

Making Decisions Under Uncertainty

What an agent should do depends on:

- The agent's **ability** — what options are available to it.
- The agent's **beliefs** — the ways the world could be, given the agent's knowledge. Sensing the world updates the agent's beliefs.
- The agent's **preferences** — what the agent actually wants and the tradeoffs when there are risks.

Decision theory specifies how to trade off the desirability and probabilities of the possible outcomes for competing actions.

Decision Variables

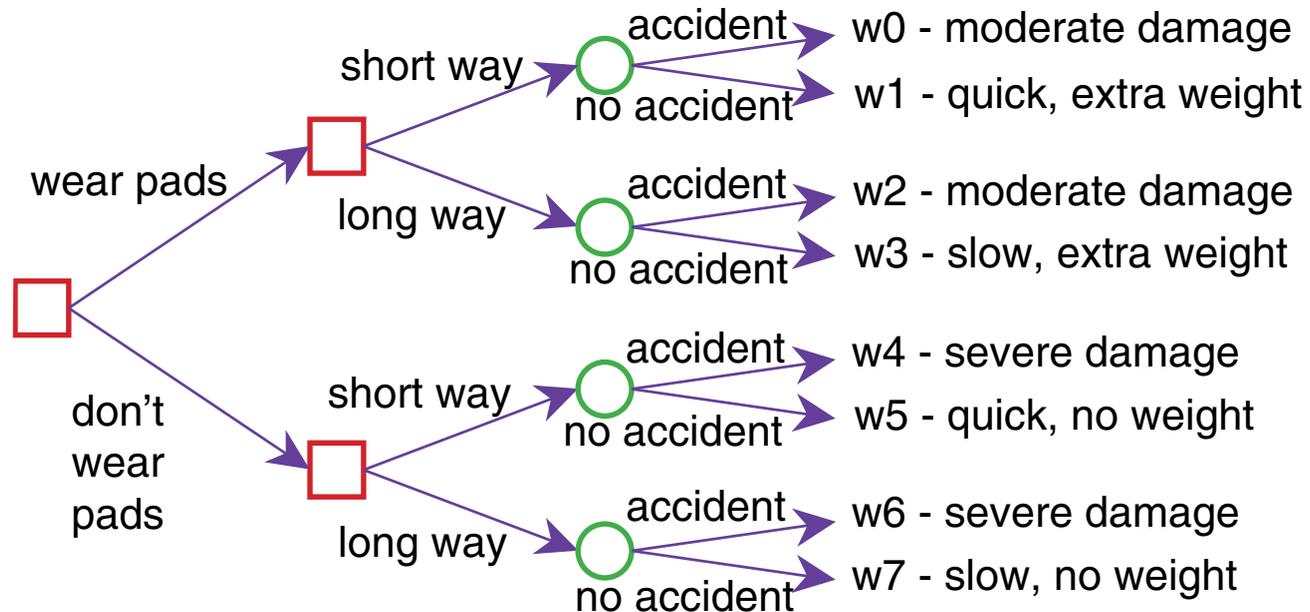
- **Decision variables** are like random variables that an agent gets to choose the value of.
- A possible world specifies the value for each decision variable and each random variable.
- For each assignment of values to all decision variables, the measures of the worlds satisfying that assignment sum to 1.
- The probability of a proposition is undefined unless you condition on the values of all decision variables.

Decision Tree for Delivery Robot

The robot can choose to wear pads to protect itself or not.

The robot can choose to go the short way past the stairs or a long way that reduces the chance of an accident.

There is one random variable of whether there is an accident.



Expected Values

- The expected value of a function of possible worlds is its average value, weighting possible worlds by their probability.
- Suppose $u(w)$ is the value of function u on world w .

