Decision Trees

Ricard Torres

ITAM

Economía Financiera, 2015

Decision trees

A decision tree is a set of nodes, some of which are connected by edges.

There is a distinguished node, the root, such that there is a unique path that connects any other node with it.

Nodes can be either:

1. Decision nodes, in which the decision maker must choose an alternative
2. Nature or chance nodes, that represent a random event
3. Terminal nodes, that are labelled with the utility the decision maker attains if he or she reaches it

At each decision node, the decision maker must choose one of a set of possible alternatives or actions; we associate each such action with an edge coming out (ie, away from the root) of the decision node.

Information sets

An information set is a set of decision nodes that the decision maker is unable to distinguish.

That is, when one of the nodes is reached, the decision maker knows he/she is somewhere within the information set, but does not know at which of the nodes in it.

The actions available at any two decision nodes that belong to the same information set are the same.

For convenience, we may think of all decision nodes as belonging to some information set, even if it contains only one node.

The edges that come out (away from the root) of a nature node correspond to the different results of the random event, and they are labelled with the corresponding probabilities.

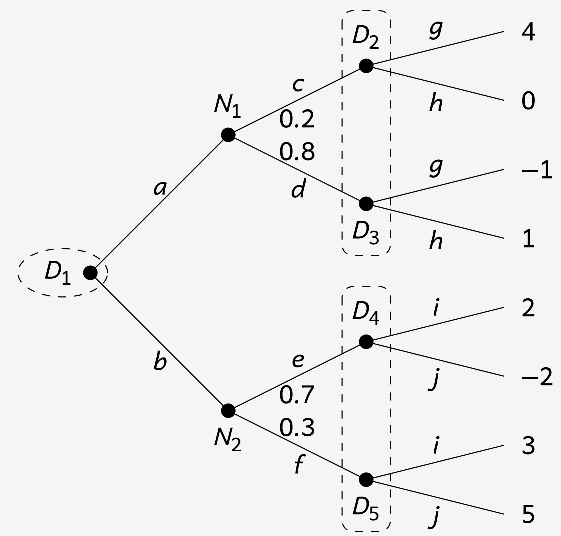

There are five decision nodes, labelled D_1 (the root) through D_5, two nature nodes, N_1 and N_2, and eight terminal nodes, labelled with utilities. There are three information sets: $\{D_1\}$, $\{D_2, D_3\}$, and $\{D_4, D_5\}$. The actions available at nodes in the same information set are the same.

Strategies

A strategy is a rule that associates with each decision node one of the actions available there.

We impose the restriction that the action chosen at any two decision nodes that belong to the same information set must be the same, otherwise we would have an inconsistency with our interpretation of information sets.

A given strategy gives rise to a certain probability distribution over terminal nodes.

So for each strategy, we may compute the corresponding expected utility.

The objective of the decision maker is to find an optimal strategy, ie, one for which the expected utility is highest.

Perfect recall

We say that a decision tree satisfies perfect recall if, whenever a decision node of follows a particular action, then all of the other decision nodes that are in the same information set follow the same action.

This implies that, for any information set and any strategy, either there is a positive probability that all nodes of the information set can be reached, or none of them can.

