Section 4

Decisions

A small quiz

Which of the following two lotteries would you prefer?:

- Lottery $A = [\$1\text{mill.}]$,
- Lottery $B = 0.1[\$5\text{mill.}] + 0.89[\$1\text{mill.}] + 0.01[\$0]$.

A small quiz

Which of the following two lotteries would you prefer?:

- Lottery $A = [\$1\text{mill.}]$,
- Lottery $B = 0.1[\$5\text{mill.}] + 0.89[\$1\text{mill.}] + 0.01[\$0]$.

What about these two?:

- Lottery $C = 0.11[\$1\text{mill.}] + 0.89[\$0]$,
- Lottery $D = 0.1[\$5\text{mill.}] + 0.9[\$0]$.

A small quiz

Which of the following two lotteries would you prefer?:

- Lottery $A = [\$1\text{mill.}]$,
- Lottery $B = 0.1[\$5\text{mill.}] + 0.89[\$1\text{mill.}] + 0.01[\$0]$.

What about these two?:

- Lottery $C = 0.11[\$1\text{mill.}] + 0.89[\$0]$,
- Lottery $D = 0.1[\$5\text{mill.}] + 0.9[\$0]$.

Is this the rational choice?

Poker again

Consider the poker example again:

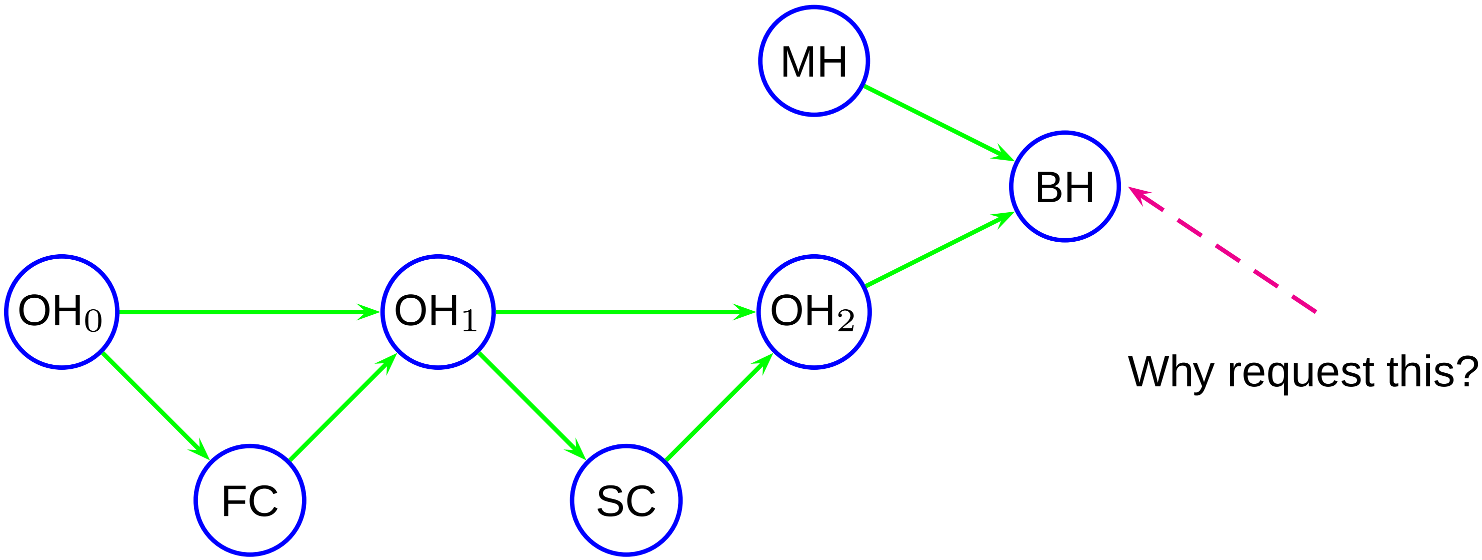

Poker again

Consider the poker example again:

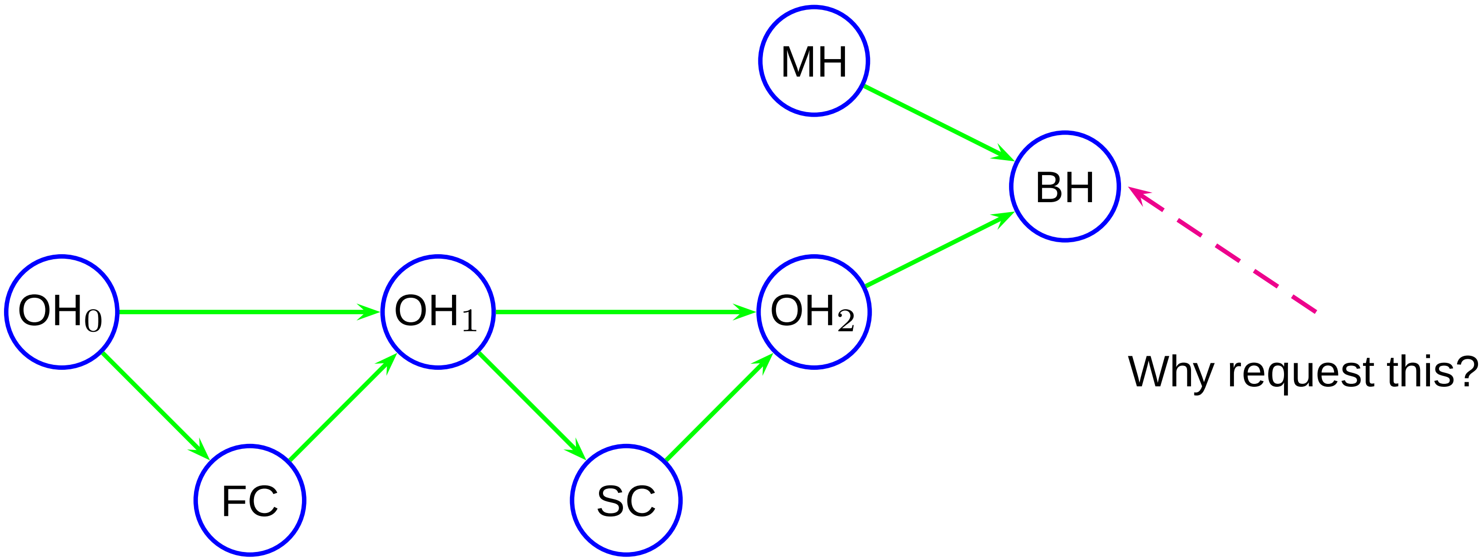

Fold or call?

- Both placed $1\$$
- She has placed $1\$$ more
- fold $\Rightarrow$ she takes the pot
- call $\Rightarrow$ place $1\$$ $\Rightarrow$ best hand takes the pot

Call or fold?

This decision problem can be represented graphically by extending the BN with a decision node and a utility node:

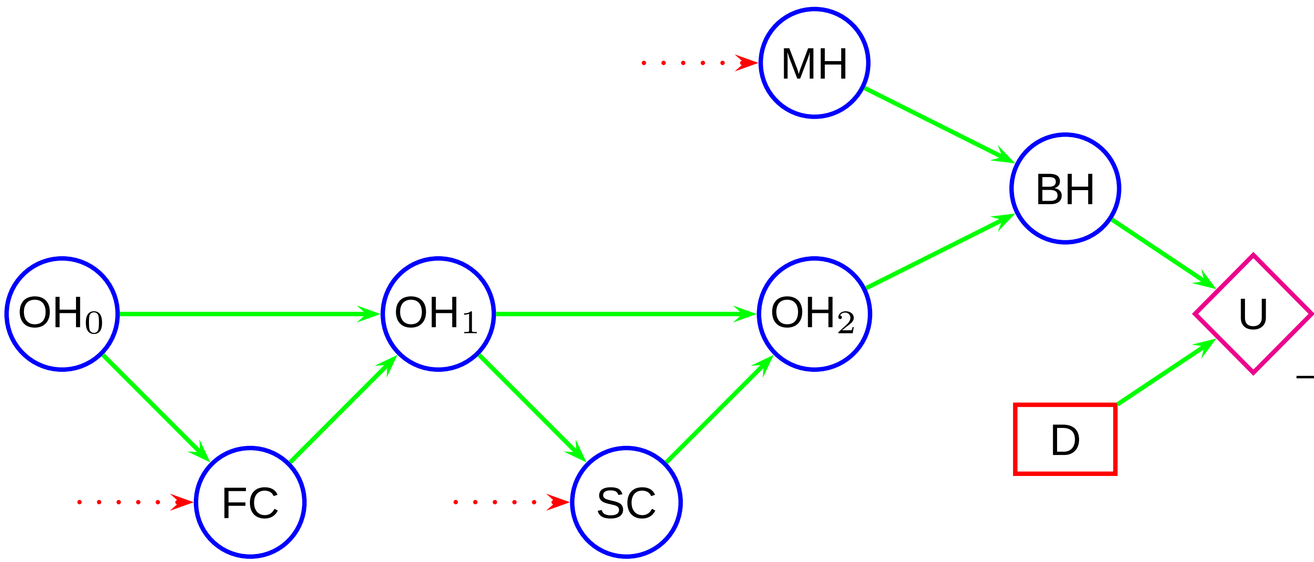

		D	
		fold	call
BH	l	-1	2
	op.	-1	-2
	draw	-1	0

U(BH, D)

Call or fold?

This decision problem can be represented graphically by extending the BN with a decision node and a utility node:

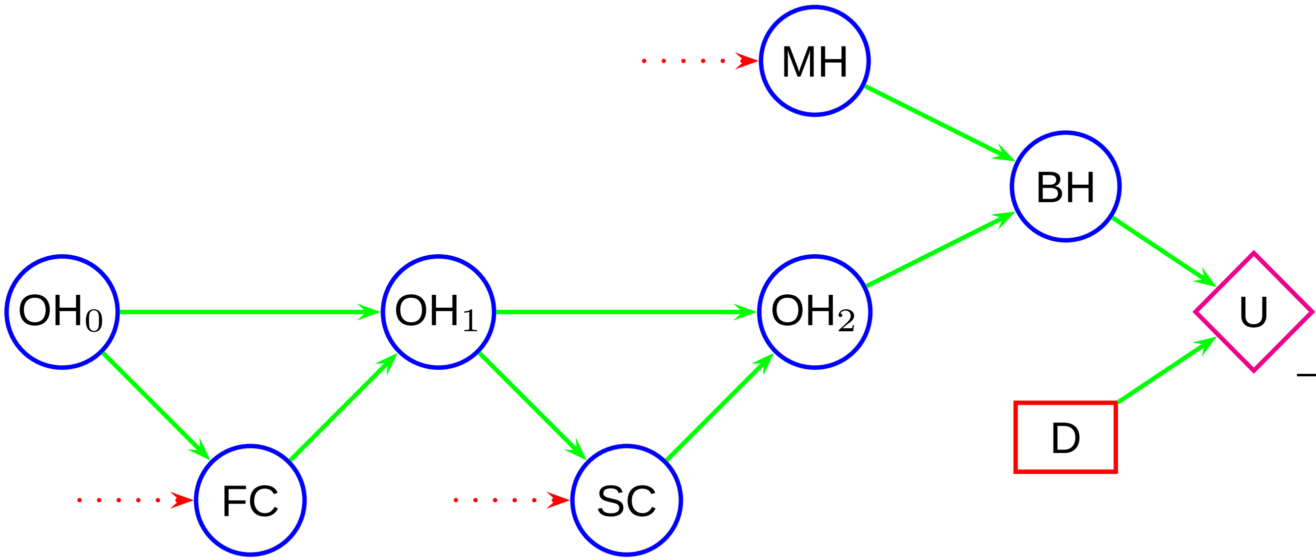

		D	
		fold	call
BH	l	-1	2
	op.	-1	-2
	draw	-1	0

U(BH, D)

The expected utility of call:

$$\text{EU}(\text{call}|e) = 2 \cdot P(\text{BH} = \text{l}|e) - 2 \cdot P(\text{BH} = \text{op.}|e) + 0 \cdot P(\text{BH} = \text{draw}|e)$$

