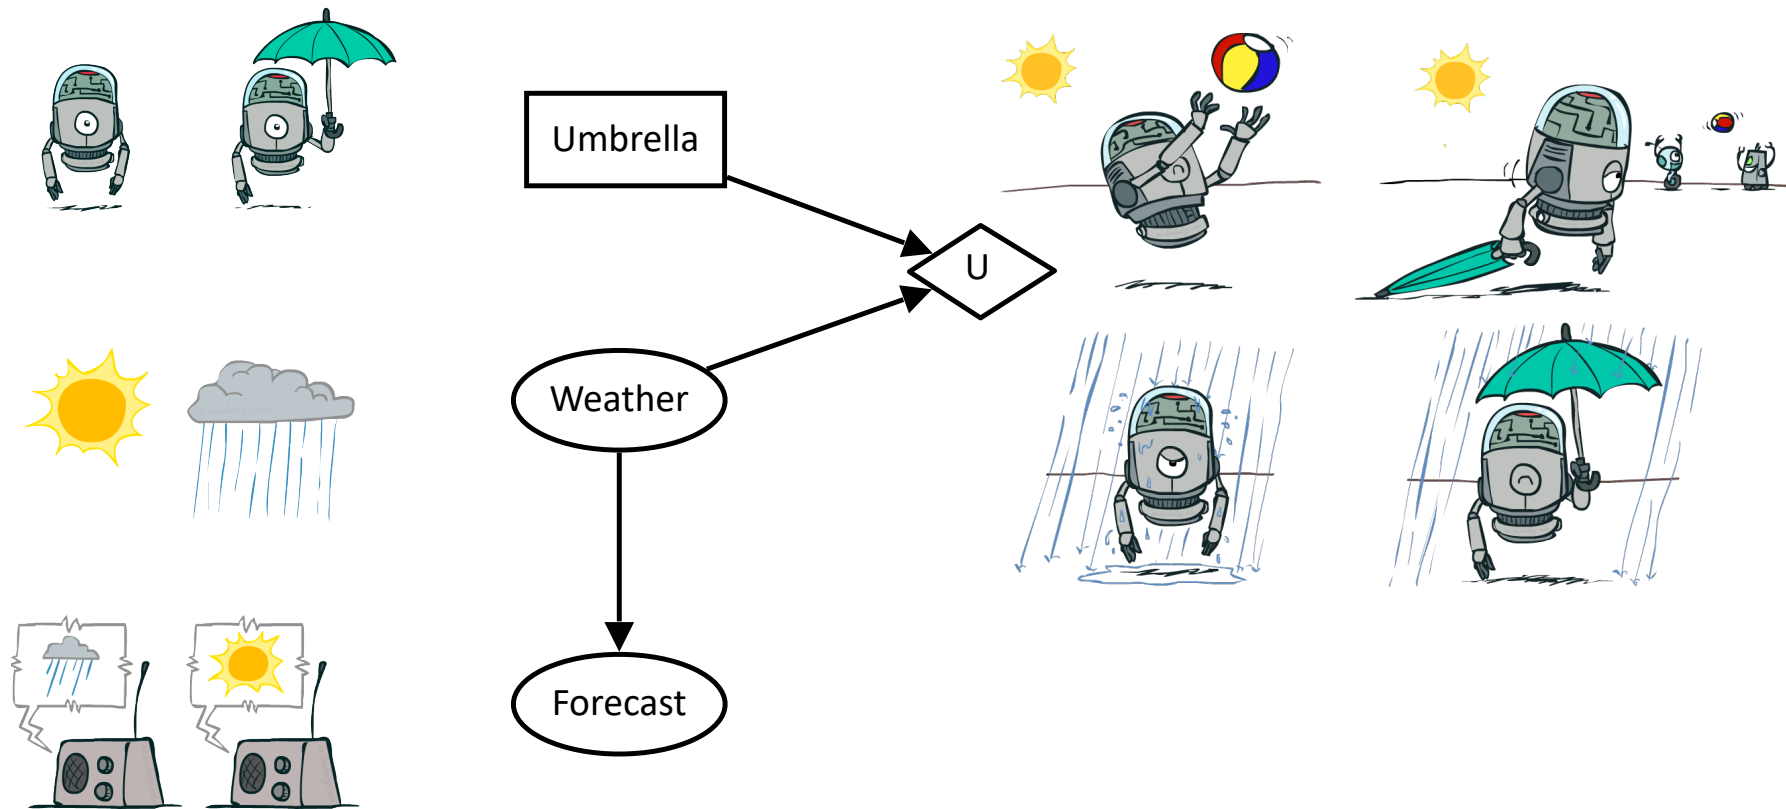
$$U(A) \geq U(B) \Leftrightarrow A \geq B$$

