11

Methods for Analyzing Decision Problems

The primary issue in dealing with a decision problem is to determine an optimal strategy, but other issues may be relevant. This chapter deals with value of information, the relevant past and future for a decision, and the sensitivity of decisions with respect to parameters.

11.1 Value of Information

As mentioned previously, there is a difference between action decisions and test decisions; action decisions may result in a state change for some of the variables, whereas test decisions are decisions to look for more evidence. A typical situation is that you may choose among some actions, but before deciding on the action you also have the option to perform some tests. The question is which test to perform, if any.

These types of decision problems can be characterized as asymmetric decision problems, since they contain at least two types of asymmetry: structural asymmetry (if you decide not to perform a test, the result is never observed), and order asymmetry (the sequence of tests may be unspecified). However, rather than looking at this as a general asymmetric decision problem we shall in this section deal directly with the problem by considering the actual value of information.

11.1.1 Test for Infected Milk?

Consider again the infected milk problem described in Example 9.1, where we assume that the farmer only has one test, which costs 6 cents and has a false positive/negative frequency of 0.01. The test situation corresponds to choosing between the two influence diagrams in Figure 11.1, where the leftmost influence diagram incurs an additional cost of 6 cents.

To establish the utilities, let us assume that the farmer has clean milk from the 49 other cows. If the farmer pours the milk into the container, he

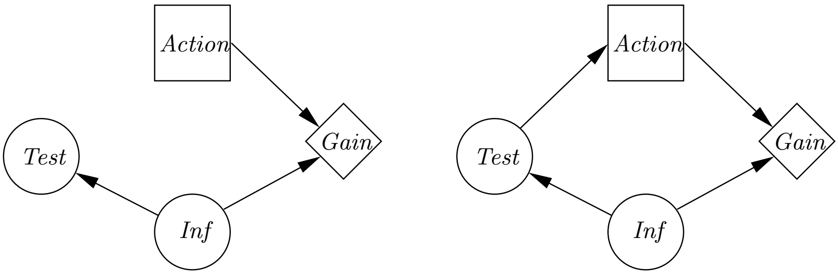

Fig. 11.1. The test scenario for infected milk corresponds to choosing between the influence diagrams, but by choosing the rightmost model you have to pay an additional 6 cents for the test.

will gain \$100 if it is not infected, and he will gain nothing if it is infected. If he throws the milk away, he will gain \$98 regardless of the state of the milk.

If the farmer does not perform a test, the probability of the milk being infected is 0.0007. The expected utility of pouring the milk into the container is

$$\begin{aligned}\mathrm{EU}(\mathit{pour}) &= P(\mathit{Inf}=\mathit{no})U(\mathit{Inf}=\mathit{no}) + P(\mathit{Inf}=\mathit{yes})U(\mathit{Inf}=\mathit{yes}) \\ &= 0.9993 \cdot 100 + 0.0007 \cdot 0 = 99.93.\end{aligned}$$

Because the expected utility of pouring the milk into the container is larger than 98, he will do this.

The reason for performing the test is that some outcome will make the farmer change the decision. To put it in another way, if the decision is the same regardless of the outcome of the test, then it is not worth the bother to perform it. Only a positive test result may change the current decision. An easy calculation yields $P(\mathit{clean}\,|\,\mathit{pos}) = 0.935$. The expected utility of pouring given a positive test result is

$$\begin{aligned}\mathrm{EU}(\mathit{pour}\,|\,\mathit{Test}=\mathit{pos}) &= P(\mathit{Inf}=\mathit{no}\,|\,\mathit{Test}=\mathit{pos})U(\mathit{Inf}=\mathit{no}) \\ &\quad + P(\mathit{Inf}=\mathit{yes}\,|\,\mathit{Test}=\mathit{pos})U(\mathit{Inf}=\mathit{yes}) \\ &= 0.935 \cdot 100 + 0.065 \cdot 0 = 93.5,\end{aligned}$$

so if the test is positive, the farmer changes his decision. The next concern is whether the test is worth its price. There are two possibilities: the test is negative and the milk is poured, or the test is positive and the milk is thrown away. The probability of the first possibility can be calculated from the specified probabilities and is 0.9893, and the second possibility has the probability 0.0107. Hence, the expected benefit of performing the test is

$$\mathrm{EU}(\mathit{Test}) = 0.9893 \cdot 100 + 0.0107 \cdot 98 = 99.98.$$

The farmer has an increase in expected utility only from 99.93 to 99.98 at the price of \$0.06, so it is not worth while to perform the test.