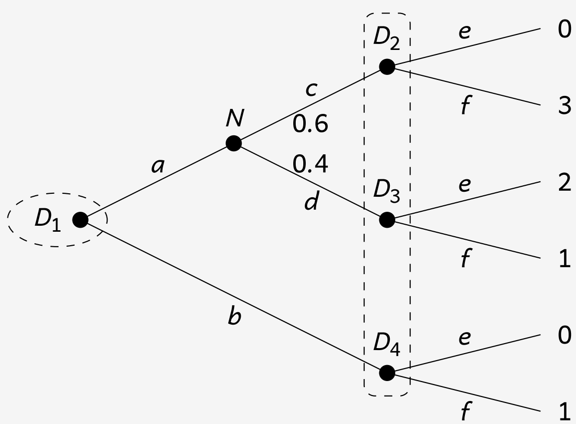

Without perfect recall

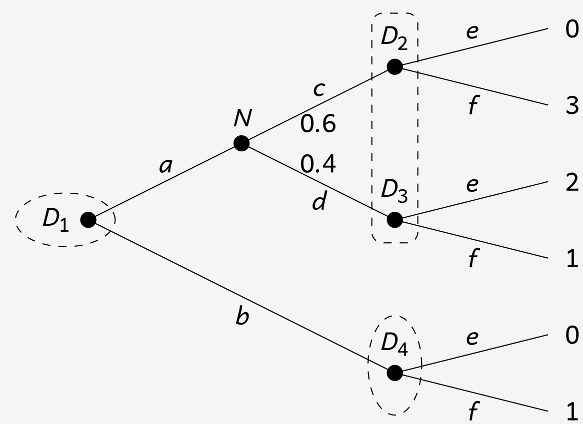

With perfect recall

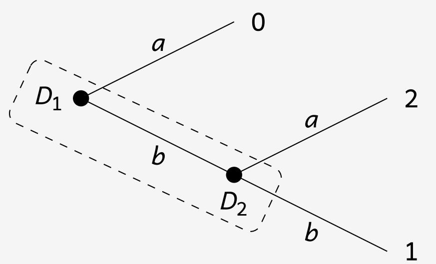

Absentmindedness (no perfect recall)

Probability assessments

A probability assessment for a given information set is a probability distribution over its nodes.

A classical result in Game Theory, Kuhn's Theorem, implies that if a decision tree satisfies perfect recall, there is a unique probability assessment that is consistent with all strategies under which the information set is reached with positive probability.

We find those probabilities by application of Bayes' formula.

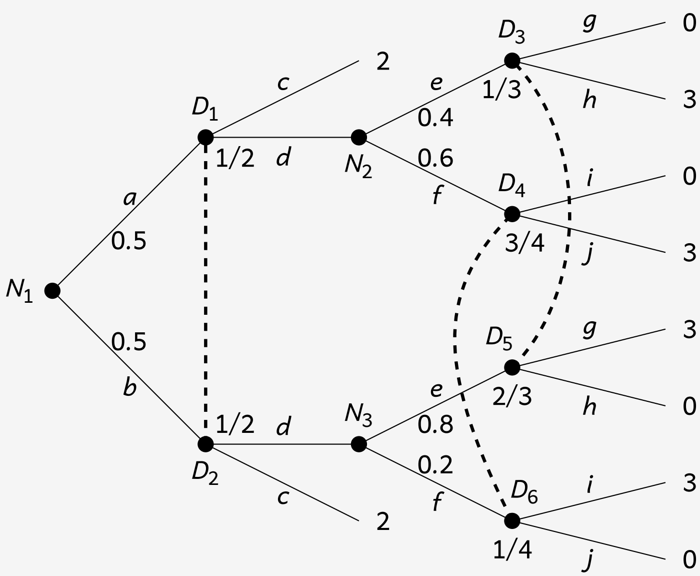

Backward induction

The procedure of backward induction consists of breaking the problem of finding an optimal strategy for the decision tree into a sequence of simple subproblems.

We proceed backwards (towards the root) from decision nodes (ie, information sets) that are closest to terminal nodes.

In order to apply this procedure, we need perfect recall to be satisfied, because we must have a unique consistent probability assessment at each information set.

1. Consider all information sets that are closest to terminal nodes.
2. In each of these information sets, find an optimal action, ie, one that yields the highest expected utility.
3. After choosing such an optimal action, substitute the nodes in the information set by the (expected) utility that corresponds to that optimal action.
4. Iterate the procedure until eventually reaching the root
5. The sequence of actions found in the different steps is an optimal strategy for the original tree.