$$U([p_1, S_1; \dots; p_n, S_n]) = p_1 U(S_1) + \dots + p_n U(S_n)$$

Decision Networks



Decision Networks



Decision Networks

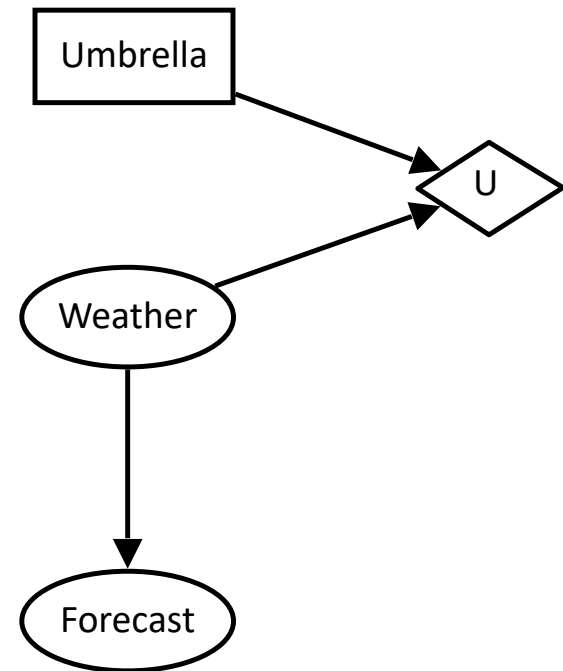
- **MEU: choose the action which maximizes the expected utility given the evidence**

- Can directly operationalize this with decision networks
 - Bayes nets, with new node types for utilities and actions
 - Lets us calculate the expected utility for each action

○ New node types:

□ ▪ Chance nodes (just like Bayes Nets)

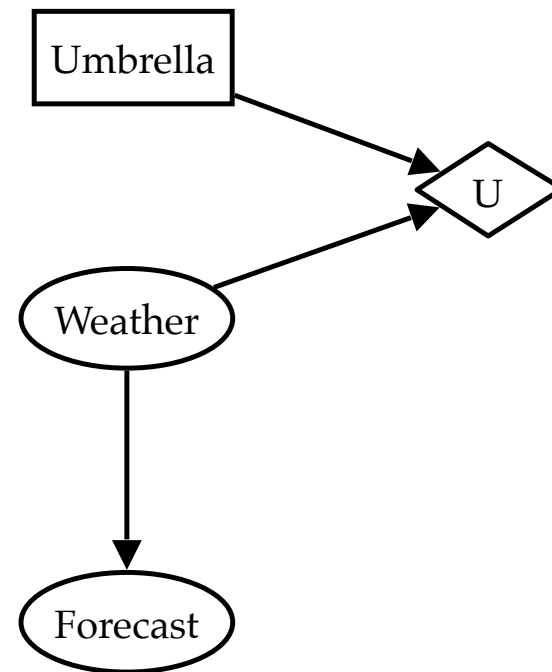
◇ ▪ Actions (rectangles, cannot have parents, act as observed evidence)



Decision Networks

- Action selection

- Instantiate all evidence
- Set action node(s) each possible way
- Calculate posterior for all parents of utility node, given the evidence
- Calculate expected utility for each action
- Choose maximizing action



