

# CS886: Influence diagrams


- January 25<sup>th</sup>
- Intro to utility theory and decision theory
- Influence diagrams (or decision networks)
- Decision making under uncertainty

# Preference Orderings

- A *preference ordering*  $\succsim$  is a ranking of all possible states of affairs (worlds)  $S$ 
  - these could be outcomes of actions, truth assts, states in a search problem, etc.
  - $s \succsim t$ : means that state  $s$  is *at least as good as*  $t$
  - $s \succ t$ : means that state  $s$  is *strictly preferred to*  $t$
- We insist that  $\succsim$  is
  - reflexive: i.e.,  $s \succsim s$  for all states  $s$
  - transitive: i.e., if  $s \succsim t$  and  $t \succsim w$ , then  $s \succsim w$
  - connected: for all states  $s, t$ , either  $s \succsim t$  or  $t \succsim s$

# Why Impose These Conditions?

- Structure of preference ordering imposes certain “rationality requirements” (it is a weak ordering)
- E.g., why transitivity?
  - Suppose you (strictly) prefer coffee to tea, tea to OJ, OJ to coffee
  - If you prefer X to Y, you’ll trade me Y plus \$1 for X
  - I can construct a “money pump” and extract arbitrary amounts of money from you



# Utilities

- Rather than just ranking outcomes, we must quantify our degree of preference
  - e.g., how much more important is **chc** than **~mess**
- A *utility function*  $U:S \rightarrow \mathbb{R}$  associates a real-valued *utility* with each outcome.
  - $U(s)$  measures your *degree* of preference for  $s$
- Note:  $U$  induces a preference ordering  $\succsim_U$  over  $S$  defined as:  $s \succsim_U t$  iff  $U(s) \geq U(t)$ 
  - obviously  $\succsim_U$  will be reflexive, transitive, connected

# Expected Utility

- Under conditions of uncertainty, each decision  $d$  induces a distribution  $\text{Pr}_d$  over possible outcomes
  - $\text{Pr}_d(s)$  is probability of outcome  $s$  under decision  $d$
- The *expected utility* of decision  $d$  is defined

$$EU(d) = \sum_{s \in S} \text{Pr}_d(s) U(s)$$

# The MEU Principle


- The *principle of maximum expected utility (MEU)* states that the optimal decision under conditions of uncertainty is that with the greatest expected utility.
- In our example
  - if my utility function is the first one, my robot should get coffee
  - if your utility function is the second one, your robot should do nothing

# Decision Problems: Uncertainty


- A *decision problem under uncertainty* is:
  - a set of *decisions*  $D$
  - a set of *outcomes* or states  $S$
  - an *outcome function*  $\text{Pr} : D \rightarrow \Delta(S)$ 
    - $\Delta(S)$  is the set of distributions over  $S$  (e.g.,  $\text{Pr}_d$ )
  - a *utility function*  $U$  over  $S$
- A *solution* to a decision problem under uncertainty is any  $d^* \in D$  such that  $\text{EU}(d^*) \succcurlyeq \text{EU}(d)$  for all  $d \in D$
- Again, for single-shot problems, this is trivial

# Expected Utility: Notes

- Note that this viewpoint accounts for both:
  - uncertainty in action outcomes
  - uncertainty in state of knowledge
  - any combination of the two



Stochastic actions



Uncertain knowledge



# Expected Utility: Notes

- Why MEU? Where do utilities come from?
  - underlying foundations of utility theory tightly couple utility with action/choice
  - a utility function can be determined by asking someone about their preferences for actions in specific scenarios (or “lotteries” over outcomes)
- Utility functions needn't be unique
  - if I multiply  $U$  by a positive constant, all decisions have same relative utility
  - if I add a constant to  $U$ , same thing
  - *$U$  is unique up to positive affine transformation*

# So What are the Complications?

## ■ Outcome space is large

- like all of our problems, states spaces can be huge
- don't want to spell out distributions like  $\Pr_d$  explicitly
- Soln: Bayes nets (or related: *influence diagrams*)