$$= \sum_{\text{BH}} \text{U}(\text{BH}, \text{call}) P(\text{BH}|e)$$

One action in general

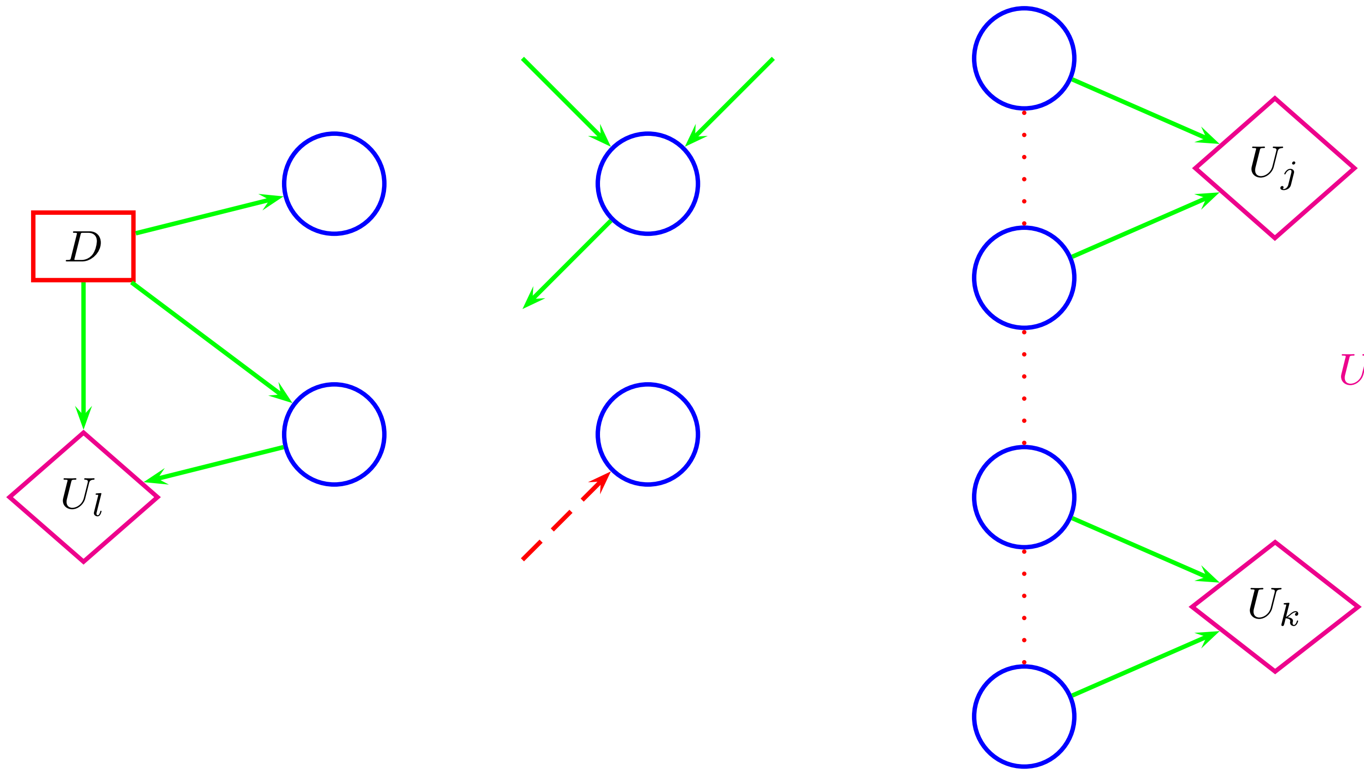

U_i has domain $\mathcal{X}_i$

One action in general

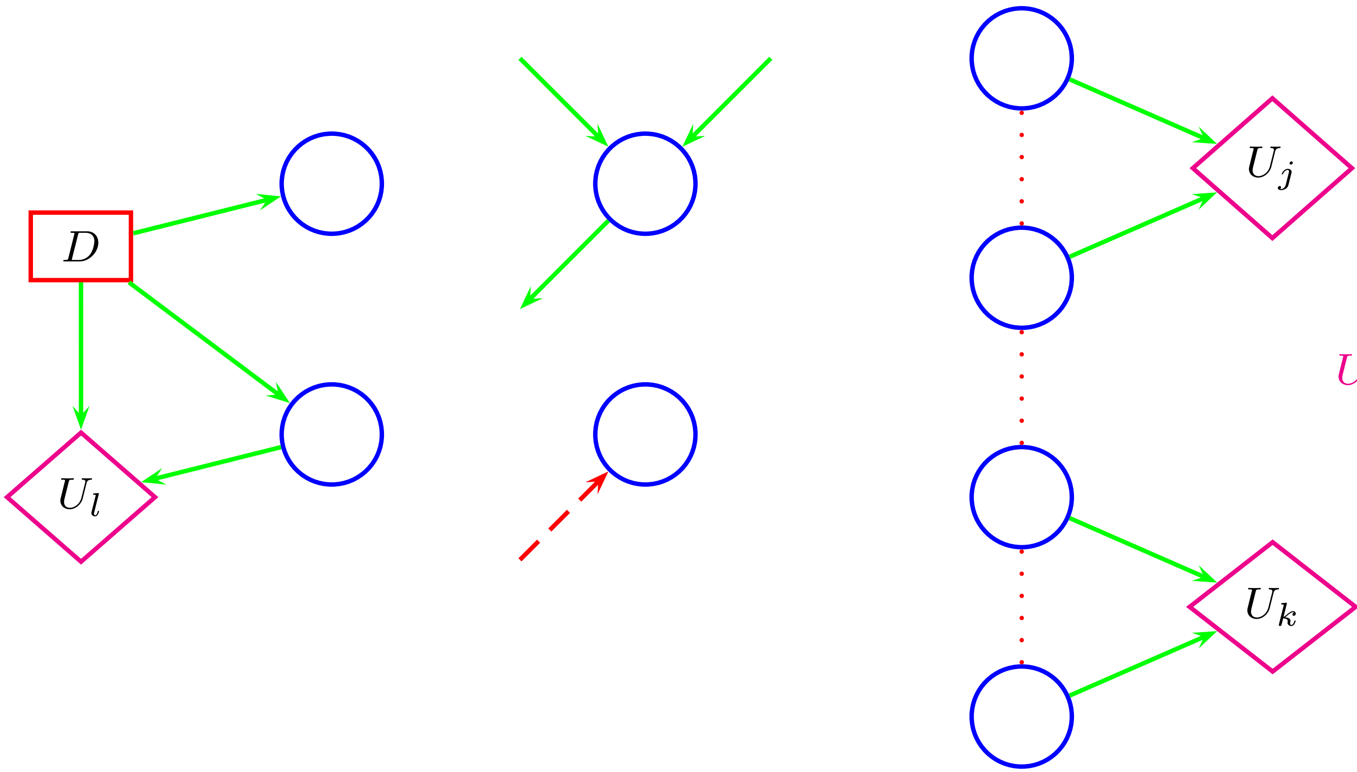

U_i has domain $\mathcal{X}_i$

$$EU(D|e) = \sum_{\mathcal{X}_1} U_1(\mathcal{X}_1) P(\mathcal{X}_1|D,e) + \cdots + \sum_{\mathcal{X}_n} U_n(\mathcal{X}_n) P(\mathcal{X}_n|D,e)$$

Choose an action with largest EU:

$$Opt(D|e) = \arg\max_{D} EU(D|e)$$

Utilities without money

Two courses: Graph algorithms (GA) and DSS

Marks: $0, 1, 2, 3, 4, 5$ (≥ 2 is a pass)

Effort: Keep pace (kp), slow down (sd), follow superficially (fs)

		Effort		
		kp	sd	fs
	0	0	0	0.1
	1	0.1	0.2	0.1
GA	2	0.1	0.1	0.4
	3	0.2	0.4	0.2
	4	0.4	0.2	0.2
	5	0.2	0.1	0

$P($GA$|$Effort$)$

		Effort		
		kp	sd	fs
	0	0	0	0.1
	1	0	0.1	0.2
DSS	2	0.1	0.2	0.2
	3	0.2	0.2	0.3
	4	0.4	0.4	0.2
	5	0.3	0.1	0

$P($DSS$|$Effort$)$

Max-score? Max-pass? Otherwise?

Marks as utilities?

		Effort		
		kp	sd	fs
	0	0	0	0.1
	1	0.1	0.2	0.1
GA	2	0.1	0.1	0.4
	3	0.2	0.4	0.2
	4	0.4	0.2	0.2
	5	0.2	0.1	0

$P($GA$|$Effort$)$

		Effort		
		kp	sd	fs
	0	0	0	0.1
	1	0	0.1	0.2
DSS	2	0.1	0.2	0.2
	3	0.2	0.2	0.3
	4	0.4	0.4	0.2
	5	0.3	0.1	0

$P($DSS$|$Effort$)$

$$EU(\text{kp,fs}) = \sum_{\text{m} \in \text{GA}} P(\text{m}|\text{kp})\text{m} + \sum_{\text{m} \in \text{DSS}} P(\text{m}|\text{fs})\text{m}$$

$$= 5.8$$

$$EU(\text{sd,sd}) = 6.1$$

$$EU(\text{fs,kp}) = \underline{6.2}$$

However, do the marks really reflect your utilities?

Subjective lotteries

I consider 2 as the worst mark (utility 0) and 5 as the best mark (utility 1). Now imagine the following lottery:

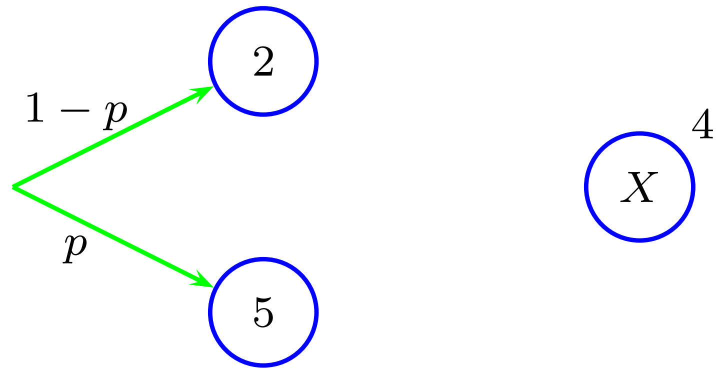

For which p am I indifferent??

Subjective lotteries

I consider 2 as the worst mark (utility 0) and 5 as the best mark (utility 1). Now imagine the following lottery:

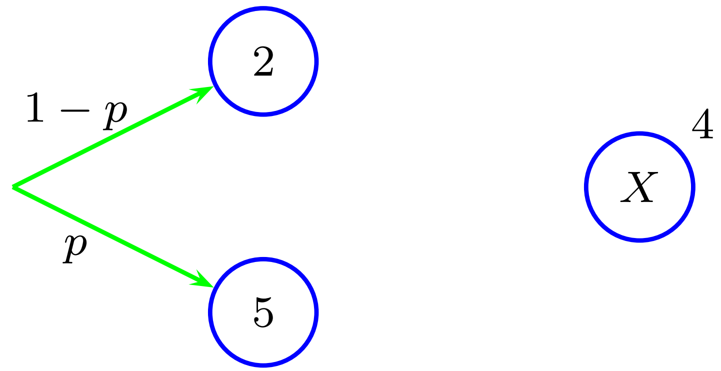

For which p am I indifferent??

$EU(Game1) = EU(Game2) => 1U(x) = (1-p)U(2) + pU(5) = p$

Subjective lotteries

I consider 2 as the worst mark (utility 0) and 5 as the best mark (utility 1). Now imagine the following lottery:

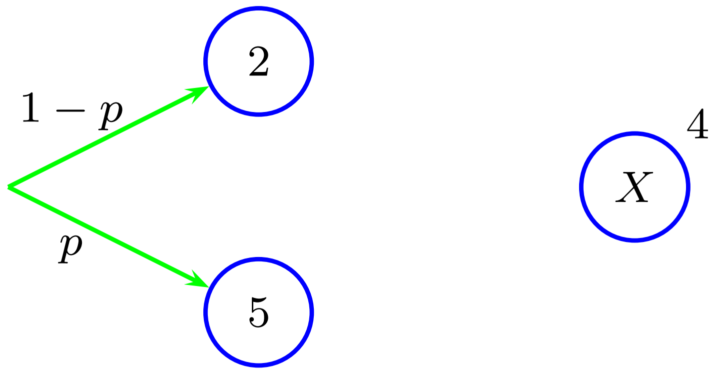

For which p am I indifferent??

$$EU(Game1) = EU(Game2) => 1U(x) = (1-p)U(2) + pU(5) = p$$

The utility table:

0	1	2	3	4	5
0.05	0.1	0	0.6	0.8	1

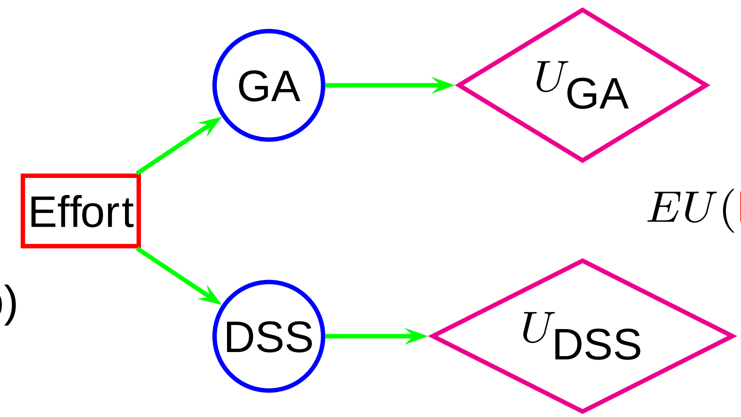

$EU(\text{Effort}) = (1.015, \underline{1.07}, 1.035)$

(kp,fs),(sd,sd),(fs,kp)

Instrumental rationality

1. ***Reflexivity.*** For any lottery A, $A \succeq A$
2. ***Completeness.*** For any pair (A, B) of lotteries, $A \succeq B$ or $B \succeq A$.
3. ***Transitivity.*** If $A \succeq B$ and $B \succeq C$, then $A \succeq C$.

Instrumental rationality

1. ***Reflexivity.*** For any lottery A, $A \succeq A$
2. ***Completeness.*** For any pair (A, B) of lotteries, $A \succeq B$ or $B \succeq A$.
3. ***Transitivity.*** If $A \succeq B$ and $B \succeq C$, then $A \succeq C$.
4. ***Preference increasing with probability.*** If $A \succeq B$ then $\alpha A + (1 - \alpha)B \succeq \beta A + (1 - \beta)B$ if and only if $\alpha \geq \beta$.

Instrumental rationality

1. ***Reflexivity.*** For any lottery A, $A \succeq A$
2. ***Completeness.*** For any pair (A, B) of lotteries, $A \succeq B$ or $B \succeq A$.
3. ***Transitivity.*** If $A \succeq B$ and $B \succeq C$, then $A \succeq C$.
4. ***Preference increasing with probability.*** If $A \succeq B$ then $\alpha A + (1-\alpha)B \succeq \beta A + (1-\beta)B$ if and only if $\alpha \geq \beta$.
5. ***Continuity.*** If $A \succeq B \succeq C$ then there exists $\alpha \in [0, 1]$ such that $B \sim \alpha A + (1-\alpha)C$

Instrumental rationality

1. ***Reflexivity.*** For any lottery A, $A \succeq A$
2. ***Completeness.*** For any pair (A, B) of lotteries, $A \succeq B$ or $B \succeq A$.
3. ***Transitivity.*** If $A \succeq B$ and $B \succeq C$, then $A \succeq C$.
4. ***Preference increasing with probability.*** If $A \succeq B$ then $\alpha A + (1 - \alpha)B \succeq \beta A + (1 - \beta)B$ if and only if $\alpha \geq \beta$.
5. ***Continuity.*** If $A \succeq B \succeq C$ then there exists $\alpha \in [0, 1]$ such that $B \sim \alpha A + (1 - \alpha)C$
6. ***Independence.*** If $C = \alpha A + (1 - \alpha)B$ and $A \sim D$, then $C \sim (\alpha D + (1 - \alpha)B)$.

Instrumental rationality

1. ***Reflexivity.*** For any lottery A, $A \succeq A$
2. ***Completeness.*** For any pair (A, B) of lotteries, $A \succeq B$ or $B \succeq A$.
3. ***Transitivity.*** If $A \succeq B$ and $B \succeq C$, then $A \succeq C$.
4. ***Preference increasing with probability.*** If $A \succeq B$ then $\alpha A + (1-\alpha)B \succeq \beta A + (1-\beta)B$ if and only if $\alpha \geq \beta$.
5. ***Continuity.*** If $A \succeq B \succeq C$ then there exists $\alpha \in [0,1]$ such that $B \sim \alpha A + (1-\alpha)C$
6. ***Independence.*** If $C = \alpha A + (1-\alpha)B$ and $A \sim D$, then $C \sim (\alpha D + (1-\alpha)B)$.

Theorem: For an individual who acts according to a preference ordering satisfying rules 1-6 above, there exists a utility function over the outcomes s.t. the expected utility is maximized.

Are you rational?

Recall:

- Lottery $A = [\$1\text{mill.}]$,
- Lottery $B = 0.1[\$5\text{mill.}] + 0.89[\$1\text{mill.}] + 0.01[\$0]$.
- Lottery $C = 0.11[\$1\text{mill.}] + 0.89[\$0]$,
- Lottery $D = 0.1[\$5\text{mill.}] + 0.9[\$0]$.

Are you rational?

Recall:

- Lottery $A = [\$1\text{mill.}]$,
- Lottery $B = 0.1[\$5\text{mill.}] + 0.89[\$1\text{mill.}] + 0.01[\$0]$.
- Lottery $C = 0.11[\$1\text{mill.}] + 0.89[\$0]$,
- Lottery $D = 0.1[\$5\text{mill.}] + 0.9[\$0]$.

Let $U(5mill) = 1, U(0) = 0, U(1mill) = u$. If you prefer A over B we get

$$u > 0.1 + 0.89u \qquad \Leftrightarrow \qquad u > \frac{10}{11}.$$

Are you rational?

Recall:

- Lottery $A = [\$1\text{mill.}]$,
- Lottery $B = 0.1[\$5\text{mill.}] + 0.89[\$1\text{mill.}] + 0.01[\$0]$.
- Lottery $C = 0.11[\$1\text{mill.}] + 0.89[\$0]$,
- Lottery $D = 0.1[\$5\text{mill.}] + 0.9[\$0]$.

Let $U(5mill) = 1, U(0) = 0, U(1mill) = u$. If you prefer A over B we get

$$u > 0.1 + 0.89u \qquad \Leftrightarrow \qquad u > \frac{10}{11}.$$

Hence,

$$EU(C) = 0.11u > 0.11\frac{10}{11} = 0.1 = EU(D),$$

and C should therefore be preferred over D.