Maximum Expected Utility

Umbrella = leave

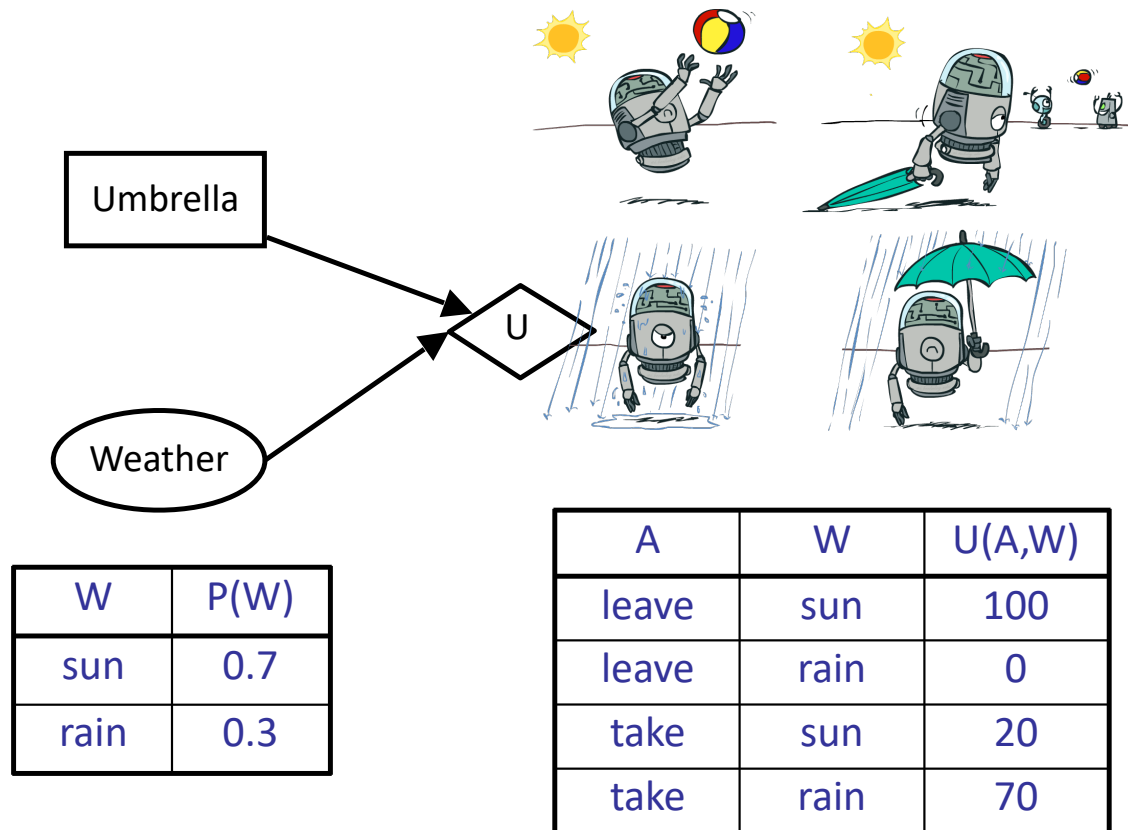
$$\begin{aligned} EU(\text{leave}) &= \sum_w P(w)U(\text{leave}, w) \\ &= 0.7 \cdot 100 + 0.3 \cdot 0 = 70 \end{aligned}$$