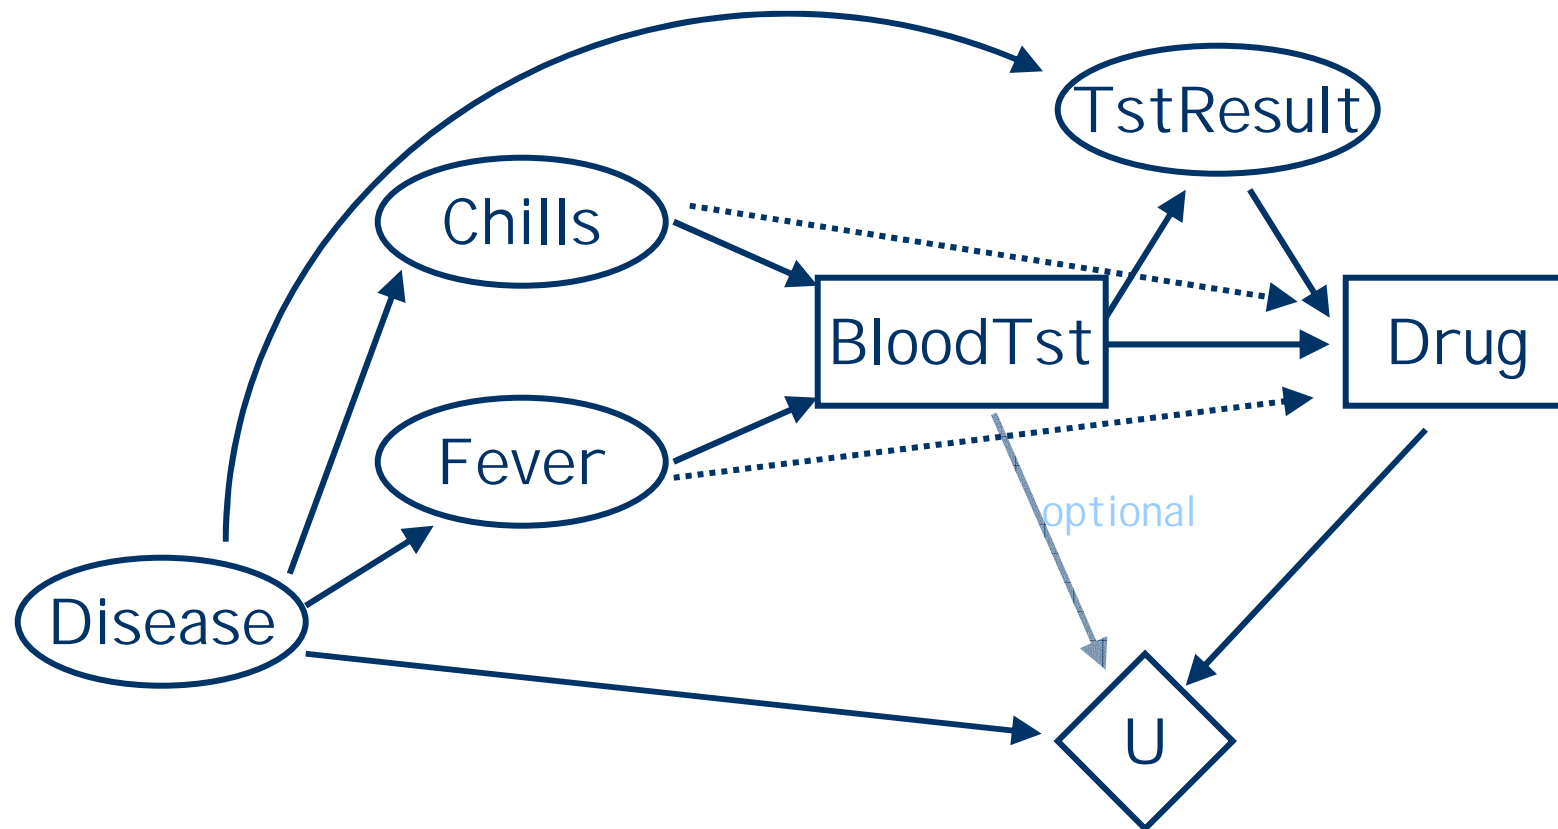
## ■ Decision space is large

- usually our decisions are not one-shot actions
- rather they involve sequential choices (like plans)
- if we treat each plan as a distinct decision, decision space is too large to handle directly
- Soln: use dynamic programming methods to *construct* optimal plans (actually generalizations of plans, called policies... like in game trees)