➤ The **expected value** of u is

$$\mathbf{E}(u) = \sum_{\omega \in \Omega} P(\omega) \times u(\omega).$$

➤ The **conditional expected value** of u given e is

$$\mathbf{E}(u|e) = \sum_{\omega \models e} P(\omega|e) \times u(\omega).$$

Utility

- Utility is a measure of desirability of worlds to an agent.
- Let U be a real-valued function such that $U(\omega)$ represents how good the world is to an agent.
- Simple goals can be specified by: worlds that satisfy the goal have utility 1; other worlds have utility 0.
- Often utilities are more complicated: **for example** some function of the amount of damage to a robot, how much energy is left, what goals are achieved, and how much time it has taken.

Single decisions

- In a single decision variable, the agent can choose $D = d_i$ for any $d_i \in \text{dom}(D)$.
- The **expected utility** of decision $D = d_i$ is $\mathbf{E}(U|D = d_i)$.
- An **optimal single decision** is the decision $D = d_{max}$ whose expected utility is maximal:

$$\mathbf{E}(U|D = d_{max}) = \max_{d_i \in \text{dom}(D)} \mathbf{E}(U|D = d_i).$$

Sequential Decisions

- An intelligent agent doesn't make a multi-step decision and carry it out without considering revising it based on future information.
- A more typical scenario is where the agent: observes, acts, observes, acts, ...
- Subsequent actions can depend on what is observed. What is observed depends on previous actions.
- Often the sole reason for carrying out an action is to provide information for future actions. For example: diagnostic tests, spying.

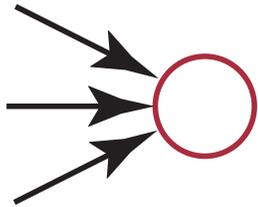
Sequential decision problems

- A **sequential decision problem** consists of a sequence of decision variables D_1, \dots, D_n .
- Each D_i has an **information set** of variables π_{D_i} , whose value will be known at the time decision D_i is made.

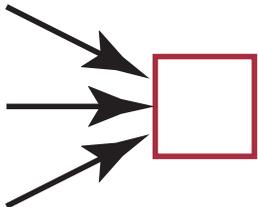
Decision Networks

- A **decision network** is a graphical representation of a finite sequential decision problem.
- Decision networks extend belief networks to include decision variables and utility.
- A decision network specifies what information is available when the agent has to act.
- A decision network specifies which variables the utility depends on.

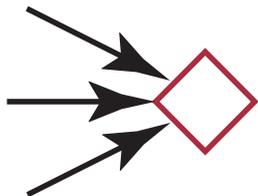
Decisions Networks



- A **random variable** is drawn as an ellipse. Arcs into the node represent probabilistic dependence.