Sequential Decision Making: Decision trees

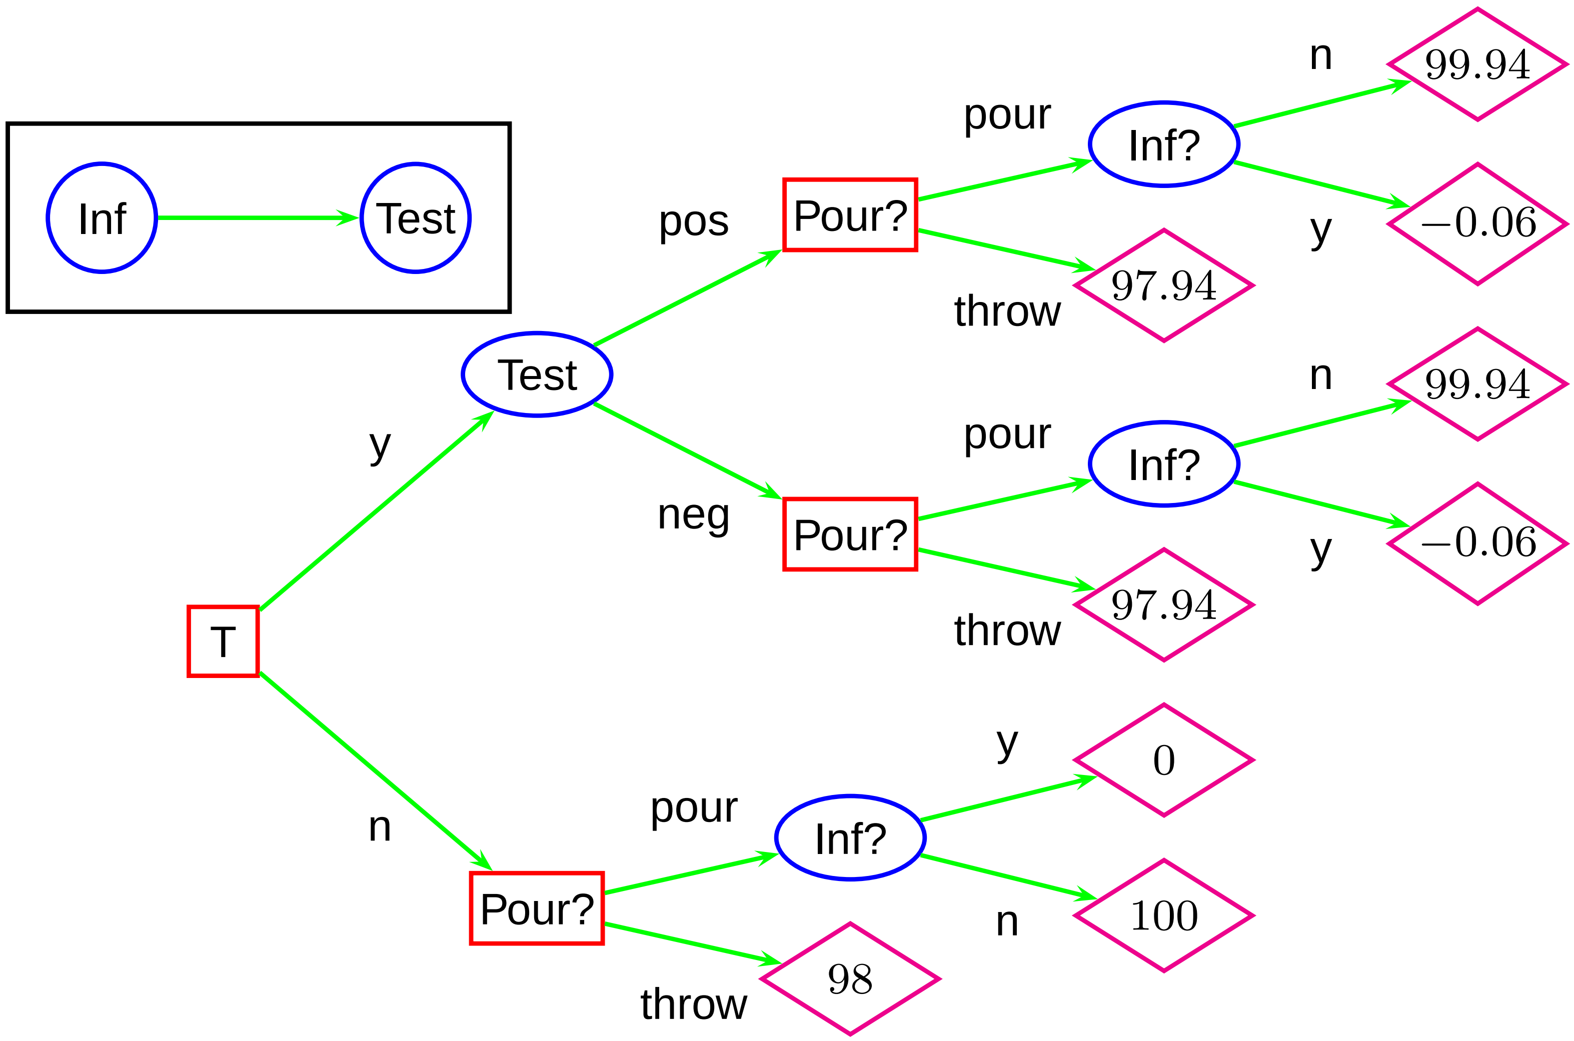

Sequential Decision Making: Decision trees

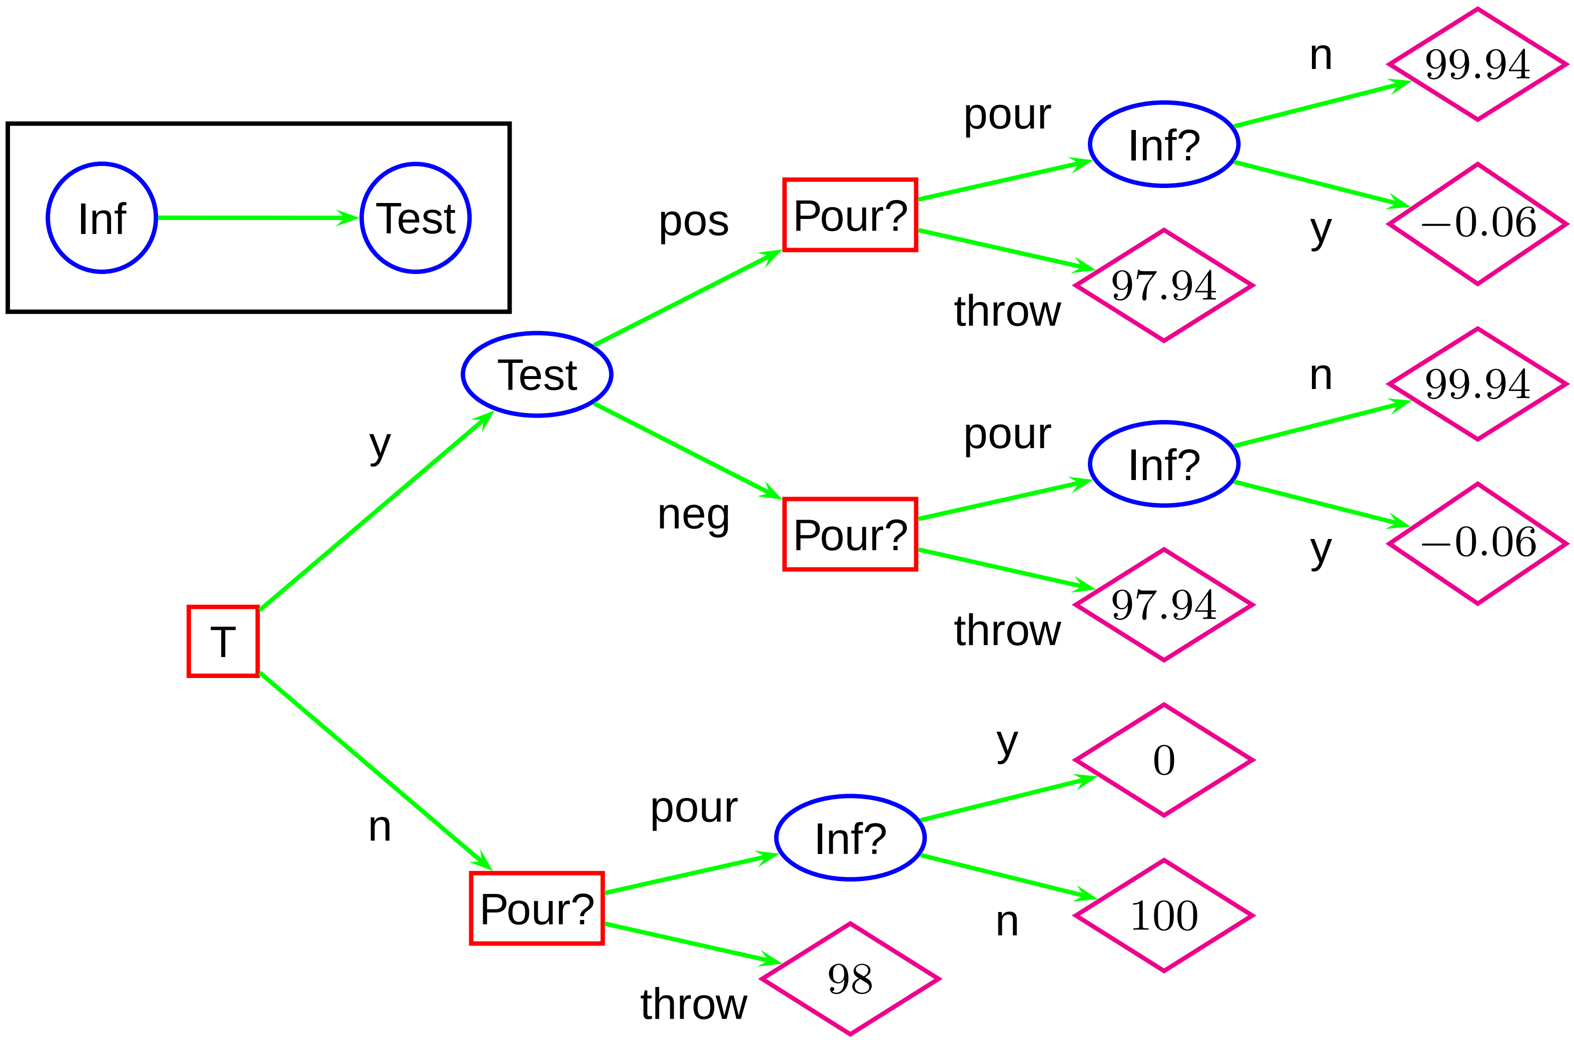

Branches from chance nodes, 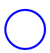, shall be labeled with the probability of the branch given the path down to the node. The probabilities can be found from the model 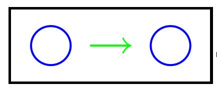.