# Decision Networks

- *Decision networks* (more commonly known as *influence diagrams*) provide a way of representing sequential decision problems
  - basic idea: represent the variables in the problem as you would in a BN
  - add decision variables – variables that you “control”
  - add utility variables – how good different states are

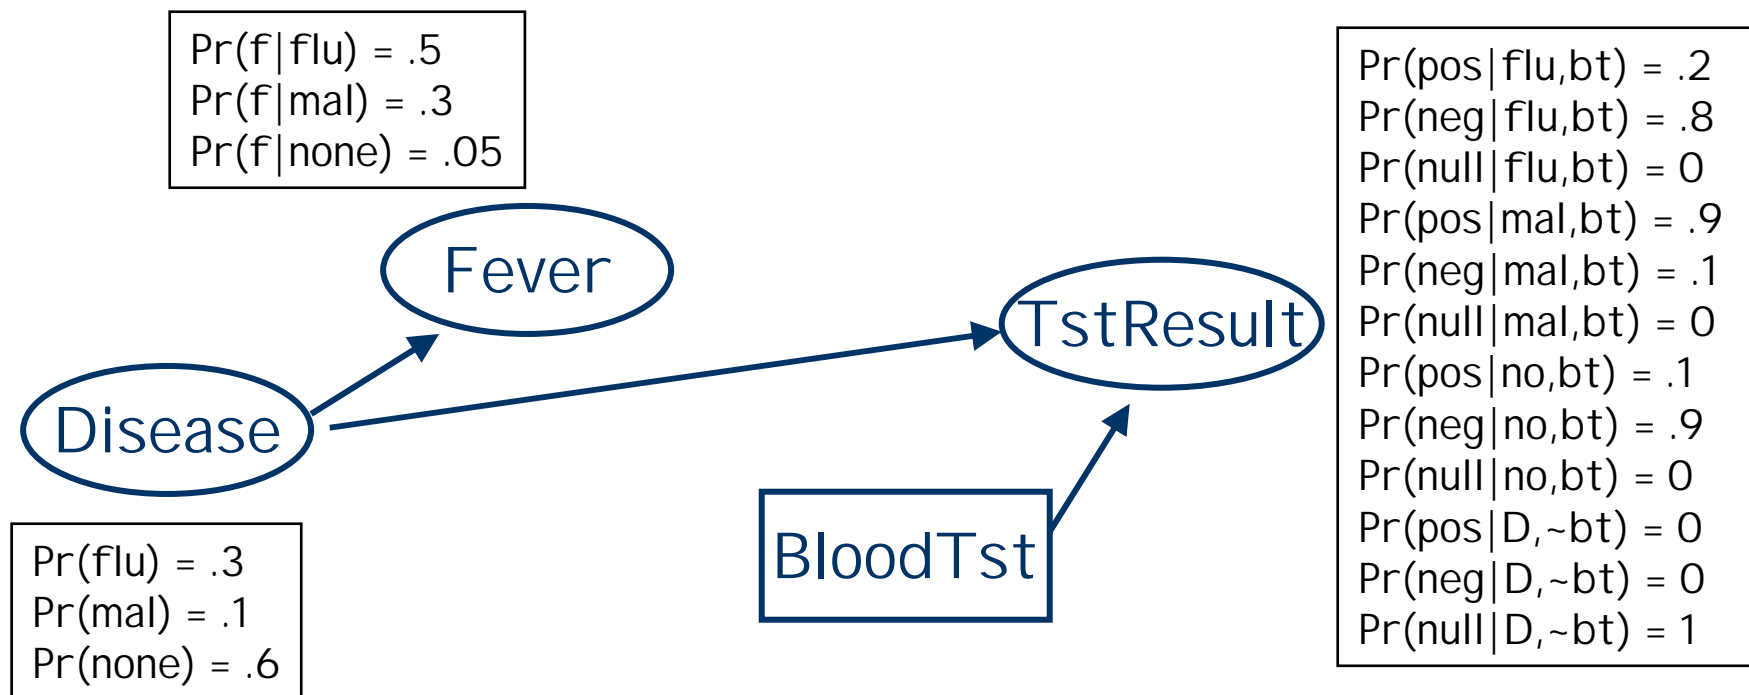
# Sample Decision Network



# Decision Networks: Chance Nodes

## ■ Chance nodes

- random variables, denoted by circles
- as in a BN, probabilistic dependence on parents