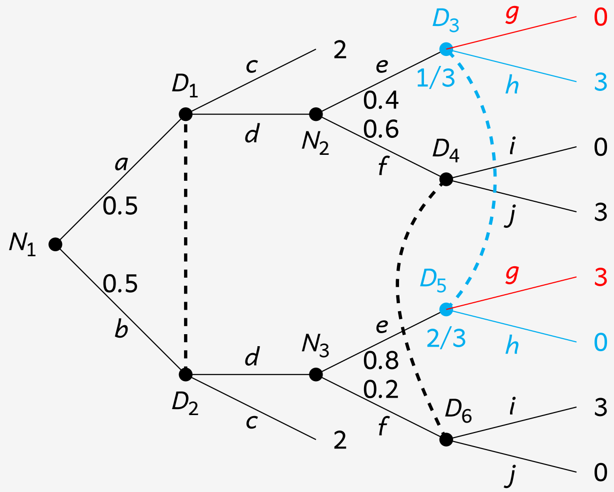

Start from the information set $\{D_3, D_5\}$. The probability assessment assigns probability 1/3 to D_3 and 2/3 to D_5. The expected utility if g is chosen is $1/3 \cdot 0 + 2/3 \cdot 3 = 2$, and if h is chosen is $1/3 \cdot 3 + 2/3 \cdot 0 = 1$, so g is an optimal action. Replace node D_3 by a terminal node with utility 0, and D_5 by a terminal node with utility 3.

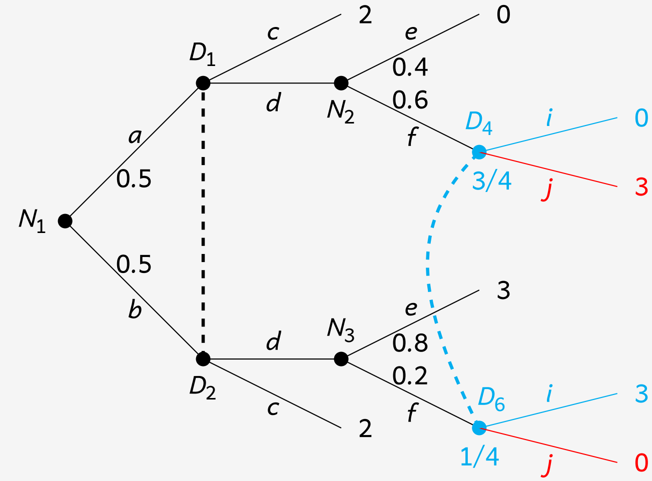

Consider the information set $\{D_4, D_6\}$. The probability assessment assigns probability 3/4 to D_4 and 1/4 to D_6. Hence, the expected utility of action i is 3/4 and the expected utility of action j is 9/4, so the latter is optimal. Replace node D_4 by a terminal node with utility 3 and node D_6 by a terminal node with utility 0.

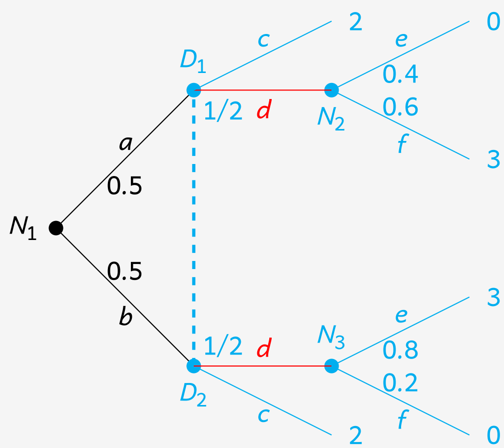

The probability assessment assigns equal probability to both decision nodes. Action c yields utility 2, and action d yields an expected utility:

$$0.5 \cdot 0.4 \cdot 0 + 0.5 \cdot 0.6 \cdot 3 + 0.5 \cdot 0.8 \cdot 3 + 0.5 \cdot 0.2 \cdot 0 = 2.1.$$

The optimal action is d.

Optimal strategy for the original tree: (d, g, j). Maximal utility: 2.1.