Solving decision trees I

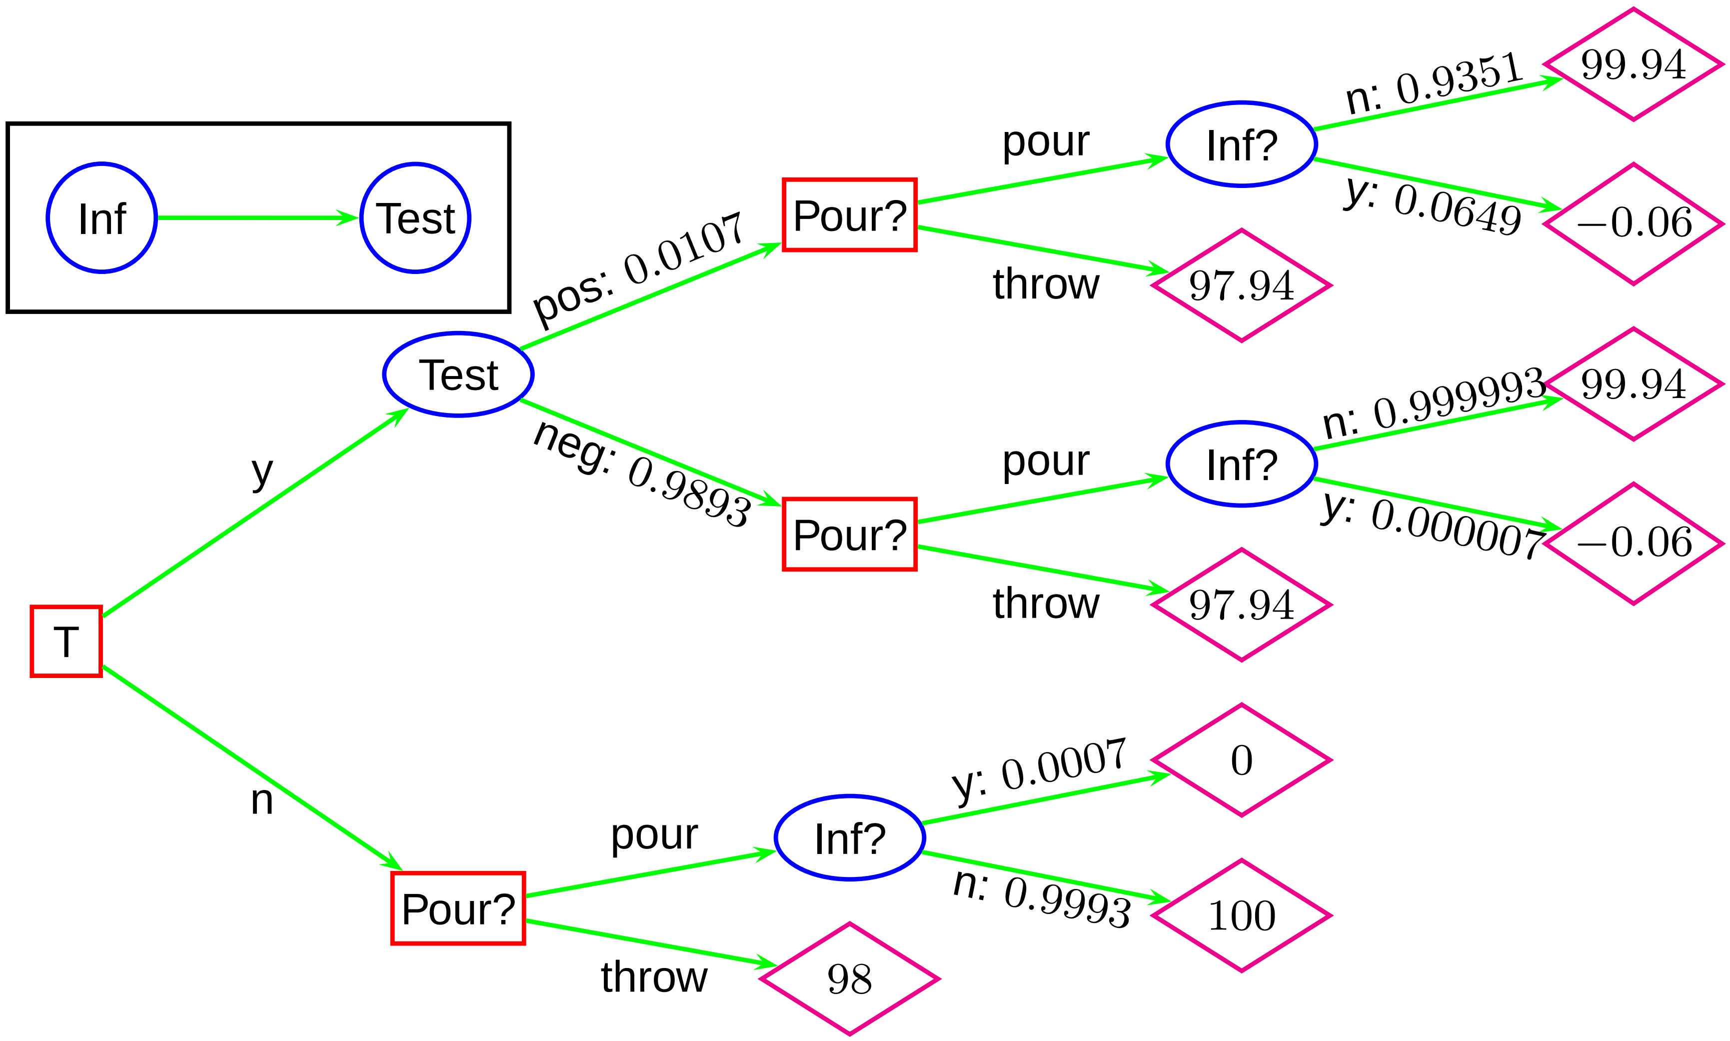

Solving decision trees I

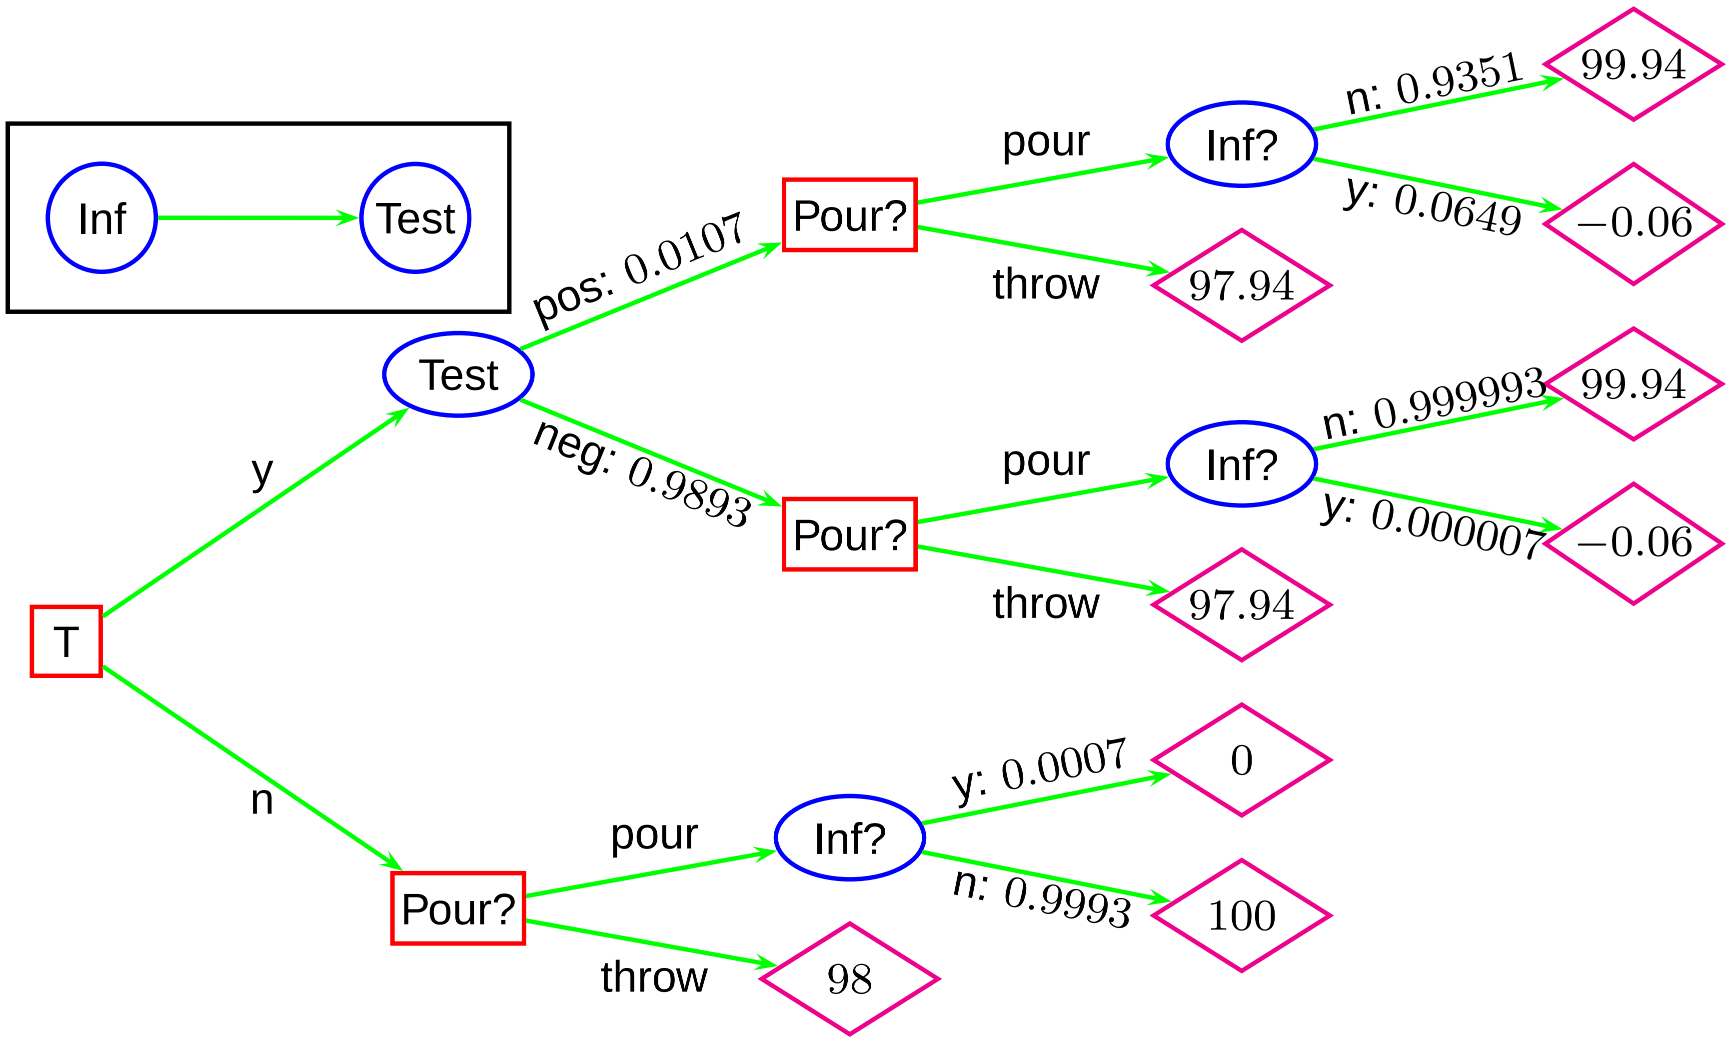

The decision tree can be solved by going from the leaves towards the root:

- Take weighted sum through chance nodes.
- Take max through decision nodes.

Solving decision trees II

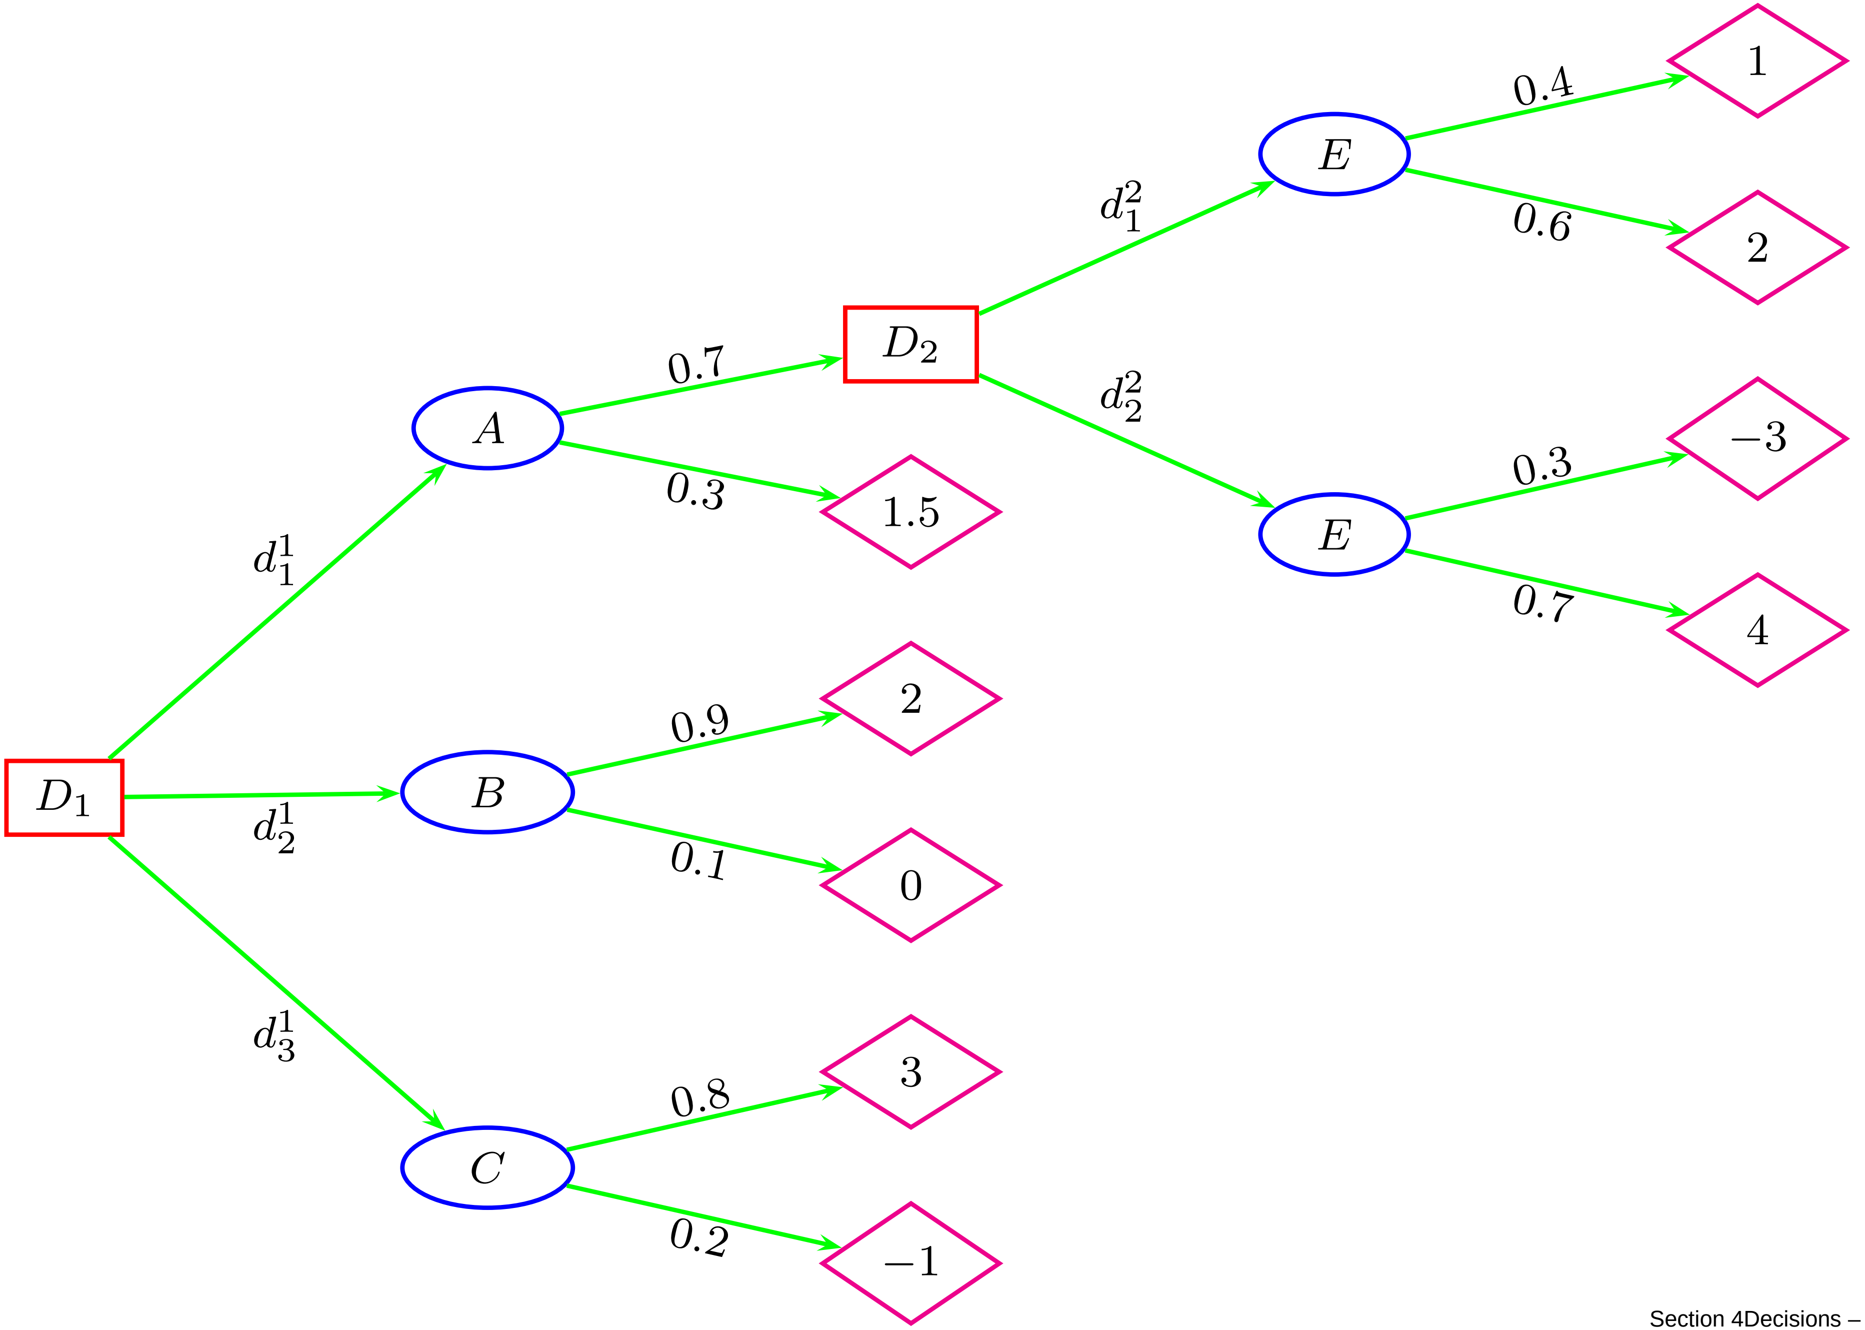

Solving decision trees II

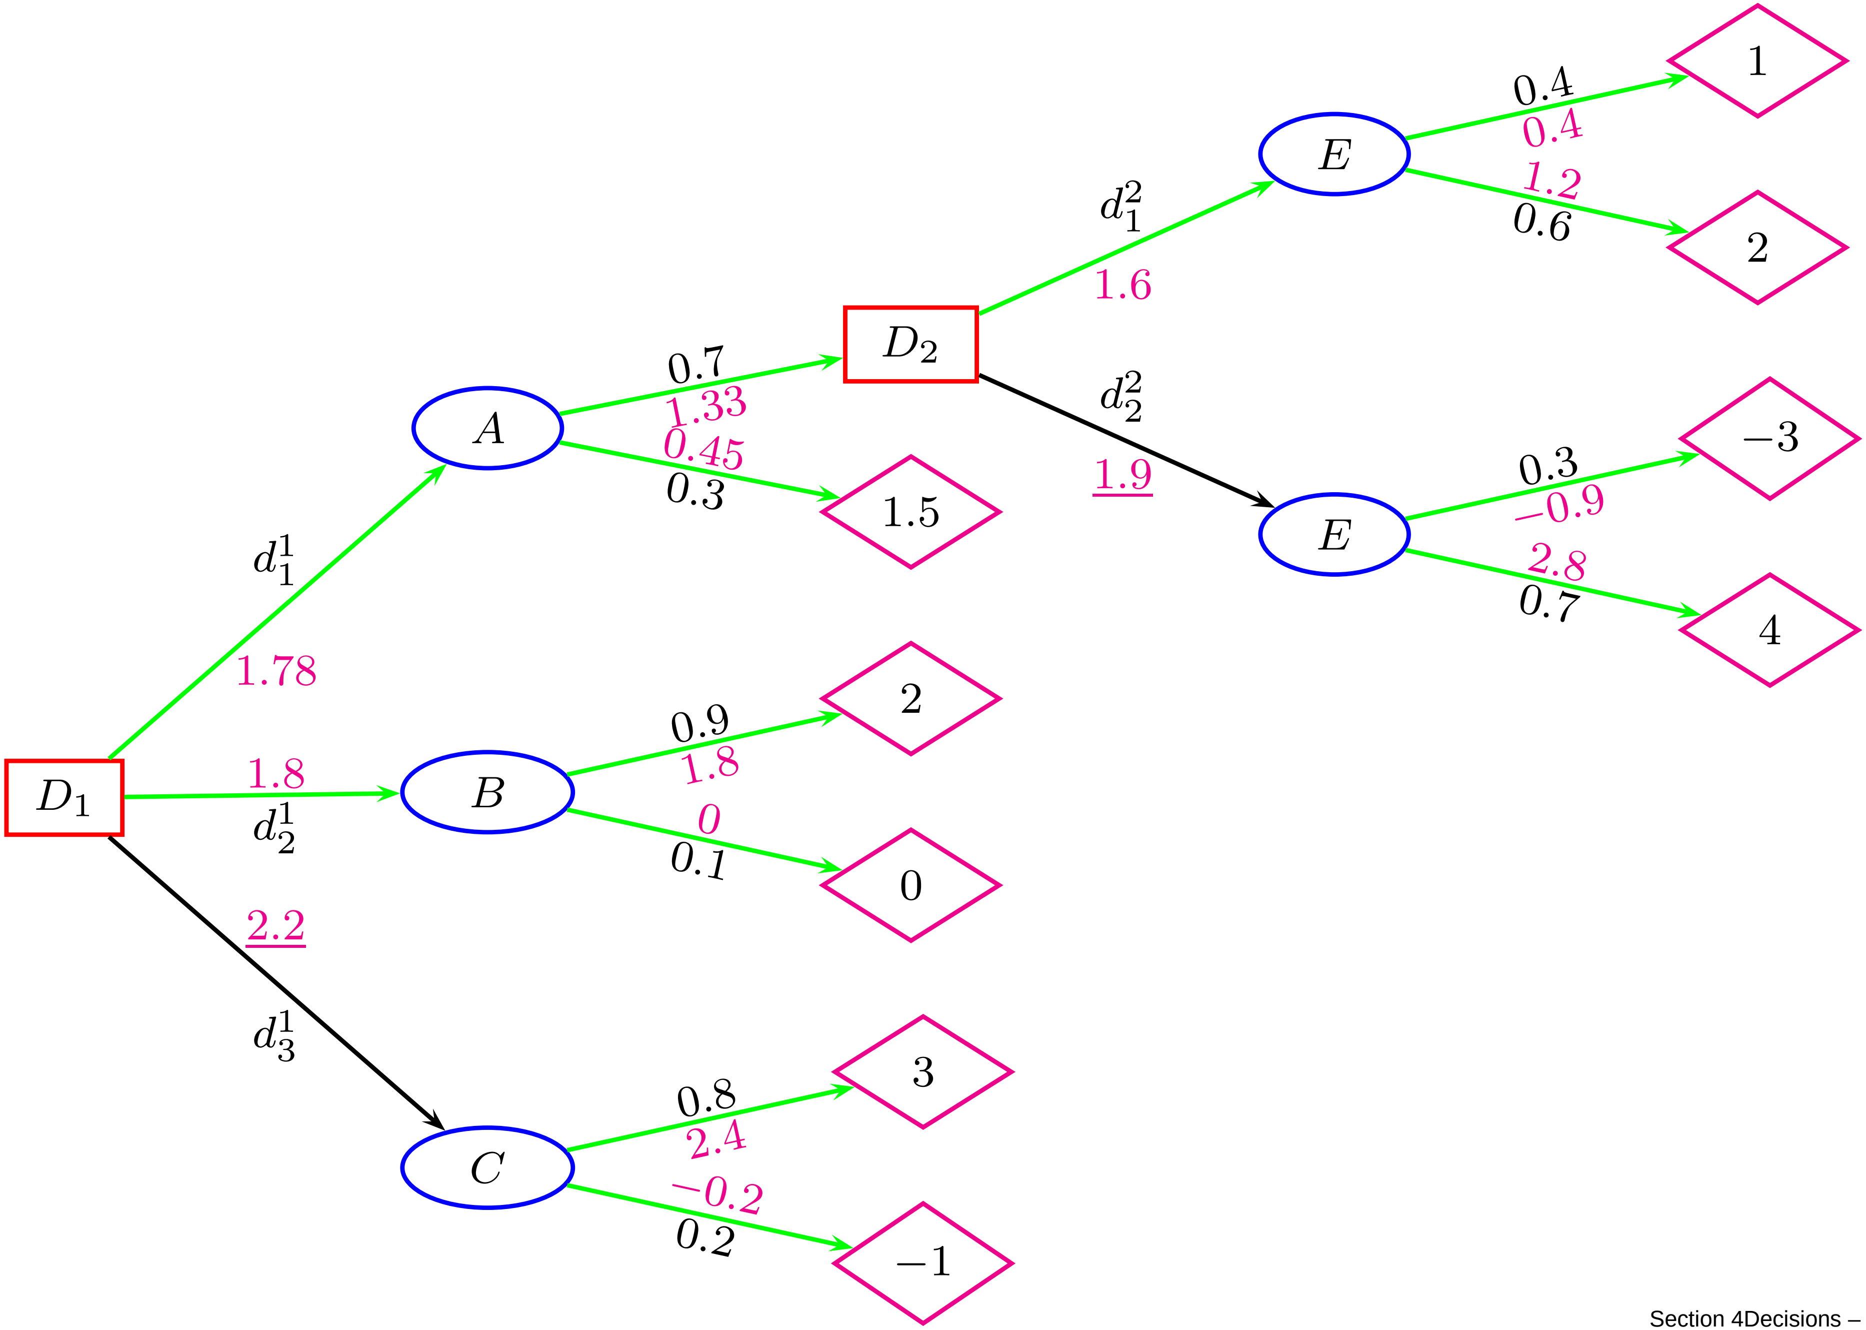

Decision trees: characteristics

Advantages:

- All scenarios are represented explicitly.
- Very few restrictions on the decision problems that can be represented.