Umbrella = take


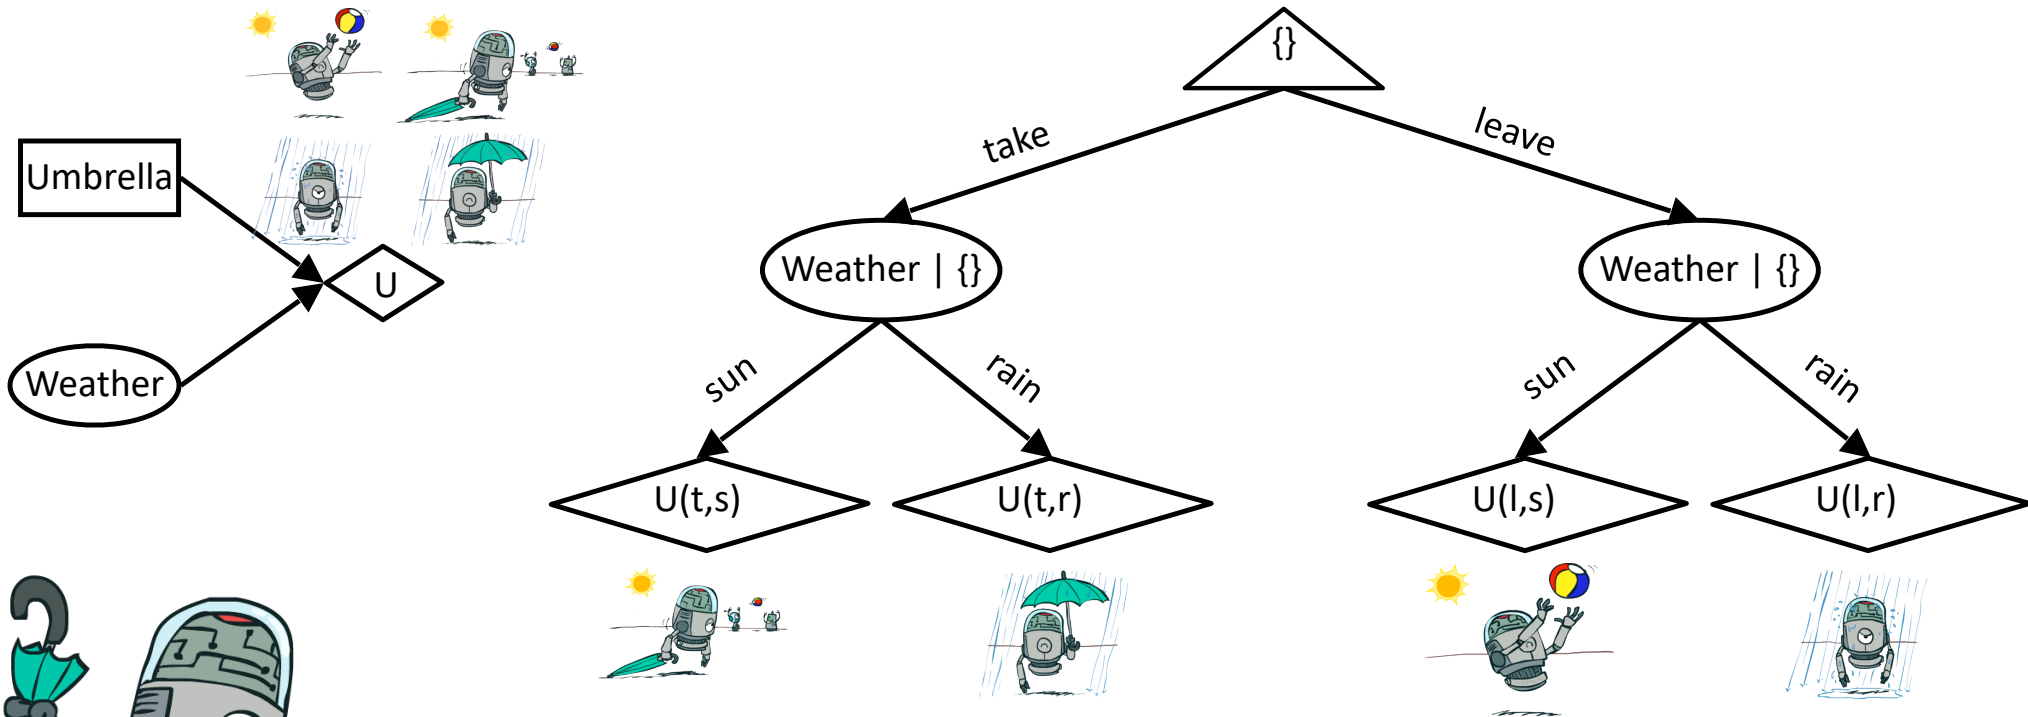
$$\begin{aligned} EU(\text{take}) &= \sum_w P(w)U(\text{take}, w) \\ &= 0.7 \cdot 20 + 0.3 \cdot 70 = 35 \end{aligned}$$