# Decision Networks: Decision Nodes

## ■ Decision nodes

- variables decision maker sets, denoted by squares
  - parents reflect *information available* at time decision is to be made
- In example decision node: the actual values of Ch and Fev will be observed before the decision to take test must be made
- agent can make *different decisions* for each instantiation of parents (i.e., policies)




# Decision Networks: Value Node

## ■ Value node

- specifies utility of a state, denoted by a diamond
- utility depends *only on state of parents* of value node
- generally: only one value node in a decision network


## ■ Utility depends only on disease and drug



$U(\text{fludrug}, \text{flu}) = 20$   
 $U(\text{fludrug}, \text{mal}) = -300$   
 $U(\text{fludrug}, \text{none}) = -5$   
 $U(\text{maldrug}, \text{flu}) = -30$   
 $U(\text{maldrug}, \text{mal}) = 10$   
 $U(\text{maldrug}, \text{none}) = -20$   
 $U(\text{no drug}, \text{flu}) = -10$   
 $U(\text{no drug}, \text{mal}) = -285$   
 $U(\text{no drug}, \text{none}) = 30$

# Decision Networks: Assumptions


- Decision nodes are totally ordered
  - decision variables  $D_1, D_2, \dots, D_n$
  - decisions are made in sequence
  - e.g., BloodTst (yes,no) decided before Drug (fd,md,no)
- *No-forgetting property*
  - any information available when decision  $D_i$  is made is available when decision  $D_j$  is made (for  $i < j$ )
  - thus all parents of  $D_i$  are parents of  $D_j$





# Policies

- Let  $Par(D_i)$  be the parents of decision node  $D_i$ 
  - $Dom(Par(D_i))$  is the set of assignments to parents
- A policy  $\delta$  is a set of mappings  $\delta_i$ , one for each decision node  $D_i$ 
  - $\delta_i : Dom(Par(D_i)) \rightarrow Dom(D_i)$
  - $\delta_i$  associates a decision with each parent asst for  $D_i$
- For example, a policy for BT might be:
  - $\delta_{BT}(c, f) = bt$
  - $\delta_{BT}(c, \sim f) = \sim bt$
  - $\delta_{BT}(\sim c, f) = bt$
  - $\delta_{BT}(\sim c, \sim f) = \sim bt$



# Value of a Policy

- *Value of a policy*  $\delta$  is the expected utility given that decision nodes are executed according to  $\delta$
- Given asst  $\mathbf{x}$  to the set  $\mathbf{X}$  of all chance variables, let  $\delta(\mathbf{x})$  denote the asst to decision variables dictated by  $\delta$ 
  - e.g., asst to  $D_1$  determined by it's parents' asst in  $\mathbf{x}$
  - e.g., asst to  $D_2$  determined by it's parents' asst in  $\mathbf{x}$  along with whatever was assigned to  $D_1$
  - etc.
- Value of  $\delta$  :

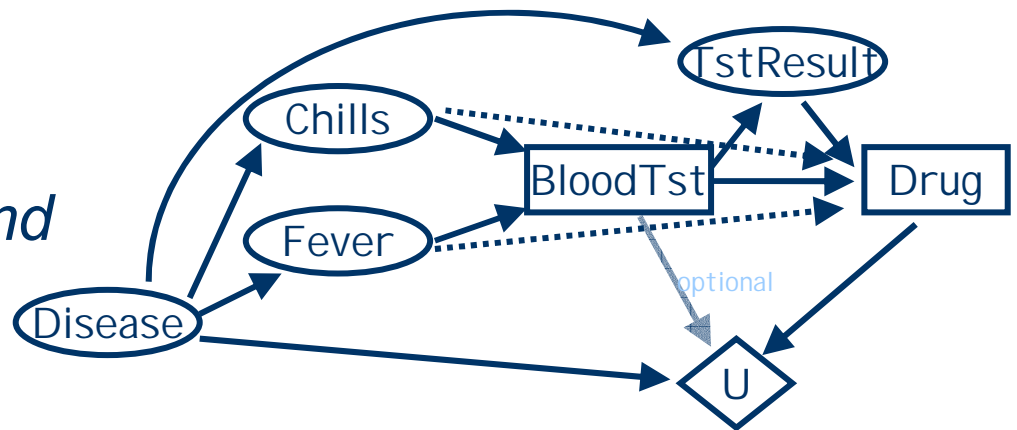
$$EU(\delta) = \sum_{\mathbf{x}} P(\mathbf{X}, \delta(\mathbf{X})) U(\mathbf{X}, \delta(\mathbf{X}))$$