Disadvantages:

- Two separate models are used: one representing the structure and one representing the uncertainties.
- The size of the decision trees grows exponentially in the number of variables.

An alternative representation

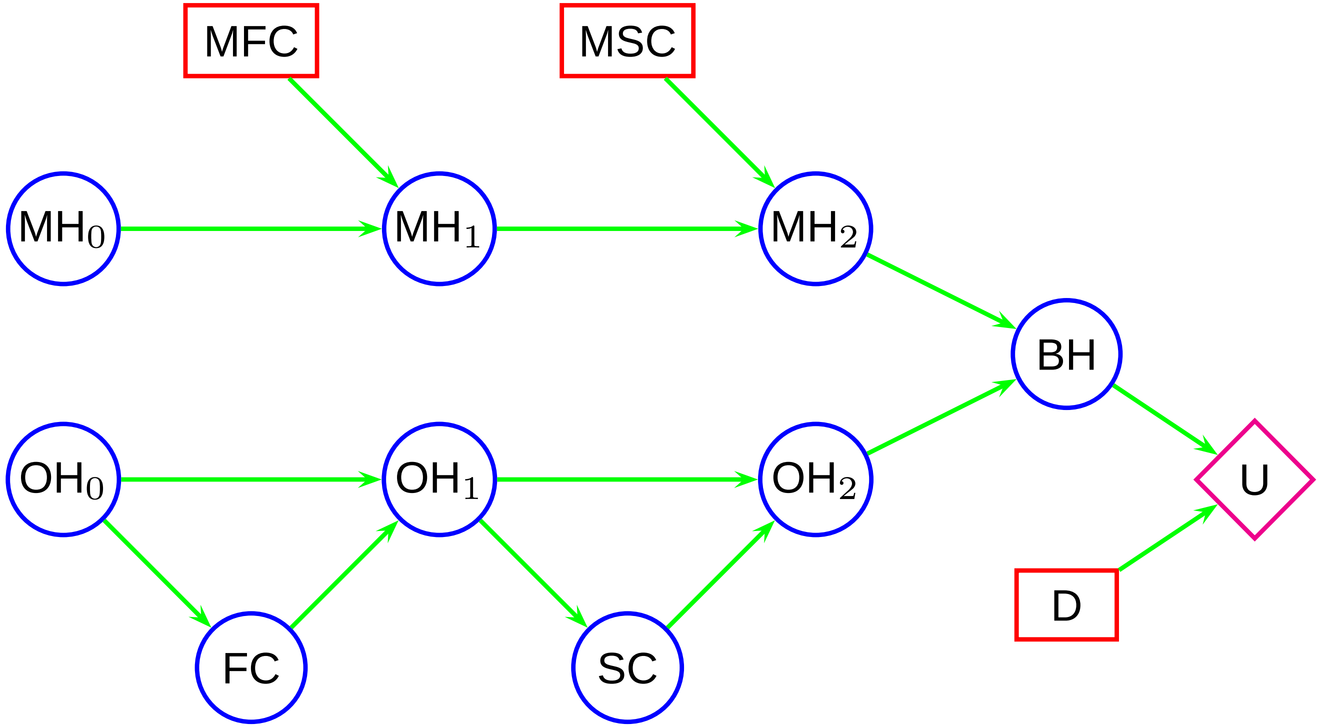

But how do we represent the sequence of decisions and observations?

Representing the decision sequence

Possible representation:

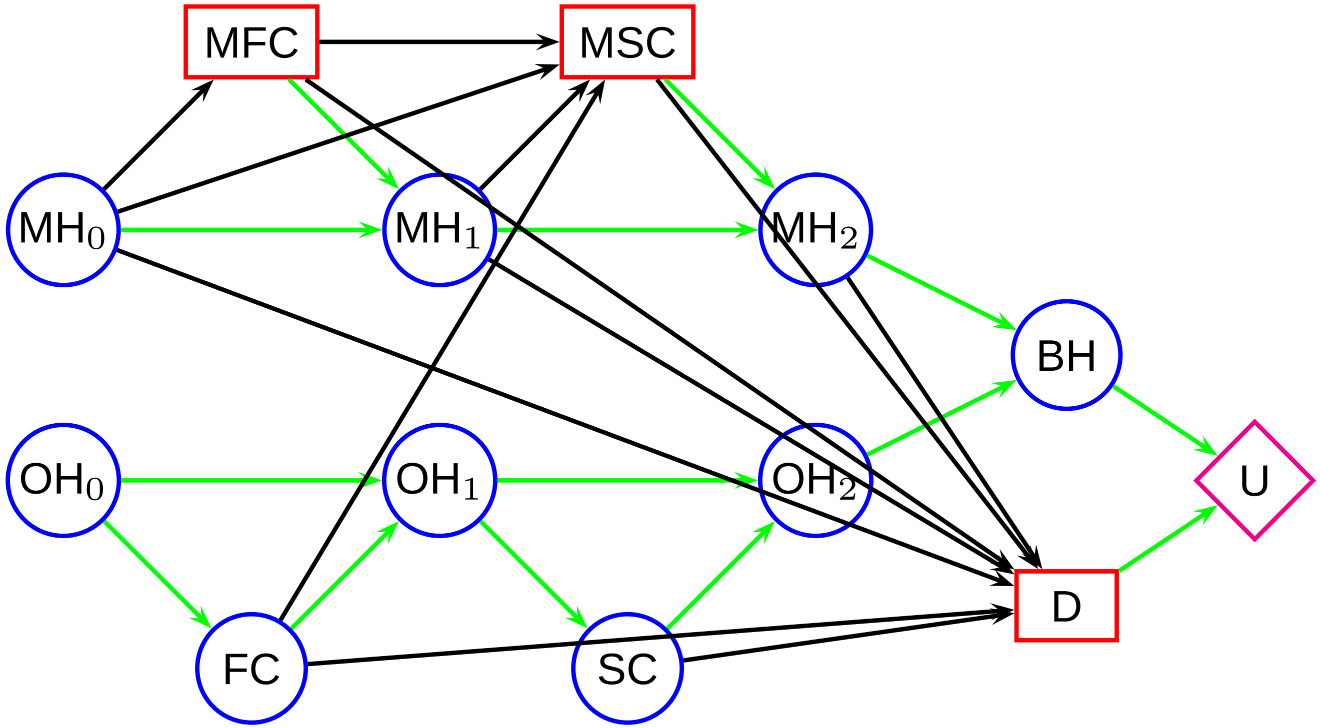

All nodes observed before a decision are parents of that decision.

- Assuming that the decision maker doesn't forget, then some links are redundant!

Representing the decision sequence

A better representation (an influence diagram):

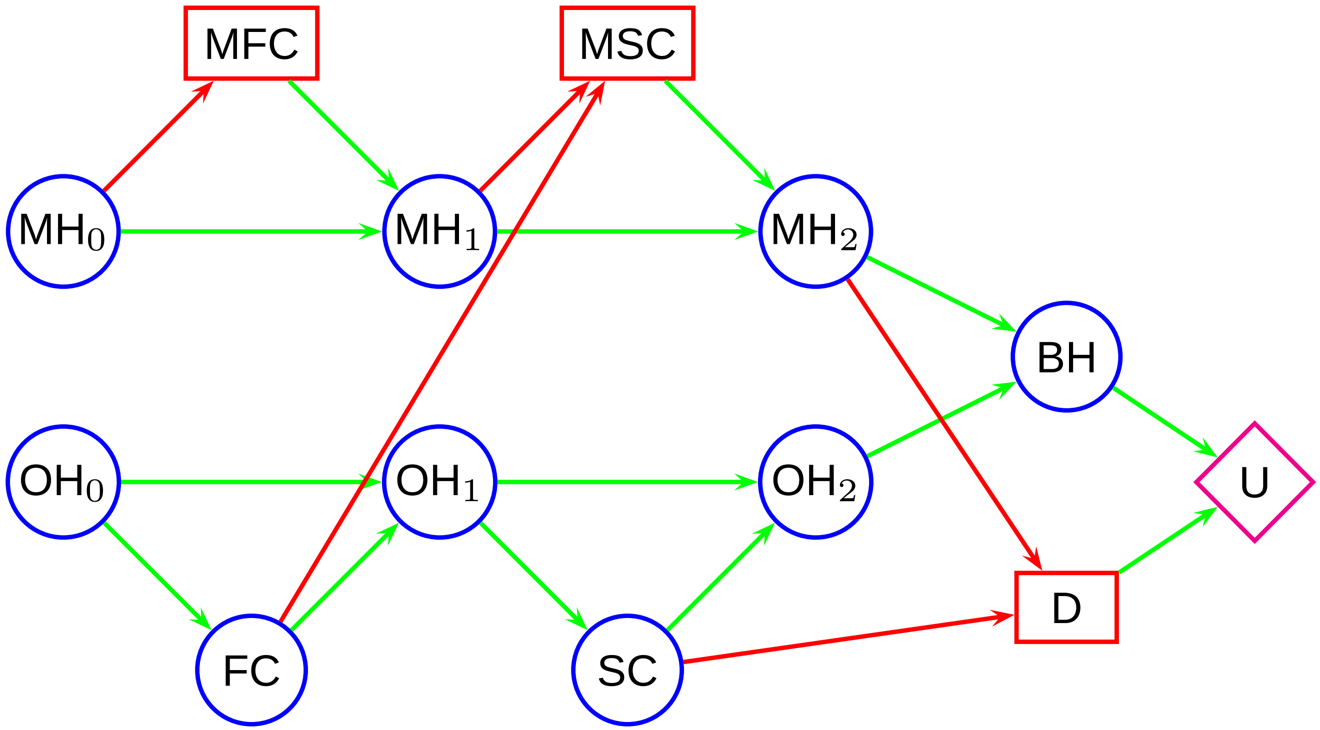

Advantages:

- You can read the sequence of decisions.
- You can read what is known at each point of decision.

Influence diagrams

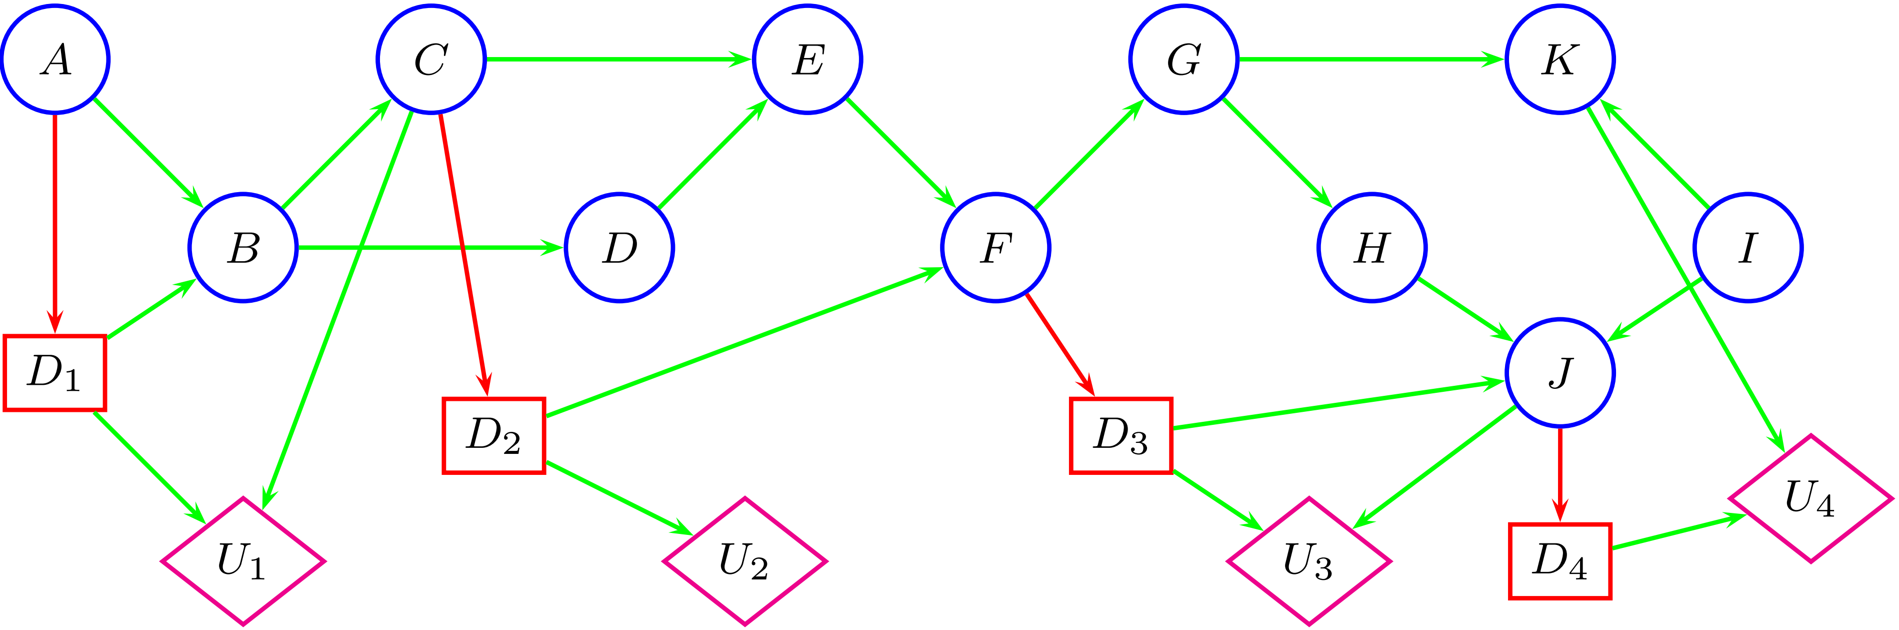

Nodes and links:

○ Chance variable $\longrightarrow$ causal links

□ Decision variable $\longrightarrow$ information links

◇ Utility function $\longrightarrow$ utility link, $U = \sum_i U_i$.

Note:

- We assume no-forgetting.
- A directed path comprising all decisions $\Rightarrow$ the scenario is well-defined.

Influence diagrams: Characteristics

Advantages:

- Grows only linearly in the number of variables.
- Requires only one model for representing both structure as well as the uncertainty model.

Disadvantages:

- The sequence of observations and decisions is the same in all scenarios (the decision problem is symmetric).

Definition: A decision problem is said to be symmetric if:

- In all decision tree representations, the number of scenarios is the same as the cardinality of the Cartesian product of the state spaces of all chance and decision variables.
- in one decision tree representation, the sequence of observations and decisions is the same in all scenarios.