Optimal decision = leave

$$MEU(\emptyset) = \max_a EU(a) = 70$$



Decisions as Outcome Trees



- Almost exactly like expectimax

Maximum Expected Utility Given Evidence

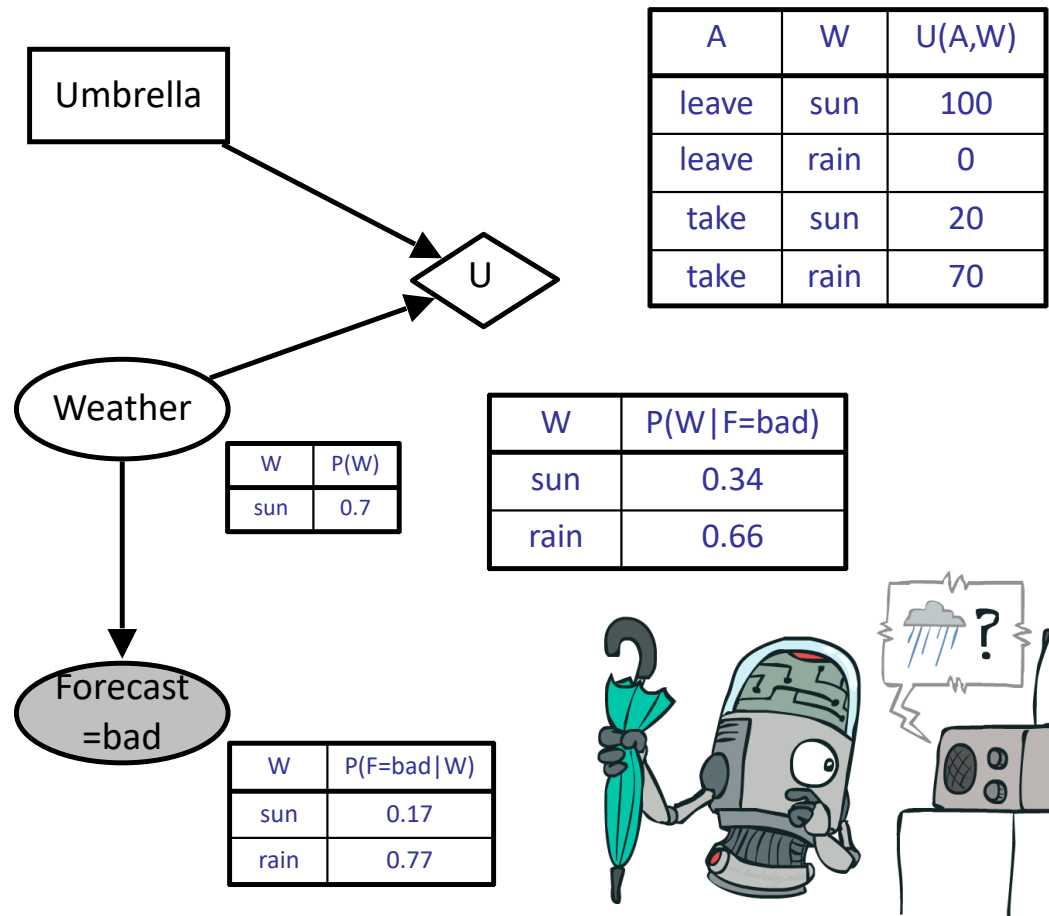
Umbrella = leave

$$EU(\text{leave}|\text{bad}) = \sum_w P(w|\text{bad})U(\text{leave}, w)$$

$$P(W) \quad P(F|W)$$

$$P(W|F) = \frac{P(W, F)}{\sum_w P(w, F)}$$

$$= \frac{P(F|W)P(W)}{\sum_w P(F|w)P(w)}$$



Maximum Expected Utility Given Evidence

Umbrella = leave

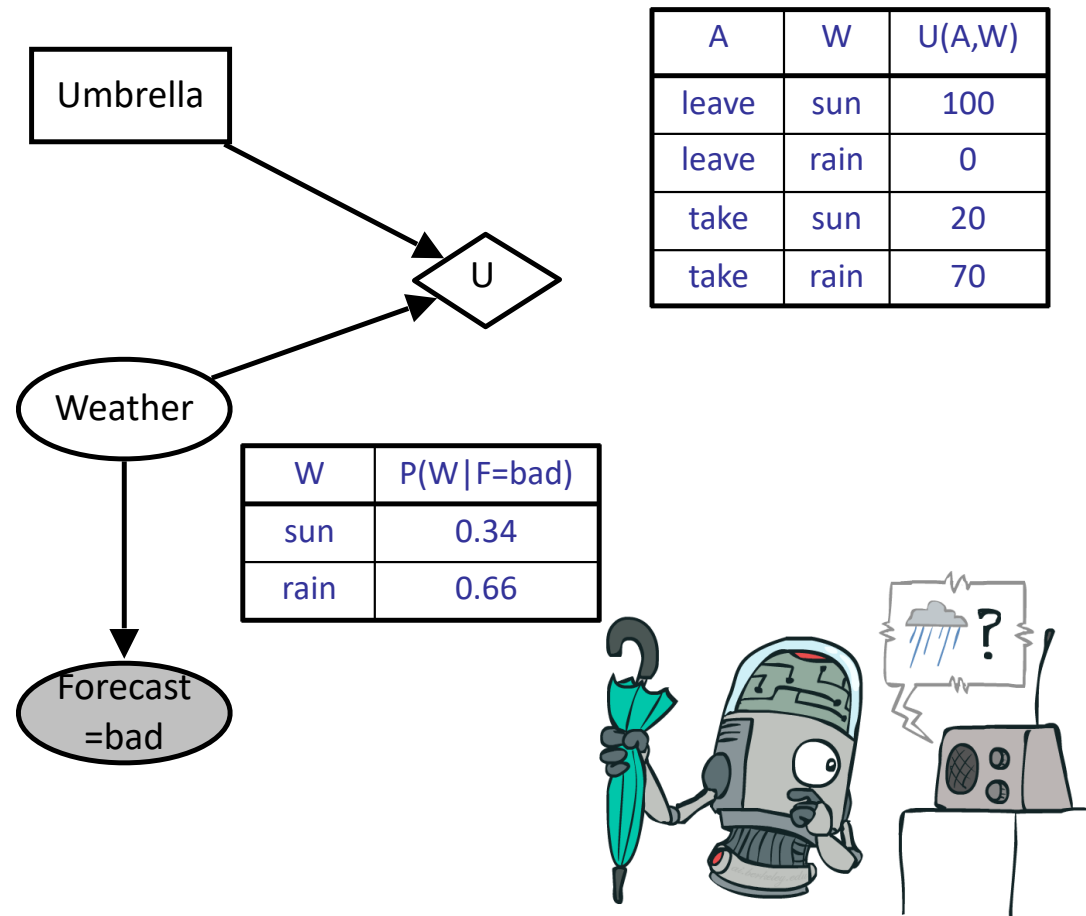
$$\begin{aligned} EU(\text{leave}|\text{bad}) &= \sum_w P(w|\text{bad})U(\text{leave}, w) \\ &= 0.34 \cdot 100 + 0.66 \cdot 0 = 34 \end{aligned}$$

Umbrella = take



$$\begin{aligned} EU(\text{take}|\text{bad}) &= \sum_w P(w|\text{bad})U(\text{take}, w) \\ &= 0.34 \cdot 20 + 0.66 \cdot 70 = 53 \end{aligned}$$

Optimal decision = take


$$MEU(F = \text{bad}) = \max_a EU(a|\text{bad}) = 53$$



Decisions as Outcome Trees




Value of Information



Value of Information

- Idea: compute value of acquiring evidence
 - Can be done directly from decision network
- Example: buying oil drilling rights
 - Two blocks A and B, exactly one has oil, worth k
 - You can drill in one location
 - Prior probabilities 0.5 each, & mutually exclusive
 - Drilling in either A or B has $EU = k/2$, $MEU = k/2$
- Question: what's the **value of information** of O?
 - Value of knowing which of A or B has oil
 - Value is expected gain in MEU from new info
 - Survey may say "oil in a" or "oil in b," prob 0.5 each
 - If we know OilLoc, MEU is k (either way)
 - Gain in MEU from knowing OilLoc?
 - $VPI(OilLoc) = k/2$
 - Fair price of information: $k/2$



