Algorithm 23.1 Finding the MEU strategy in a decision tree

Procedure MEU-for-Decision-Trees (
  $\mathcal{T}$ // Decision tree
)
  $L \leftarrow$ *Leaves*$(\mathcal{T})$
  for each node $v \in L$
    Remove v from L
    Add v's parents to L
    if v is a leaf **then**
      $\mathrm{MEU}_v \leftarrow U(v)$
    else if v belongs to nature **then**
      $\mathrm{MEU}_v \leftarrow \sum_{c \in \mathcal{C}[v]} P_v(c) \mathrm{MEU}_{succ(v,c)}$
    else // v belongs to the Agent
      $\sigma(v) \leftarrow \arg\max_{c \in \mathcal{C}[v]} \mathrm{MEU}_{succ(v,c)}$
      $\mathrm{MEU}_v \leftarrow \mathrm{MEU}_{succ(v,succ(v,))}$
  return (σ)

The optimal action for the agent is the one leading to the child whose MEU is largest, and the MEU accruing to the agent is the MEU associated with that child. The algorithm is shown in algorithm 23.1.

23.2 Influence Diagrams

The decision-tree representation is a significant improvement over representing the problem as a set of abstract outcomes; however, much of the structure of the problem is still not made explicit. For example, in our simple Entrepreneur scenario, the agent's utility if he founds the company depends only on the market demand M, and not on the results of the survey S. In the decision tree, however, the utility values appear in four separate subtrees: one for each value of the S variable, and one for the subtree where the survey is not performed. An examination of the utility values shows that they are, indeed, identical, but this is not apparent from the structure of the tree.

The tree also loses a subtler structure, which cannot be easily discerned by an examination of the parameters. The tree contains four nodes that encode a probability distribution over the values of the market demand M. These four distributions are different. We can presume that neither the survey nor the agent's decision has an effect on the market demand itself. The reason for the change in the distribution presumably arises from the effect of conditioning the distribution on different observations (or no observation) on the survey variable S. In other words, these distributions represent $P(M \mid s^0)$, $P(M \mid s^1)$, $P(M \mid s^2)$, and $P(M)$ (in the branch where the survey was not performed). These interactions between these different parameters are obscured by the decision-tree representation.