# Optimal Policies

- An *optimal policy* is a policy  $\delta^*$  such that  $EU(\delta^*) \geq EU(\delta)$  for all policies  $\delta$
- We can use the dynamic programming principle yet again to avoid enumerating all policies
- We can also use the structure of the decision network to use variable elimination to aid in the computation

# Computing the Best Policy

- We can work backwards as follows
- First compute optimal policy for Drug (last dec'n)
  - for each asst to parents (C,F,BT,TR) and for each decision value (D = md,fd,none), *compute the expected value* of choosing that value of D
  - set policy choice for each value of parents to be the value of D that has max value
  - eg:  $\delta_D(c,f,bt,pos) = md$

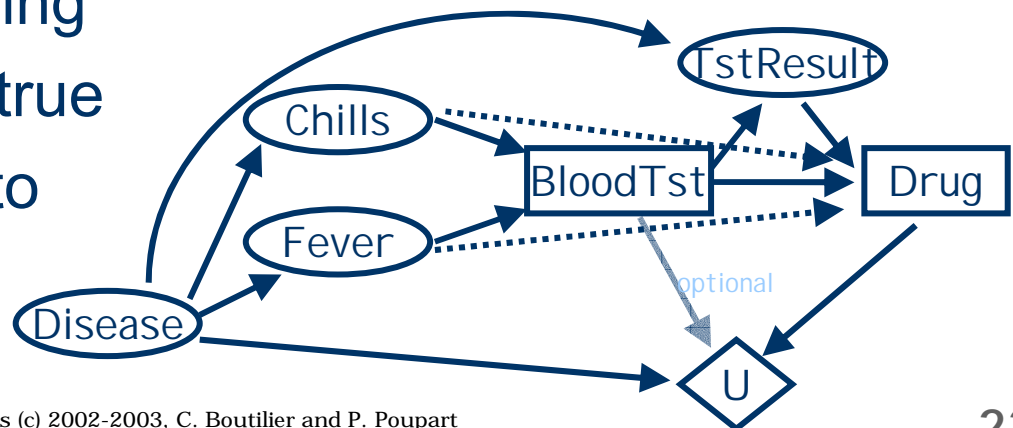


# Computing the Best Policy

- Next compute policy for BT given policy  $\delta_D(C, F, BT, TR)$  just determined for Drug
  - since  $\delta_D(C, F, BT, TR)$  is fixed, we can treat Drug as a normal random variable with deterministic probabilities
  - i.e., for any instantiation of parents, value of Drug is fixed by policy  $\delta_D$
  - this means we can solve for optimal policy for BT just as before
  - only uninstantiated vars are random vars (once we fix *its* parents)

# Computing the Best Policy

- How do we compute these expected values?
  - suppose we have asst  $\langle c, f, bt, pos \rangle$  to parents of *Drug*
  - we want to compute EU of deciding to set  $Drug = md$
  - we can run variable elimination!
- Treat  $C, F, BT, TR, Dr$  as evidence
  - this reduces factors (e.g.,  $U$  restricted to  $bt, md$ : depends on *Dis*)
  - eliminate remaining variables (e.g., only *Disease* left)
  - left with factor:  $U() = \sum_{Dis} P(Dis | c, f, bt, pos, md) U(Dis)$
- We now know EU of doing  $Dr=md$  when  $c, f, bt, pos$  true
- Can do same for  $fd, no$  to decide which is best