Symmetric decision trees

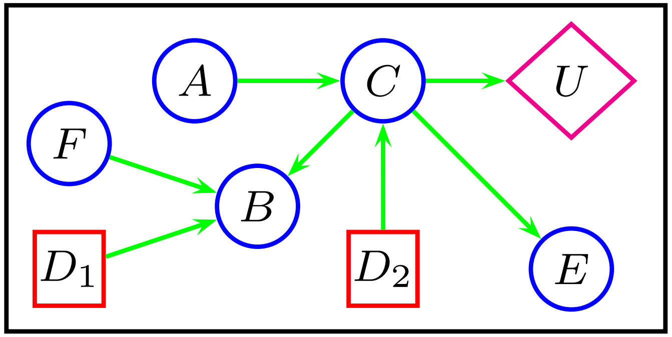

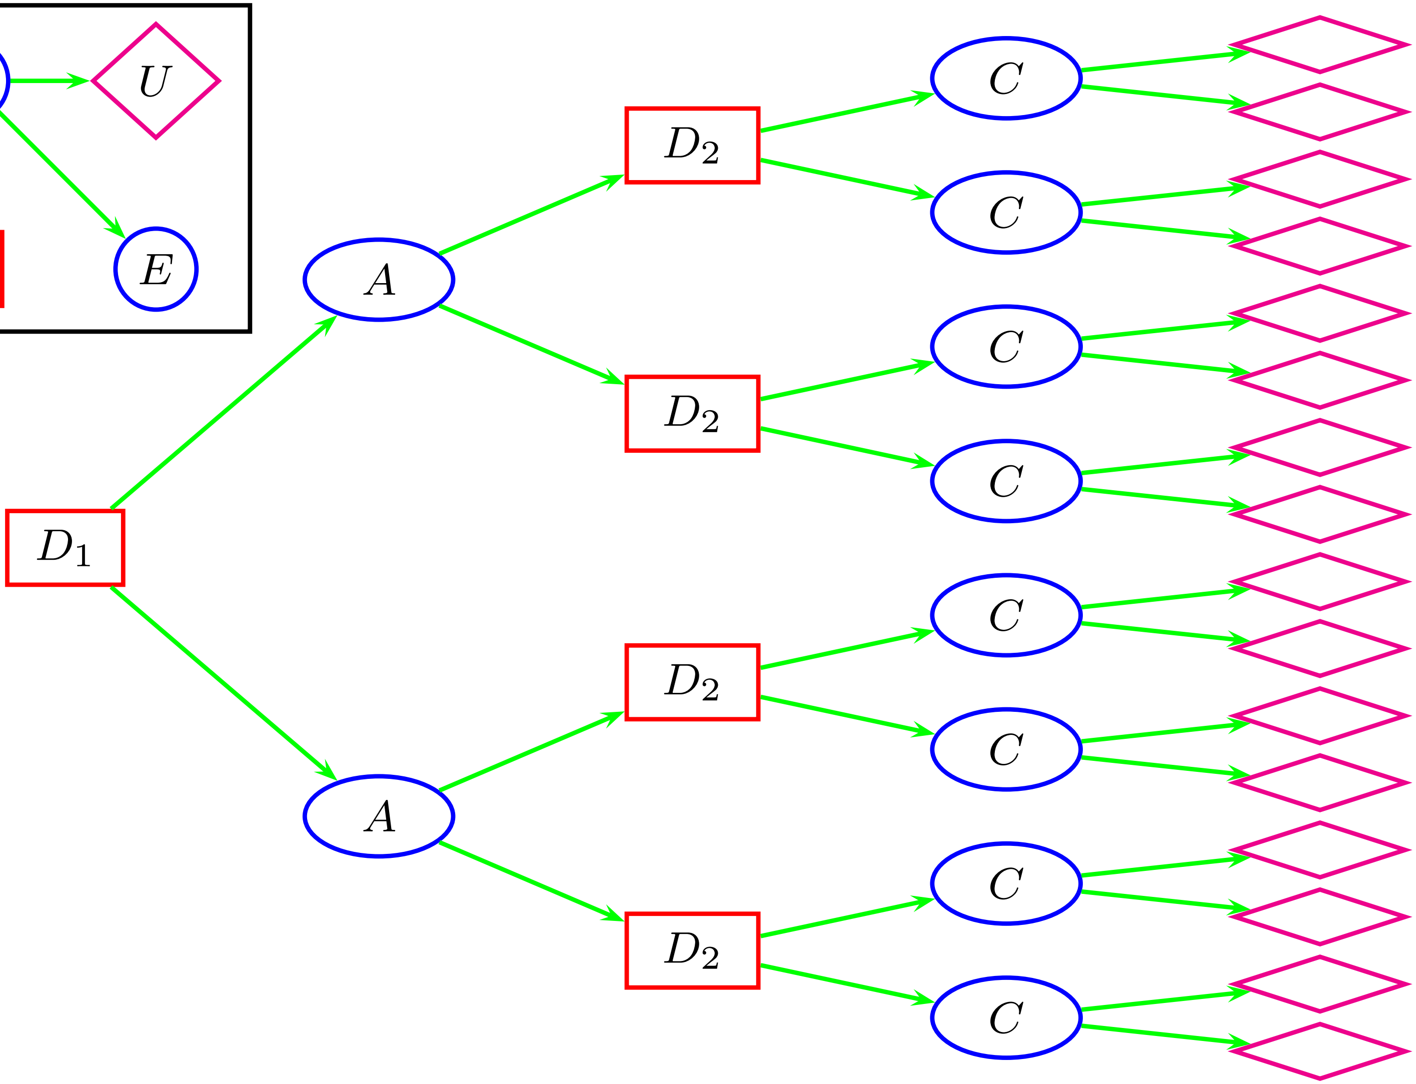

Symmetric decision trees

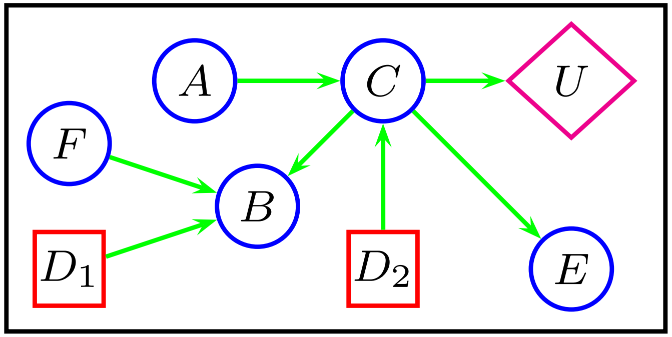

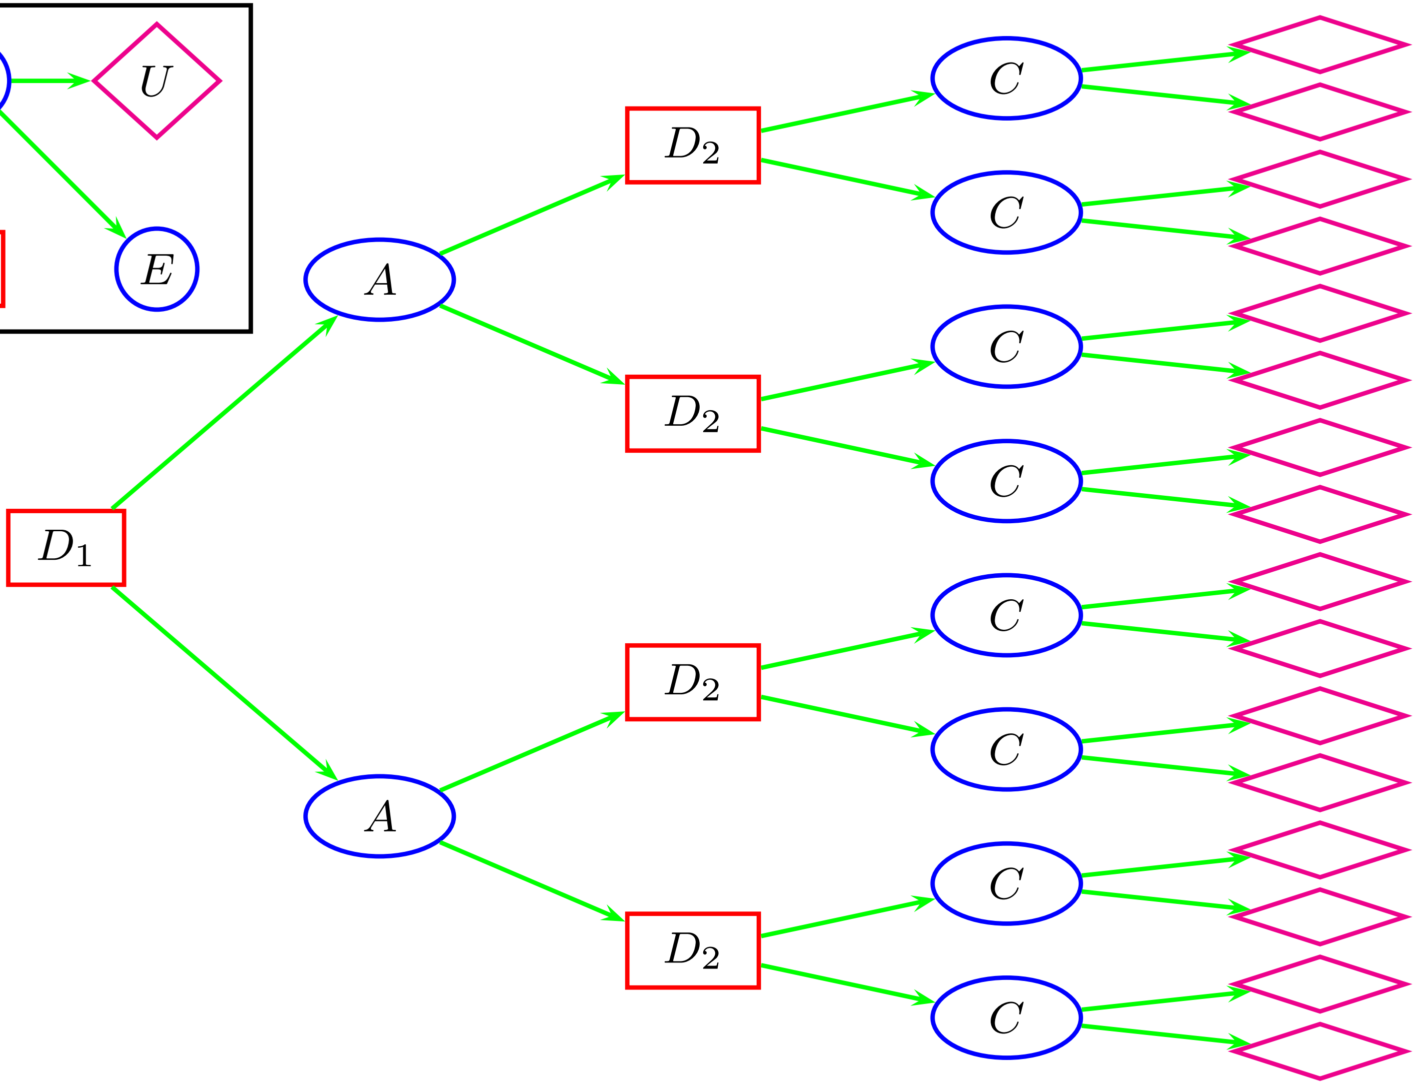

The sequence of observations and decisions is the same in all scenarios:

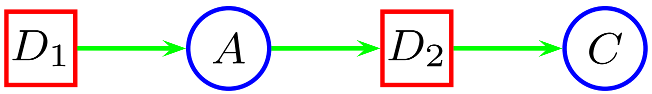

Optimal strategy I

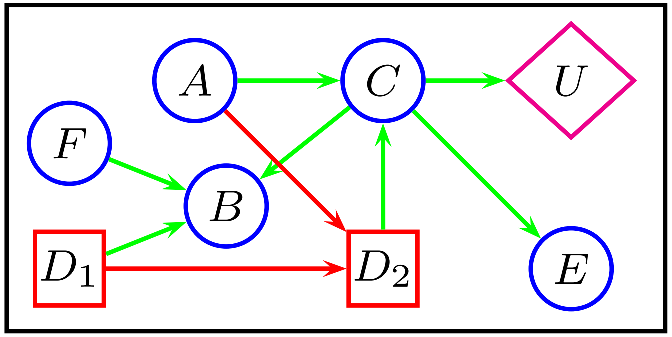

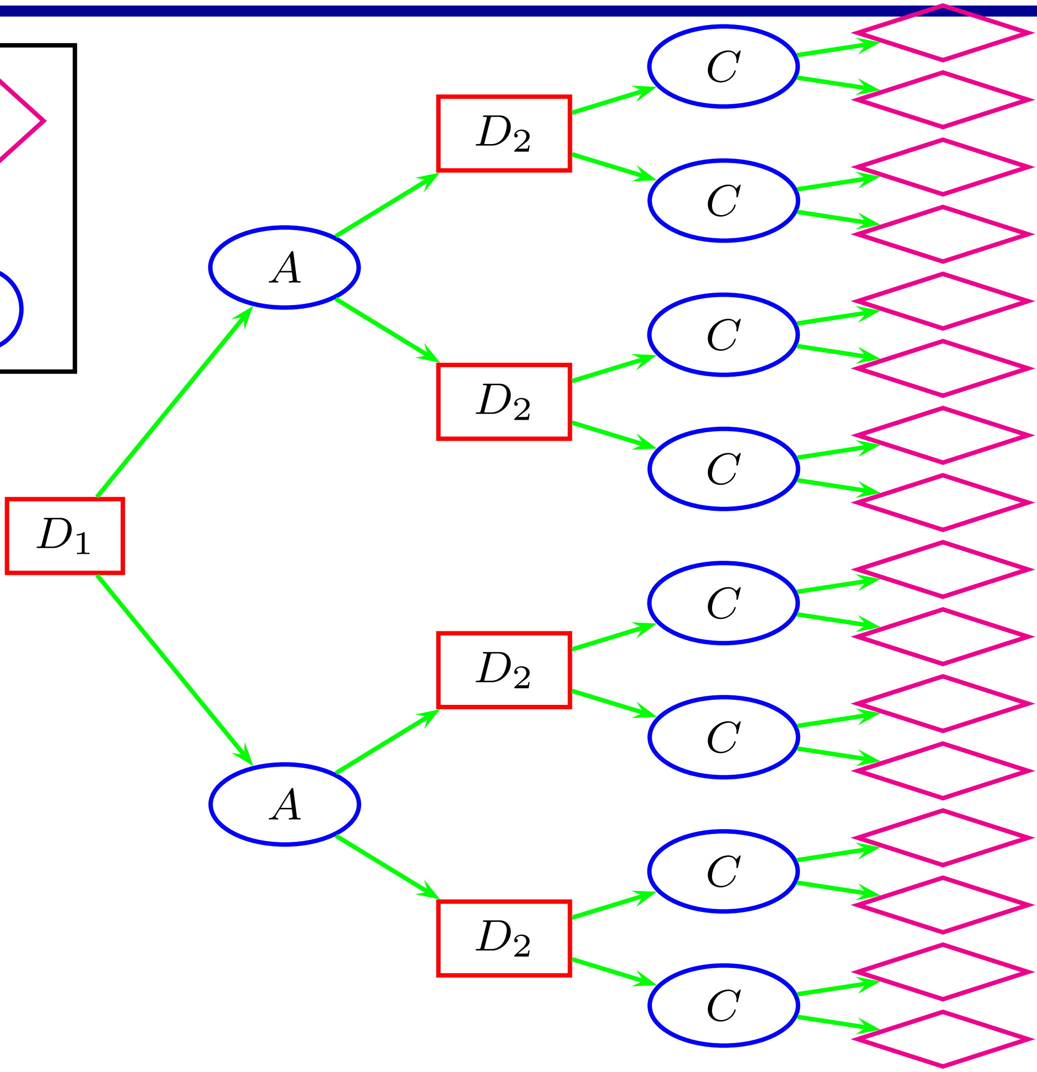

Optimal strategy I

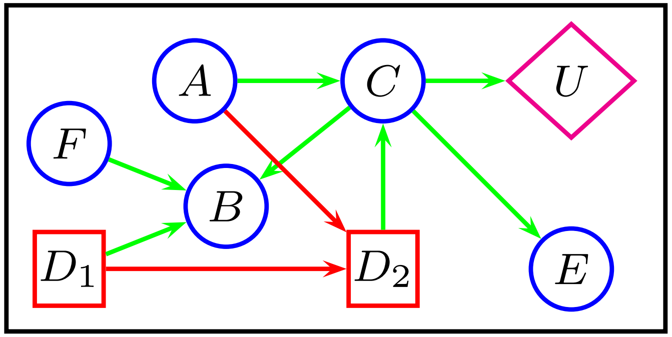

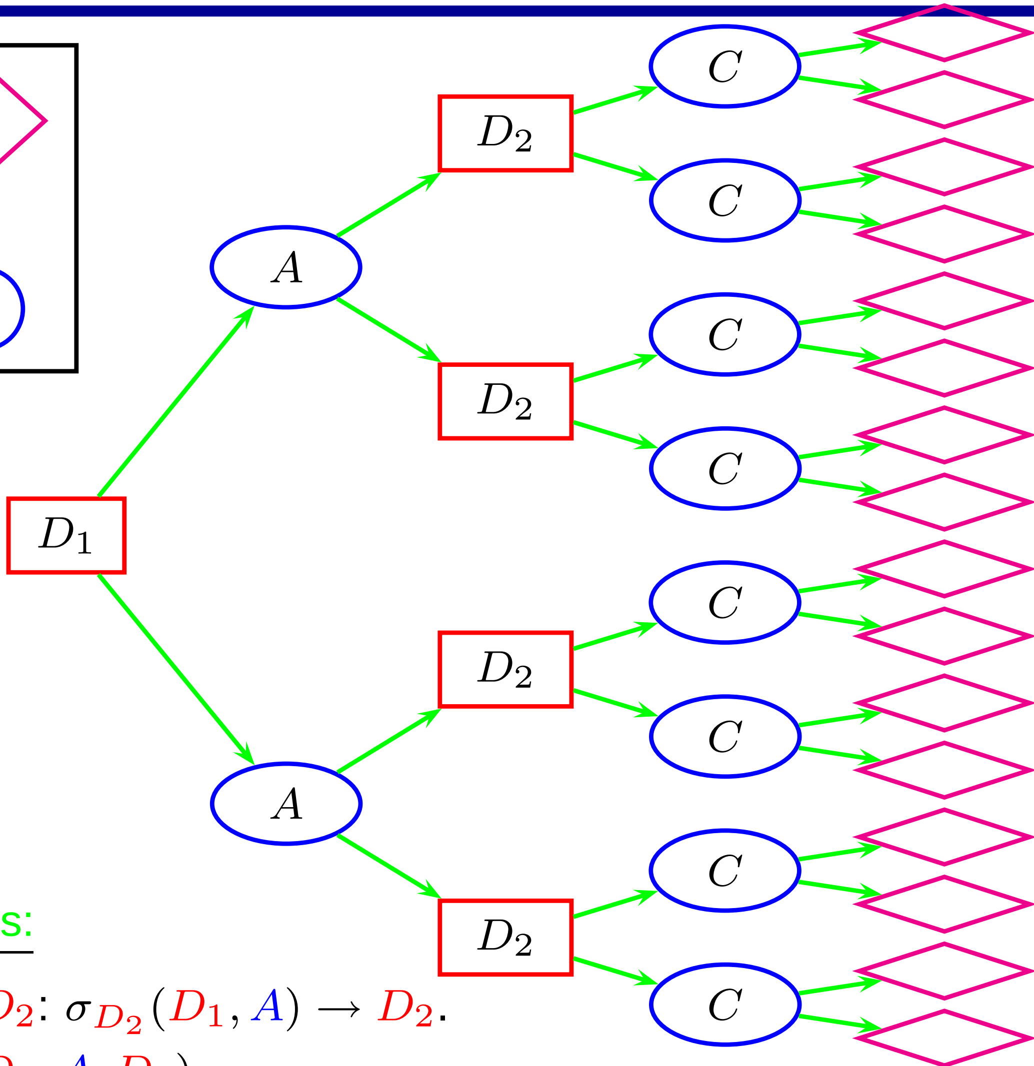

Solution for influence diagrams:

1. Determine a policy for D_2: $\sigma_{D_2}(D_1, A) \rightarrow D_2$.
   For this we need $P(C|D_1, A, D_2)$.
2. Use σ_{D_2} for determining a policy for D_1: $\sigma_{D_1} \rightarrow D_1$.
   For this we need $P(A|D_1)$.

All probabilities can be achieved from the model without folding out the decision tree.

Optimal strategy II

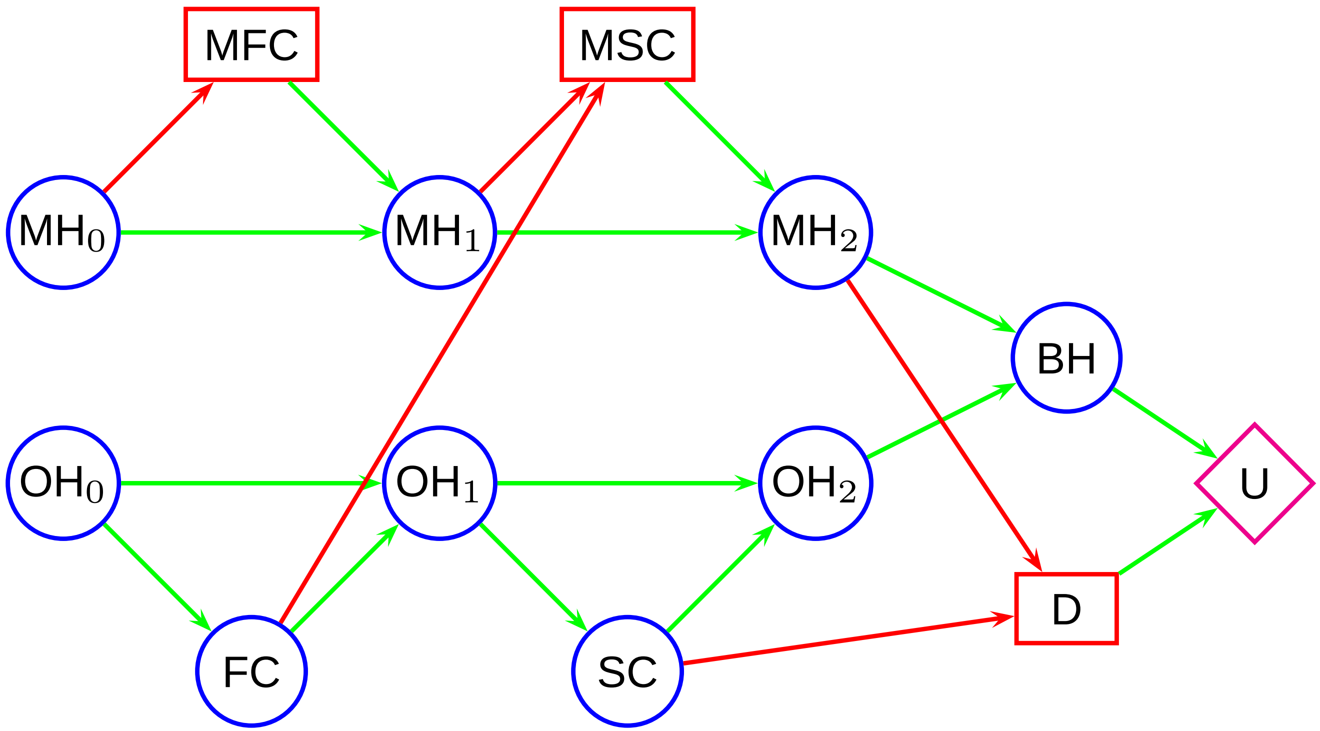

The policy for D: $\sigma_{\text{D}}(\text{MH}_0, \text{MFC}, \text{FC}, \text{MH}_1, \text{MSC}, \text{SC}, \text{MH}_2) \rightarrow \text{D}$

We request: $P(\text{BH}|\text{MH}_0, \text{MFC}, \text{FC}, \text{MH}_1, \text{MSC}, \text{SC}, \text{MH}_2, \text{D})$

Optimal strategy II

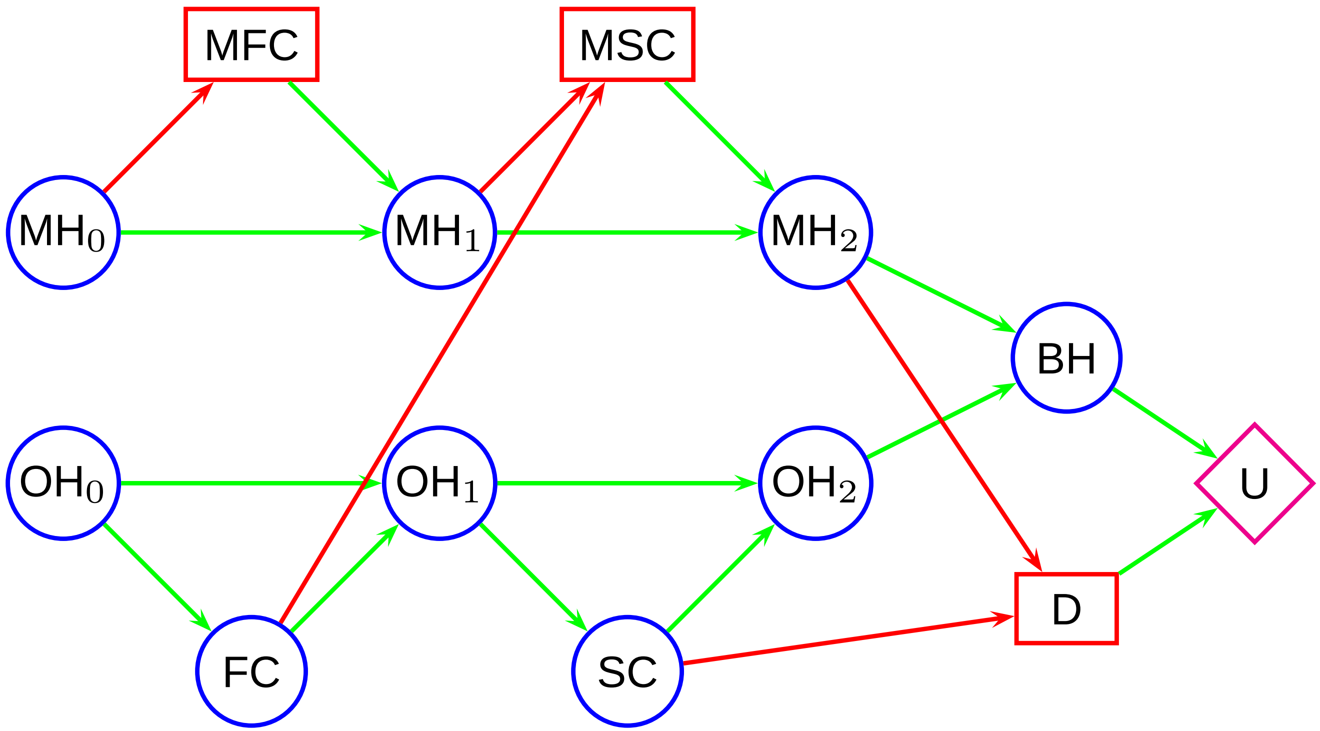

The policy for D: $\sigma_{D}(MH_0, MFC, FC, MH_1, MSC, SC, MH_2) \rightarrow D$

We request: $P(BH | MH_0, MFC, FC, MH_1, MSC, SC, MH_2, D)$

From d-separation we can find the relevant past!

Fishing in the north sea

Based on measurements, T, a quota for fishing volume, FV, for next year is decided. The amount of fish, V, and the quota determines the utility.

Fishing in the north sea

Based on measurements, T, a quota for fishing volume, FV, for next year is decided. The amount of fish, V, and the quota determines the utility.

A five year period:

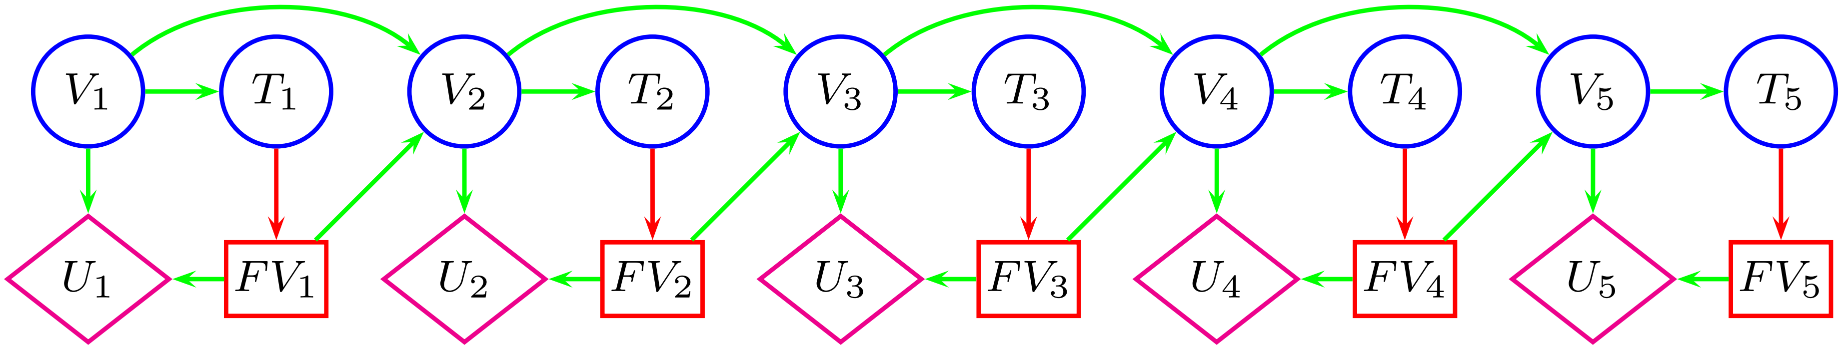

Unfortunately, the optimal policy for FV_5 depends on the entire past:

$$\sigma_{FV_5}(T_1, FV_1, T_2, FV_2, T_3, FV_3, T_4, FV_4, T_5)$$

This is intractable!

Information blocking

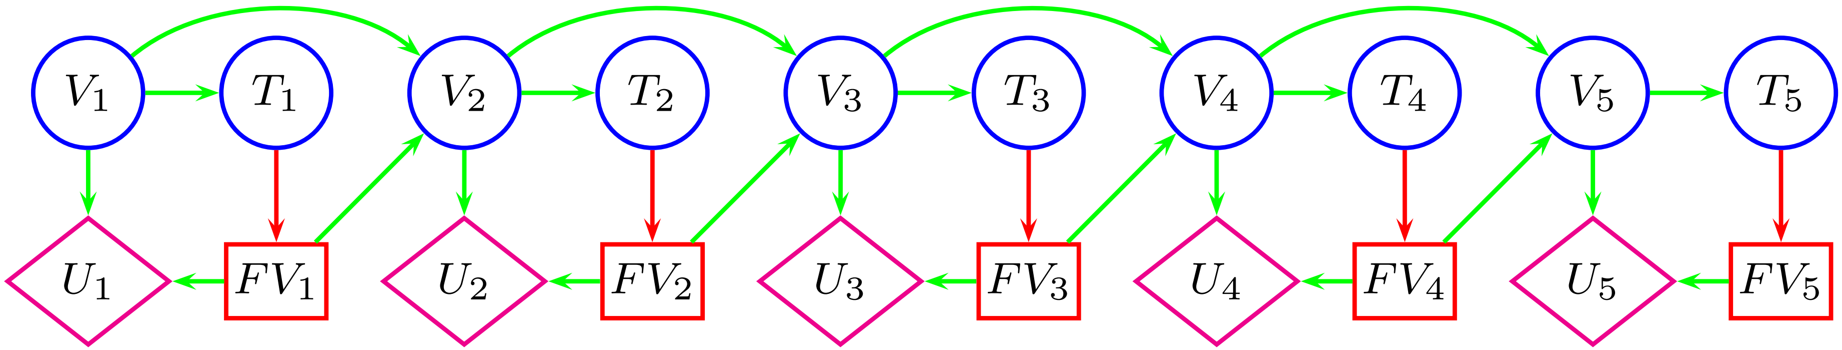

To make the calculations tractable we may use an approximation instead:

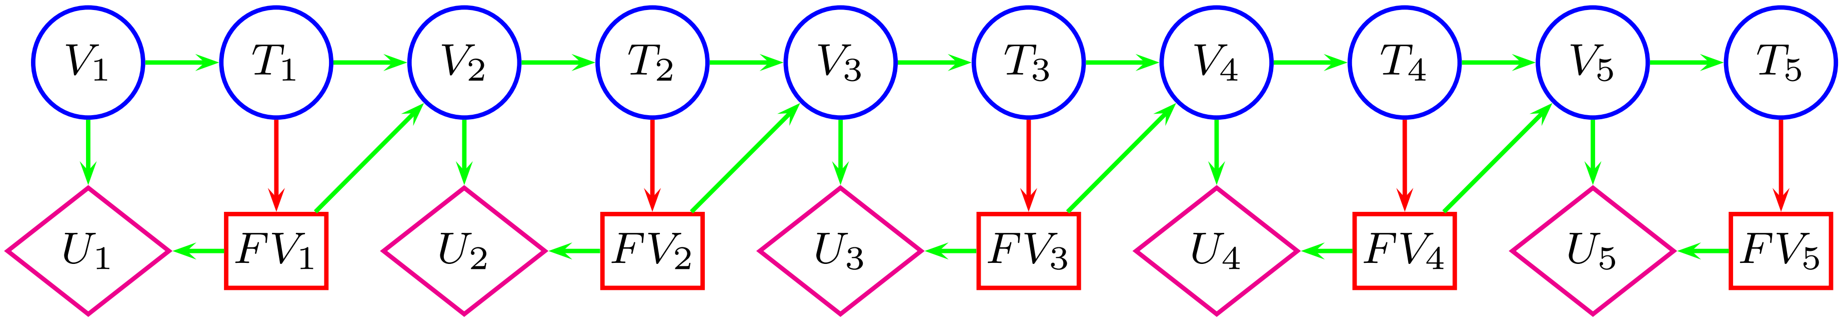

The probability $P(V_2|T_1, FV_1)$ is taken from the initial model.