Value of Perfect Information

MEU with no evidence

$$MEU(\emptyset) = \max_a EU(a) = 70$$

MEU if forecast is bad

$$MEU(F = \text{bad}) = \max_a EU(a|\text{bad}) = 53$$

MEU if forecast is good

$$MEU(F = \text{good}) = \max_a EU(a|\text{good}) = 89.4$$


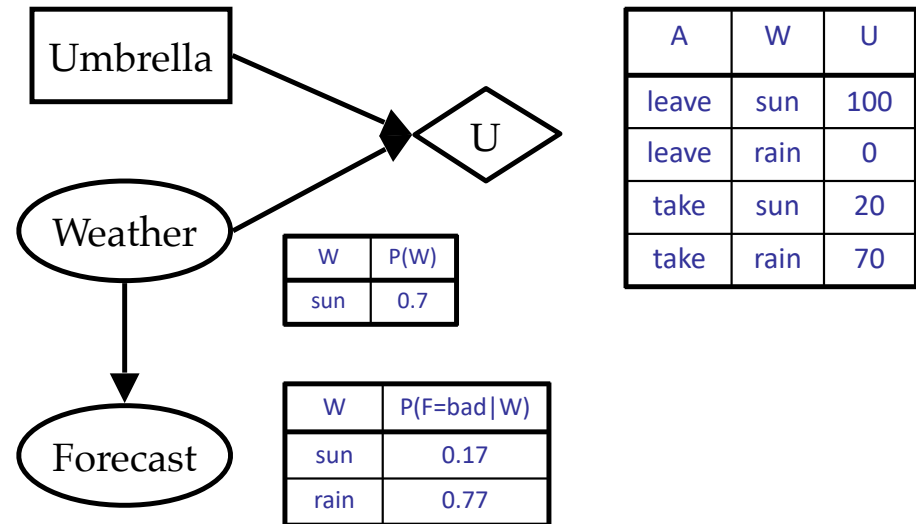
Forecast distribution

F	P(F)
bad	0.35
good	0.65



$$\begin{aligned} VPI &= 0.35 (53) + 0.65 (89.4) - 70 \\ VPI &= 6.66 \end{aligned}$$

$$VPI(E'|e) = \left(\sum_{e'} P(e'|e) MEU(e, e') \right) - MEU(e)$$



Value of Information

- Assume we have evidence $E=e$. Value if we act now:

$$MEU(e) = \max_a \sum_s P(s|e) U(s, a)$$




- We see new evidence $E' = e'$. Value if we act then:

$$MEU(e, e') = \max_a \sum_s P(s|e, e') U(s, a)$$

- BUT E' is a random variable whose value is **unknown**, so we don't know what e' will be.

- Expected $MEU(e, E') = \sum_{e'} P(e'|e) MEU(e, e')$

- Value $VPI(E'|e) = MEU(e, E') - MEU(e)$
by revealing E' first then acting, as opposed to acting now:



VPI Properties

- Nonnegative

$$\forall E', e : \text{VPI}(E'|e) \geq 0$$



- Nonadditive

(think of observing E_j twice)