# Computing Expected Utilities

- The preceding illustrates a general phenomenon
  - computing expected utilities with BNs is quite easy
  - utility nodes are just factors that can be dealt with using variable elimination

$$\begin{aligned} EU &= \sum_{A,B,C} P(A,B,C) U(B,C) \\ &= \sum_{A,B,C} P(C|B) P(B|A) P(A) U(B,C) \end{aligned}$$

- Just eliminate variables in the usual way



# Optimizing Policies: Key Points

- If a decision node  $D$  has no decisions that follow it, we can find its policy by instantiating each of its parents and computing the expected utility of each decision for each parent instantiation
  - no-forgetting means that all other decisions are instantiated (they must be parents)
  - its easy to compute the expected utility using VE
  - the number of computations is quite large: we run expected utility calculations (VE) for each parent instantiation together with each possible decision  $D$  might allow
  - policy: choose max decision for each parent instant'n



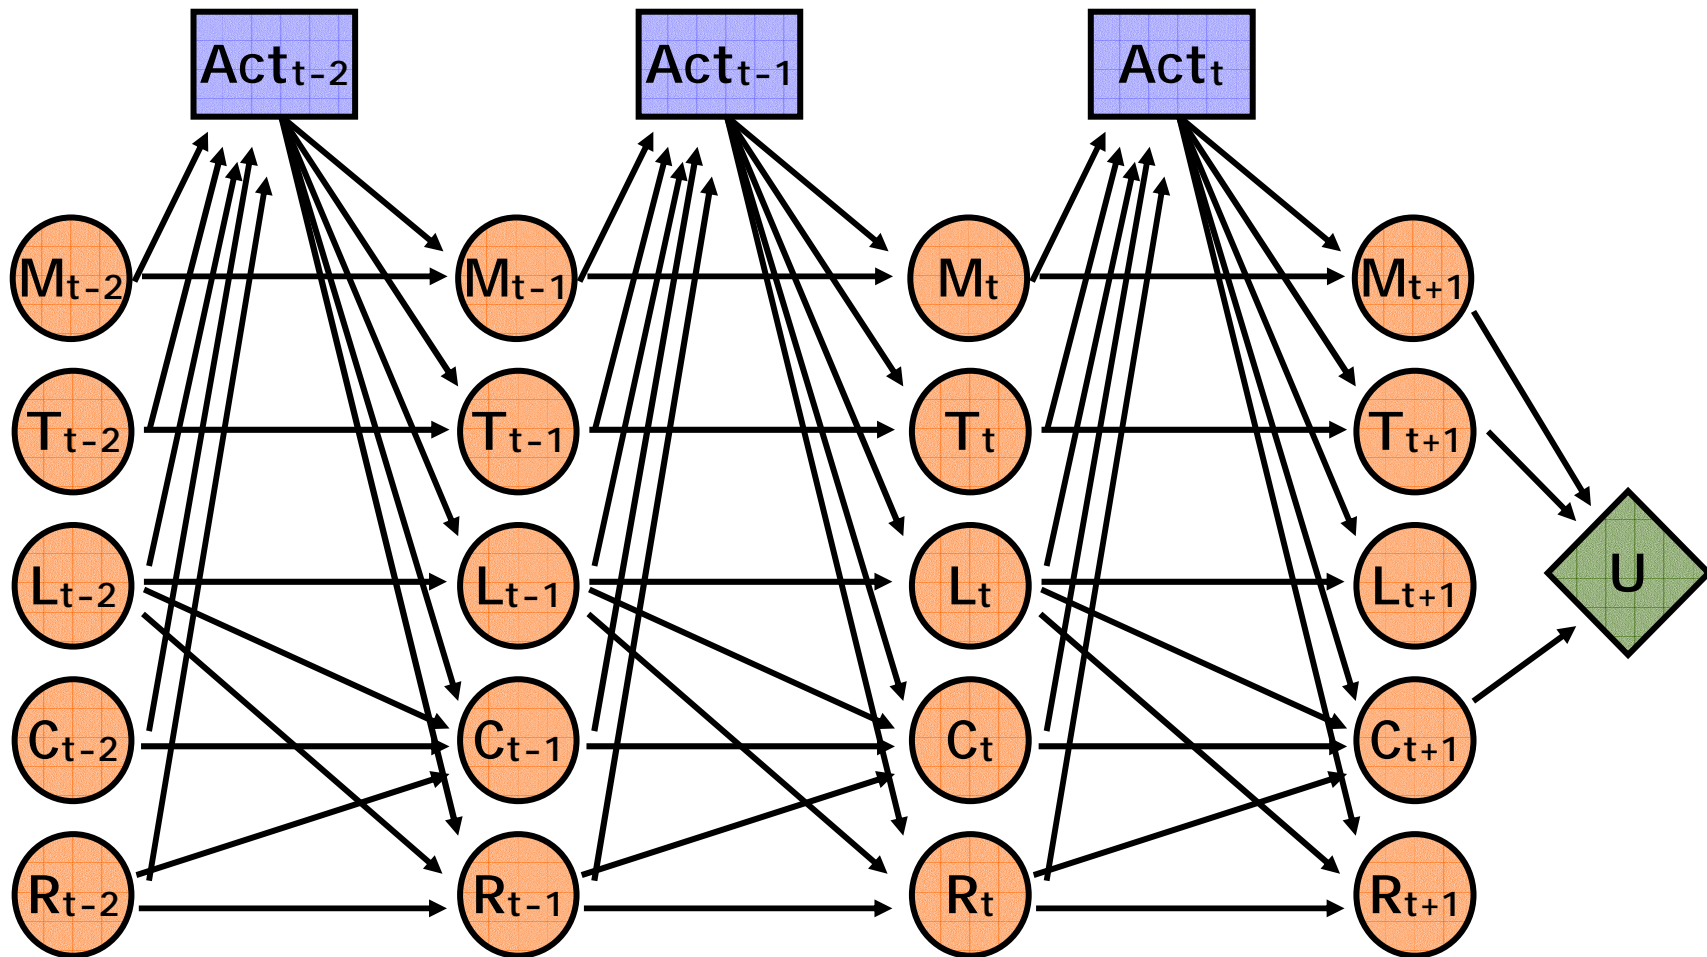
# Optimizing Policies: Key Points

- When a decision D node is optimized, it can be treated as a random variable
  - for each instantiation of its parents we now know what value the decision should take
  - just treat policy as a new CPT: for a given parent instantiation  $\mathbf{x}$ , D gets  $\delta(\mathbf{x})$  with probability 1 (all other decisions get probability zero)
- If we optimize from last decision to first, at each point we can optimize a specific decision by (a bunch of) simple VE calculations
  - it's successor decisions (optimized) are just normal nodes in the BNs (with CPTs)