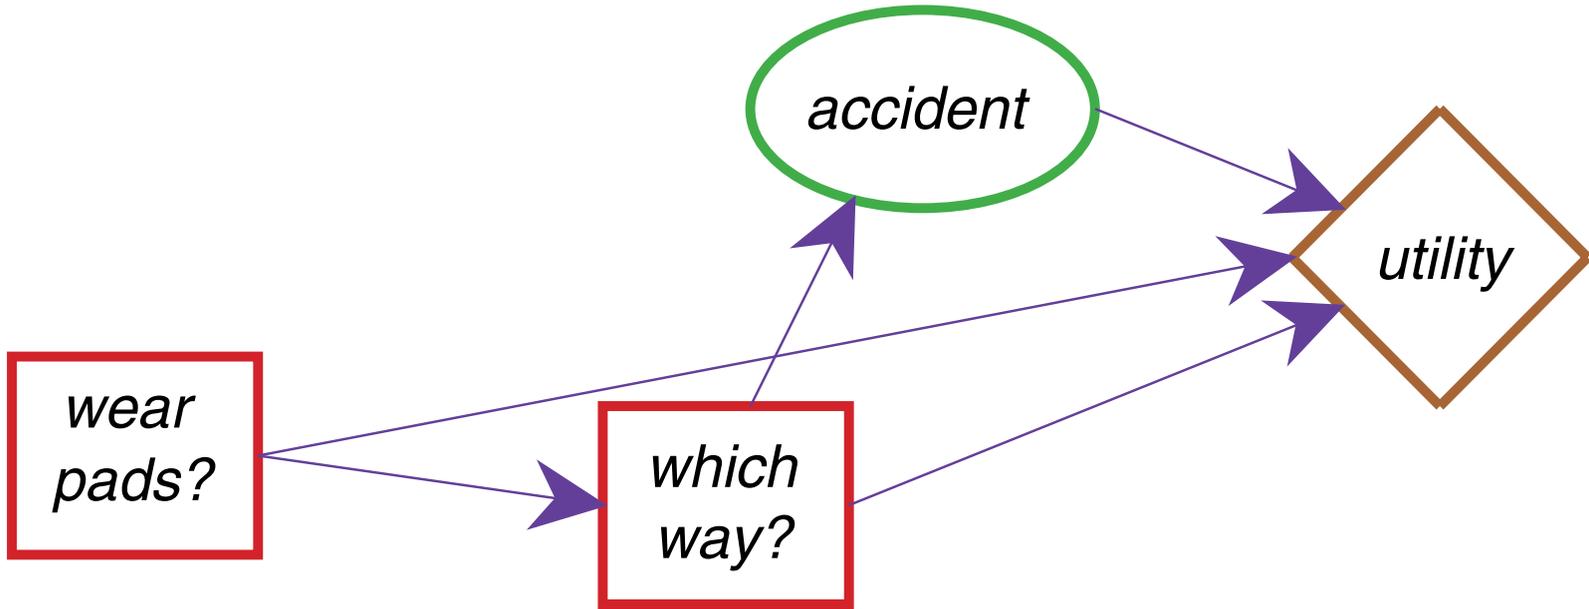


- A **decision variable** is drawn as a rectangle. Arcs into the node represent information available when the decision is made.