The dangers of non-observed nodes

Temporal links between non-observed nodes are dangerous!

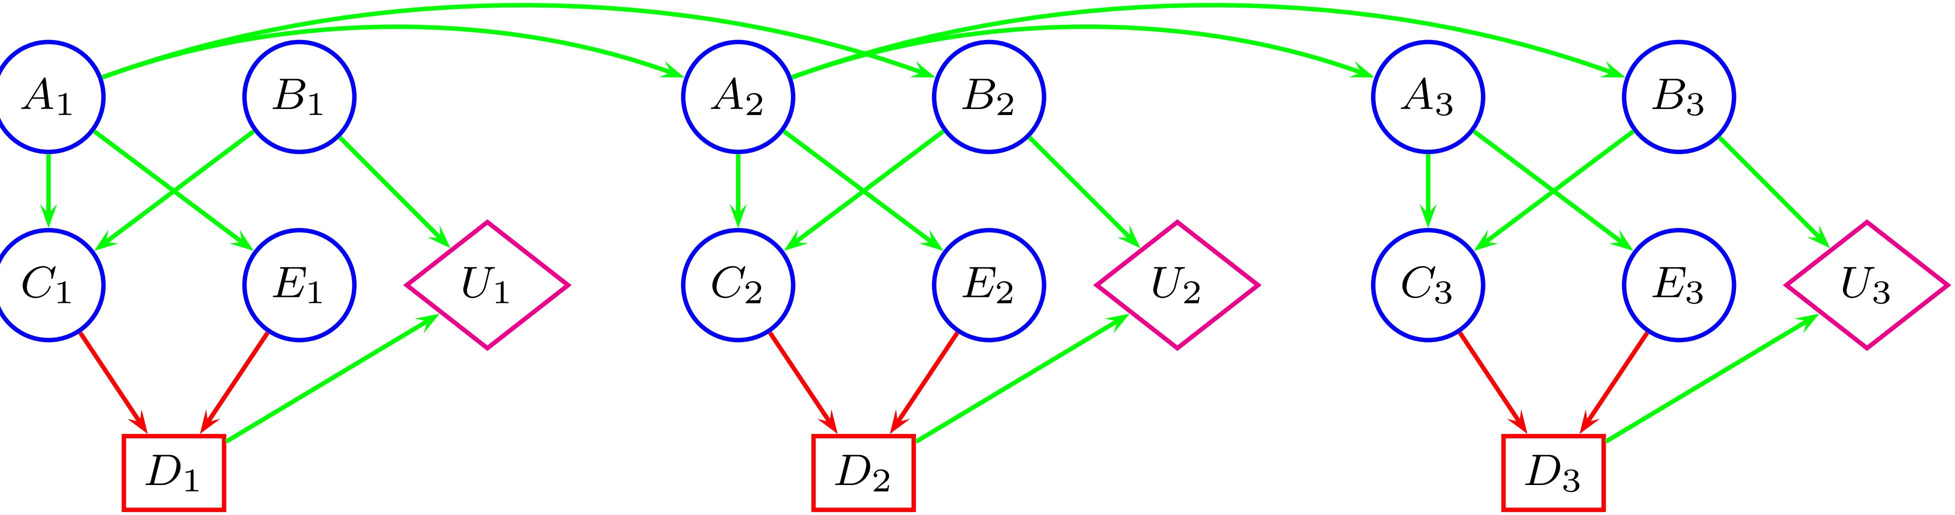

The dangers of non-observed nodes

Temporal links between non-observed nodes are dangerous!

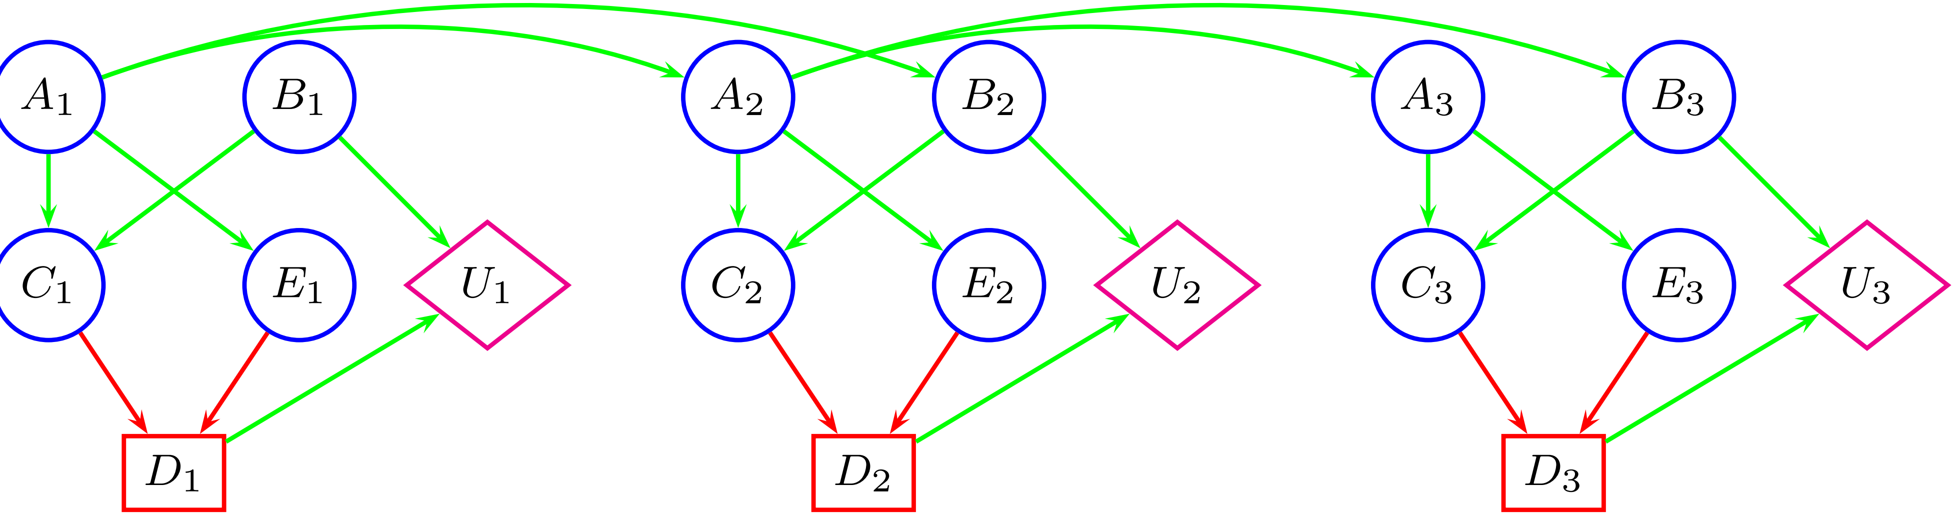

We may introduce history variables to summarize the past:

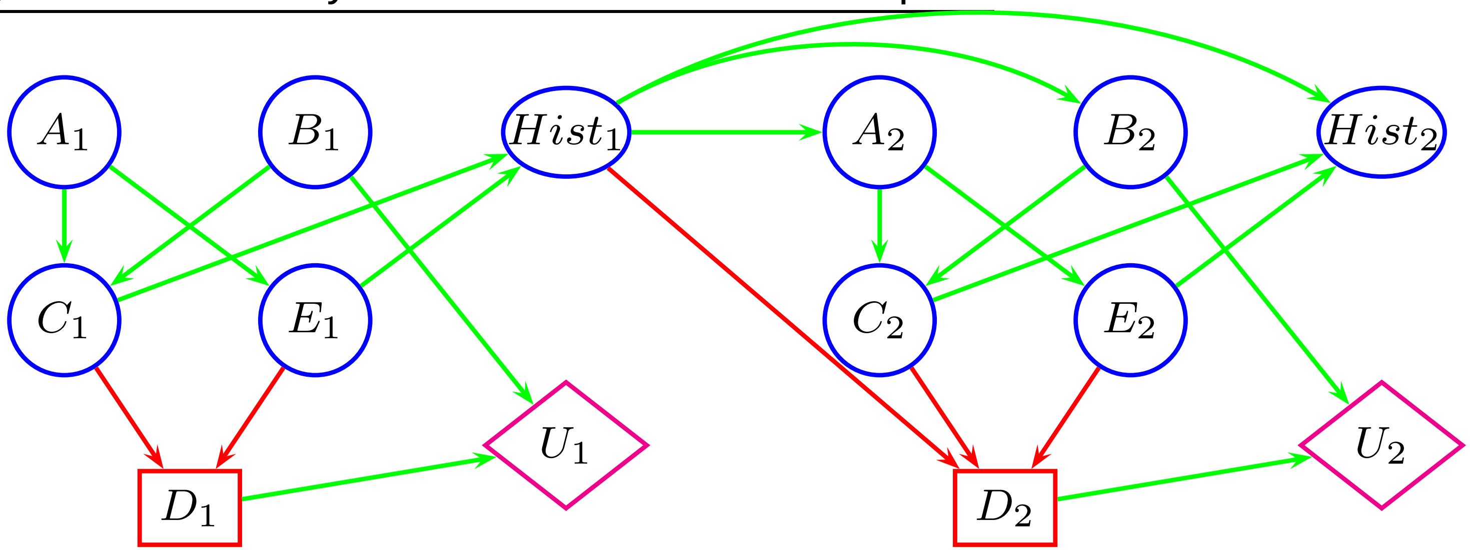

When are ID's suitable for repeated use

- The sequence of decisions $D_1, D_2, \ldots, D_n$ is fixed.
- The chance variables in I_i are always observed after D_i and before D_{i+1}.
- The decision maker remembers the past.
- The decision problem is symmetric.

  The decision-observation sequence is independent of the actual observations and decisions.

A cause of asymmetry: Test decisions

Take your temperature before deciding on aspirin.

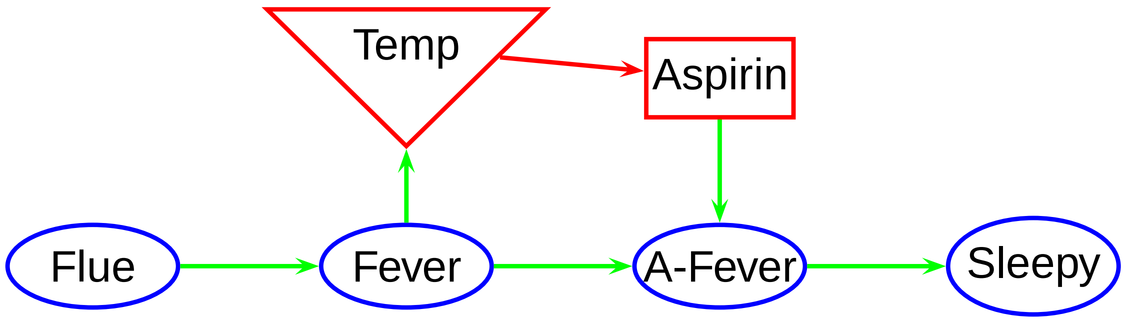

You only observe the test result (Fever) if you decide to take your temperature

A cause of asymmetry: Test decisions

But these problems can still be modeled in influence diagrams:

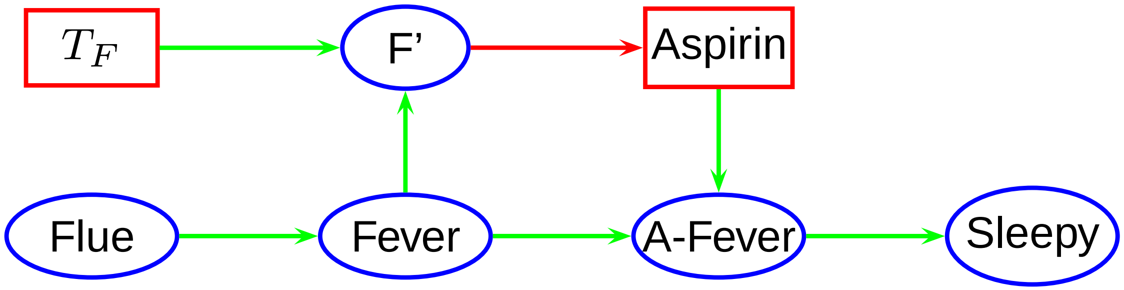

		Fever	
		y	n
T_F	y	$(1, 0, 0)$	$(0, 1, 0)$
	n	$(0, 0, 1)$	$(0, 0, 1)$

$P(F' = (\text{y}, \text{n}, \text{no-t}) | \text{Fever}, T_F)$

Transformation of test-decisions in general

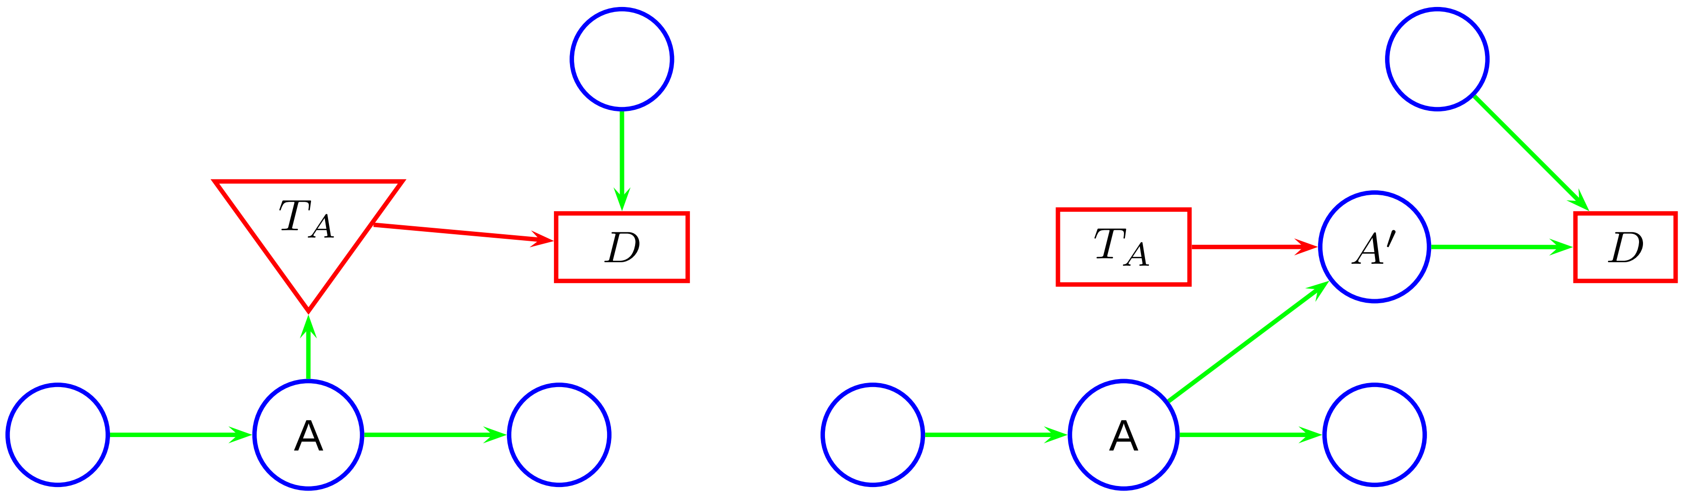

		A			
		a_1	a_2	$\cdots$	a_n
T_A	y	$(1, 0, \ldots, 0)$	$(0, 1, 0, \ldots, 0)$	$\cdots$	$(0, \ldots, 0, 1, 0)$
	n	$(0, \ldots, 1)$	$(0, \ldots, 1)$	$\cdots$	$(0, \ldots, 1)$

$$P(F' = (a_1, \ldots, a_n, no - t) | \text{A}, T_A)$$