# Decision Network Notes

- Decision networks commonly used by decision analysts to help structure decision problems
- Much work put into computationally effective techniques to solve these
  - common trick: replace the decision nodes with random variables at outset and solve a plain Bayes net (a subtle but useful transformation)
- Complexity much greater than BN inference
  - we need to solve a number of BN inference problems
  - one BN problem for each setting of decision node parents and decision node value

# DBN-Decision Nets for Planning



# DBN Decision Networks

- In example on previous slide:
  - we assume the state (of the variables at any stage) is fully observable
    - hence all time  $t$  vars point to time  $t$  decision
  - this means the state at time  $t$  d-separates the decision at time  $t-1$  from the decision at time  $t-2$
  - so we ignore “no-forgetting” arcs between decisions
    - once you *know* the state at time  $t$ , what you *did* at time  $t-1$  to get there is irrelevant to the decision at time  $t-1$
- If the state were not fully observable, we could not ignore the “no-forgetting” arcs

# A Detailed Decision Net Example

- Setting: you want to buy a used car, but there's a good chance it is a "lemon" (i.e., prone to breakdown). Before deciding to buy it, you can take it to a mechanic for inspection. S/he will give you a report on the car, labelling it either "good" or "bad". A good report is positively correlated with the car being sound, while a bad report is positively correlated with the car being a lemon.
- The report costs \$50 however. So you could risk it, and buy the car without the report.
- Owning a sound car is better than having no car, which is better than owning a lemon.