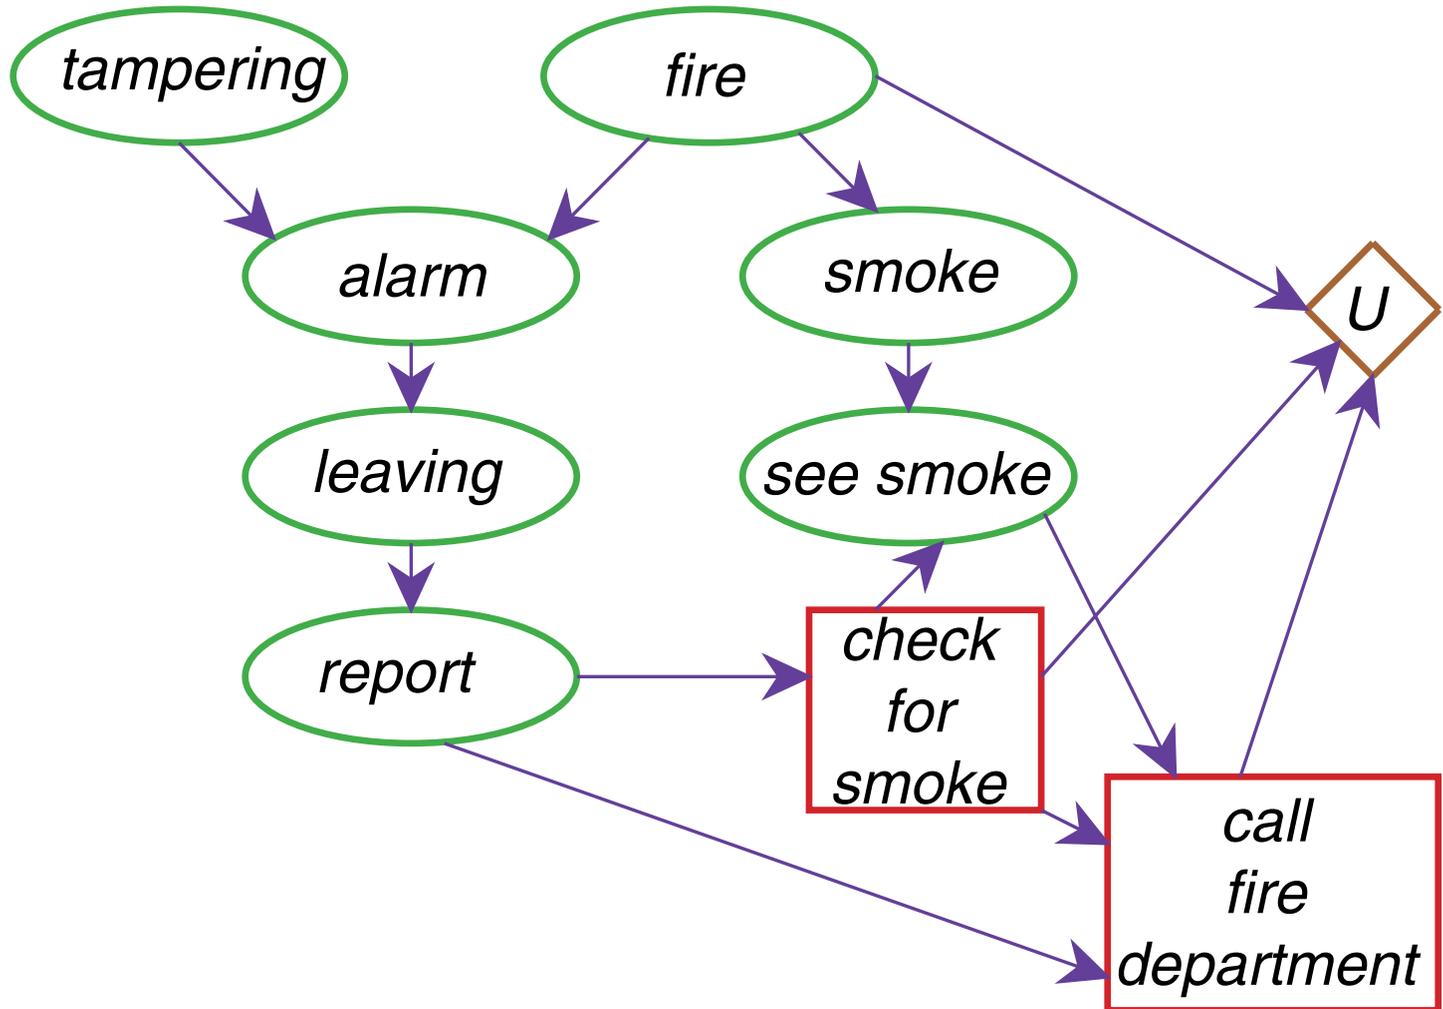
- A **value** node is drawn as a diamond. Arcs into the node represent values that the value depends on.

Example Decision Network



This shows explicitly which nodes affect whether there is an accident.

Decision Network for the Alarm Problem



What should an agent do?

- What an agent should do at any time depends on what it will do in the future.
- What an agent does in the future depends on what it did before.

Policies

- A policy specifies what an agent should do under each circumstance.
- A **policy** is a sequence $\delta_1, \dots, \delta_n$ of **decision functions**

$$\delta_i : \text{dom}(\pi_{D_i}) \rightarrow \text{dom}(D_i).$$

This policy means that when the agent has observed $O \in \text{dom}(\pi_{D_i})$, it will do $\delta_i(O)$.

Expected Value of a Policy

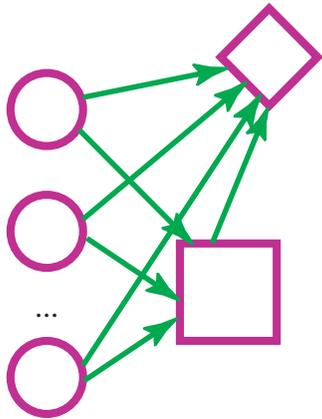
➤ Possible world ω **satisfies** policy δ , written $\omega \models \delta$ if the world assigns the value to each decision node that the policy specifies.