$$\text{VPI}(E_j, E_k|e) \neq \text{VPI}(E_j|e) + \text{VPI}(E_k|e)$$



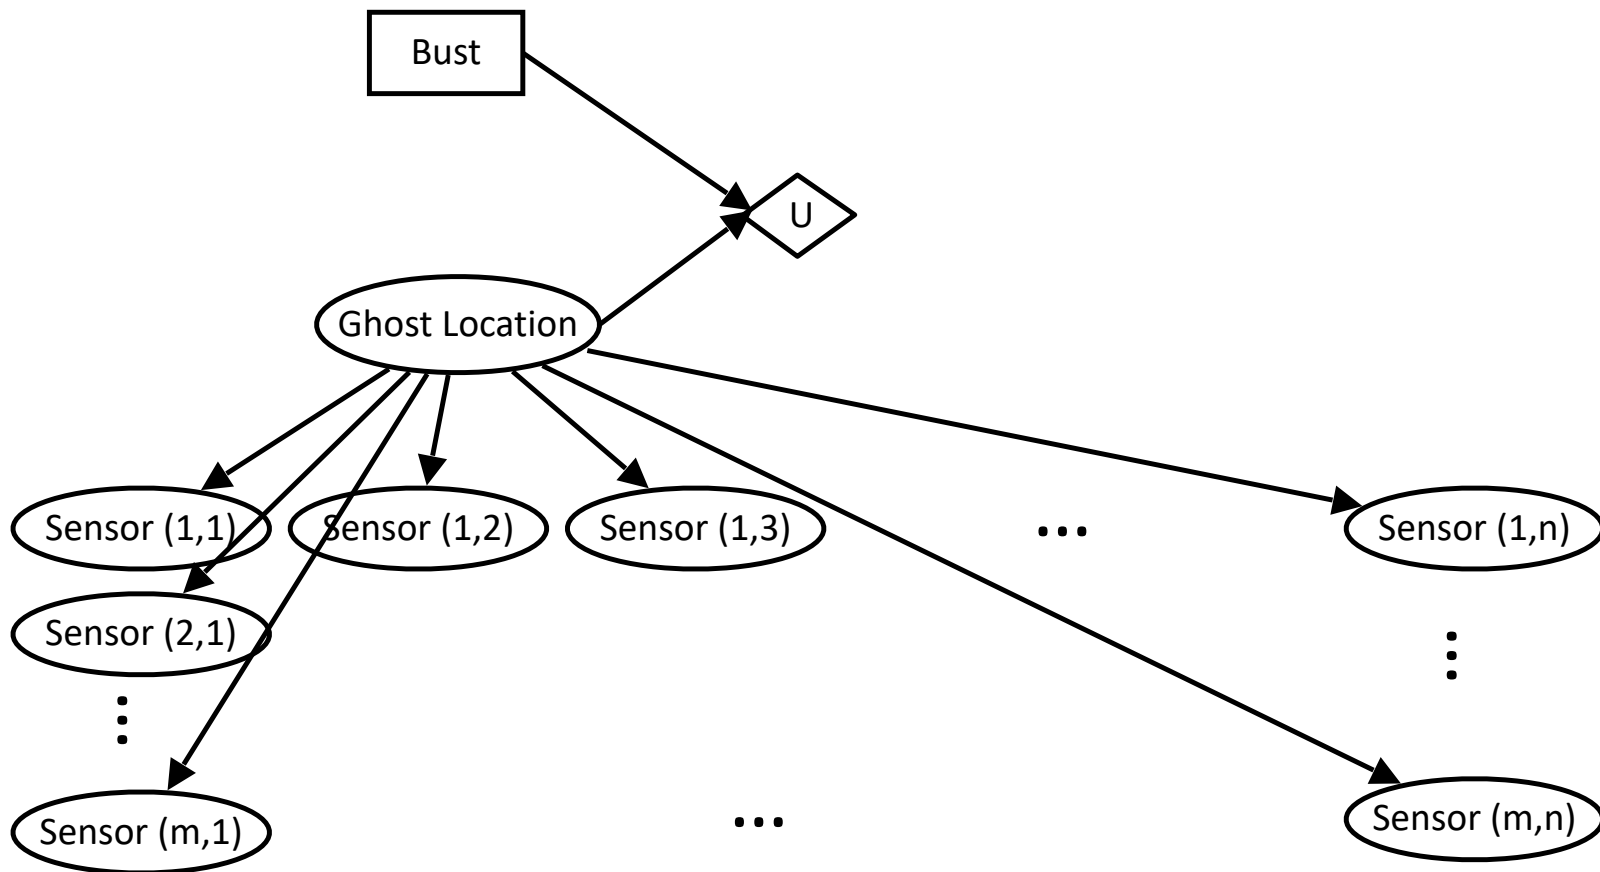
- Order independent

$$\text{VPI}(E_j, E_k|e) = \text{VPI}(E_j|e) + \text{VPI}(E_k|e, E_j)$$

$$= \text{VPI}(E_k|e) + \text{VPI}(E_j|e, E_k)$$





Ghostbusters Decision Network



Quick VPI Questions

- The soup of the day is either clam chowder or split pea, but you wouldn't order either one. What's the value of knowing which it is?
- There are two kinds of plastic forks at a picnic. One kind is slightly sturdier. What's the value of knowing which?
- You're playing the lottery. The prize will be \$0 or \$100. You can play any number between 1 and 100 (chance of winning is 1%). What is the value of knowing the winning number?



Value of Imperfect Information?





- No such thing
- Information corresponds to the observation of a node in the decision network
- If data is “noisy” that just means we don’t observe the original variable, but another variable which is a noisy version of the original one

VPI Question

- VPI(OilLoc) ?
- VPI(ScoutingReport) ?
- VPI(Scout) ?
- VPI(Scout | ScoutingReport) ?

○ Generally:

If $\text{Parents}(U) \quad Z \mid \text{CurrentEvidence}$
Then $\text{VPI}(Z \mid \text{CurrentEvidence}) = 0$



||