100



# Evaluate Last Decision: Buy (1)

- $EU(B|I,R) = \sum_L P(L|I,R,B) U(L,B)$
- $I = i, R = g$ :
  - $EU(\text{buy}) = P(I|i, g) U(I,\text{buy}) + P(\sim I|i, g) U(\sim I,\text{buy}) - 50$   
 $= .18 * -600 + .82 * 1000 - 50 = 662$
  - $EU(\sim \text{buy}) = P(I|i, g) U(I,\sim \text{buy}) + P(\sim I|i, g) U(\sim I,\sim \text{buy}) - 50$   
 $= -300 - 50 = -350$  (-300 indep. of lemon)
  - So optimal  $\delta_{Buy}(i,g) = \text{buy}$

## Evaluate Last Decision: Buy (2)

■  $I = i, R = b$ :

- $EU(\text{buy}) = P(I|i, b) U(I, \text{buy}) + P(\sim I|i, b) U(\sim I, \text{buy}) - 50$   
 $= .89 * -600 + .11 * 1000 - 50 = -474$
- $EU(\sim \text{buy}) = P(I|i, b) U(I, \sim \text{buy}) + P(\sim I|i, b) U(\sim I, \sim \text{buy}) - 50$   
 $= -300 - 50 = -350$  (-300 indep. of lemon)
- So optimal  $\delta_{Buy}(i, b) = \sim \text{buy}$



## Evaluate Last Decision: Buy (3)

- $I = \sim i$ ,  $R = n$  (note: no inspection cost subtracted)
  - $EU(\text{buy}) = P(I|\sim i, n) U(I, \text{buy}) + P(\sim I|\sim i, n) U(\sim I, \text{buy})$   
 $= .5 * -600 + .5 * 1000 = 200$
  - $EU(\sim \text{buy}) = P(I|\sim i, n) U(I, \sim \text{buy}) + P(\sim I|\sim i, n) U(\sim I, \sim \text{buy})$   
 $= -300 - 50 = -350$  (-300 indep. of lemon)
  - So optimal  $\delta_{Buy}(\sim i, n) = \text{buy}$
- So optimal policy for Buy is:
  - $\delta_{Buy}(i, g) = \text{buy}$  ;  $\delta_{Buy}(i, b) = \sim \text{buy}$  ;  $\delta_{Buy}(\sim i, n) = \text{buy}$
- Note: we don't bother computing policy for  $(i, \sim n)$ ,  $(\sim i, g)$ , or  $(\sim i, b)$ , since these occur with probability 0

# Evaluate First Decision: Inspect

■  $EU(I) = \sum_{L,R} P(L,R|I) U(L, \delta_{Buy}(I,R))$

- where  $P(R,L|I) = P(R|L,I)P(L|I)$
- $EU(i) = .1 \cdot -600 + .4 \cdot -300 + .45 \cdot 1000 + .05 \cdot -300 - 50$   
 $= 237.5 - 50 = 187.5$
- $EU(\sim i) = P(I|\sim i, n) U(I, buy) + P(\sim I|\sim i, n) U(\sim I, buy)$   
 $= .5 \cdot -600 + .5 \cdot 1000 = 200$
- So optimal  $\delta_{Inspect}(\sim i) = buy$



	$P(R,L   I)$	$\delta_{Buy}$	$U(L, \delta_{Buy})$
g,I	0.1	buy	$-600 - 50 = -650$
b,I	0.4	~buy	$-300 - 50 = -350$
g,~I	0.45	buy	$1000 - 50 = 950$
b,~I	0.05	~buy	$-300 - 50 = -350$

# Value of Information

- So optimal policy is: don't inspect, buy the car
  - $EU = 200$
  - Notice that the EU of inspecting the car, then buying it iff you get a good report, is 237.5 less the cost of the inspection (50). So inspection not worth the improvement in EU.
  - But suppose inspection cost \$25: then it would be worth it ( $EU = 237.5 - 25 = 212.5 > EU(\sim i)$ )
  - The *expected value of information* associated with inspection is 37.5 (it improves expected utility by this amount ignoring cost of inspection). How? Gives opportunity to change decision ( $\sim$ buy if bad).
  - You should be willing to pay up to \$37.5 for the report