➤ The **expected utility of policy δ** is

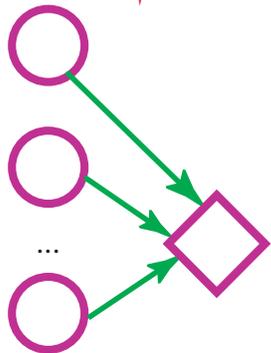
$$\mathbf{E}(U|\delta) = \sum_{\omega \models \delta} U(\omega) \times \mu(\omega),$$

➤ An **optimal policy** is one with the highest expected utility.

Finding the optimal policy



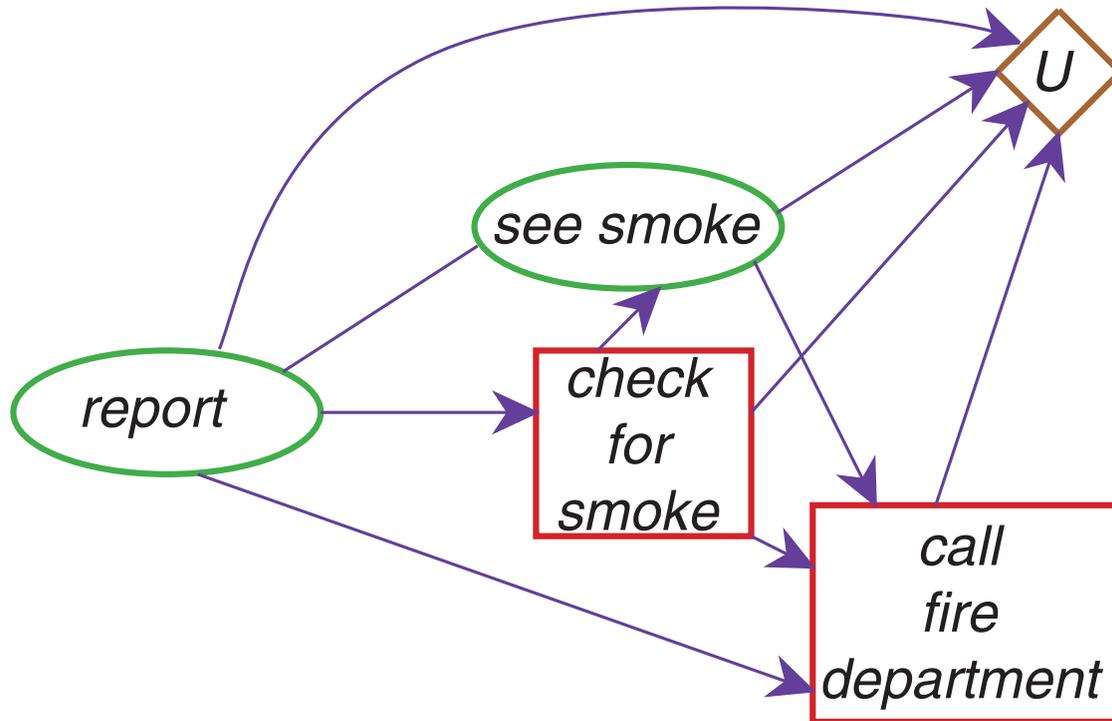
- If value node is only connected to a decision node and (some of) its parents
➡ select a decision to maximize value for each assignment to the parent.



- If it isn't of this form, eliminate the non-observed variables.
- Replace decision node with value node.
- Repeat till there are no more decision nodes.

Reduced Alarm Example

Eliminate the non-observed variables for the final decision.



Complexity of finding the optimal policy

- If there are k binary parents, there are 2^k optimizations.
- If there are b possible actions, there are b^{2^k} policies.
- The dynamic programming algorithm is much more efficient than searching through policy space.

Value of Information

- The value of information X for decision D is the utility of the the network with an arc from X to D minus the utility of the network without the arc.
- The value of information is always non-negative.
- It is positive only if the agent changes its action depending on X .