11.1.2 Myopic Hypothesis-Driven Data Request

In the preceding example, we attached a value to the various information scenarios, namely the expected utility of the optimal action. The driving force for evaluating the information scenario was how the distribution of the variable *Infected?* was affected by the test. We call this kind of data request *hypothesis-driven*: the distribution of a hypothesis variable H is the target of the analysis. To formulate it in more general terms, there is a *value function* V attached to the distribution $P(H)$. Usually, the value function is a maximal utility for a decision variable D:

$$V(P(H)) = \max_{d \in D} \sum_{h \in H} U(d,h) P(h \mid d).$$

Note that here we use $V(P(H))$ rather than $\mathrm{EU}(D)$ to emphasize that we are looking at the decision problem in a value-of-information context. If test T with cost C_T yields the outcome t, then the value of the new information scenario is

$$V(P(H \mid t)) = \max_{d \in D} \sum_{h \in H} U(d,h) P(h \mid t, d).$$

Since the outcome of T is not known, we can calculate only the *expected value*:

$$EV(T) = \sum_{t \in T} V(P(H \mid t)) \cdot P(t \mid d).$$

The *expected benefit* of performing test T is

$$EB(T) = EV(T) - V(P(H)).$$

The *expected profit* is

$$EP(T) = EB(T) - C_T.$$

The hard part in the calculations is the calculation of $P(H \mid T, D)$. This will usually require one propagation per state of T and D. Very often, the action has no impact on the hypothesis, and this reduces the work.

If there are several possible tests to perform, we are faced with a new problem. We may calculate the expected profit of each test, but we cannot be sure that the best choice is the one with the highest expected positive profit. A proper analysis of the data-request situation should consist in an analysis of all possible sequences of tests (including the empty sequence). To avoid such an intractable analysis, the so-called *myopic* approximation is often used: If you are allowed to perform at most one test, which one will you choose? The answer is the one with the highest expected profit if it is positive.

The myopic approach does not guarantee an optimal sequence (see also Section 10.5.4 in a troubleshooting context). Sometimes a single test does not yield anything by itself, whereas its outcome may be crucial for selecting a second very informative test.

Now, assume you have the tests $T_1, \ldots, T_m$, let H be the hypothesis variable, and assume that the action has no impact on H. To calculate the expected profit for all tests, you need $P(H \mid T_i)$ for each T_i. This can be achieved by propagating each possible outcome of each possible test. It can also be achieved in a simpler way. By propagating the states of H rather than the states of the tests, we get $P(T_i \mid H)$ for all T_i. Bayes' rule yields

$$P(H \mid T_i) = P(T_i \mid H)\frac{P(H)}{P(T_i)}.$$

Because $P(T_i)$ and $P(H)$ are available initially, we do not need more propagations than there are states in H.

The junction tree framework can also be used to perform some types of value of information analysis. For example, consider the influence diagram in Figure 11.2, where the variable C is observed prior to D_3.

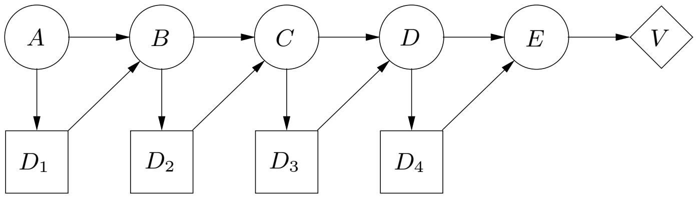

Fig. 11.2. An influence diagram.

The observation may improve the decision D_3 and yield a higher expected utility. The observation has a cost, though, but since it does not affect the strategy, it is not part of the model. Assume now that we wish to analyze how much the observation actually improves the expected utility. The situation in which C is not observed is reflected in the influence diagram in Figure 11.3. If the difference in MEU between the two influence diagrams is smaller than the cost of observing, then it does not pay to perform the test.

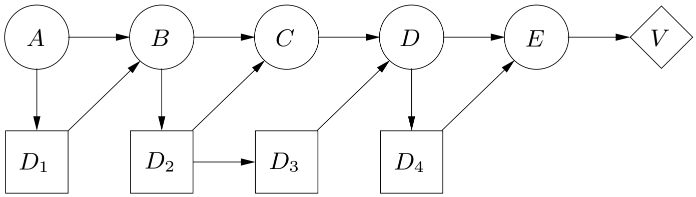

Fig. 11.3. An influence diagram for the scenario from Figure 11.2 but with C not observed.

If we assume that the cost of observing is not dependent on the timing, the MEU cannot get higher by delaying an observation that must eventually be performed. Therefore, the only option we have is either to observe as soon as possible or never to observe.

Using a method similar to propagation of variables as described in Section 5.2, the calculation of the various MEUs can be joined in one strong junction tree. Perform a strong triangulation for the influence diagram modeling that the observations have not been performed (that is, with the chance variables under analysis as members of $\mathcal{I}_n$) and construct the strong junction tree. When solving the influence diagram corresponding to an observation of the chance node C just before deciding on D_i, you use the same strong junction tree. However, you defer the elimination of C until D_i has been eliminated. Figure 11.4 shows the influence diagram from Figure 10.5, where an observation is optional for several variables as indicated by the dashed arrows. The reader may check that you can solve all influence diagrams corresponding to all combinations of possible observations through delayed elimination in the strong junction tree in Figure 10.8.

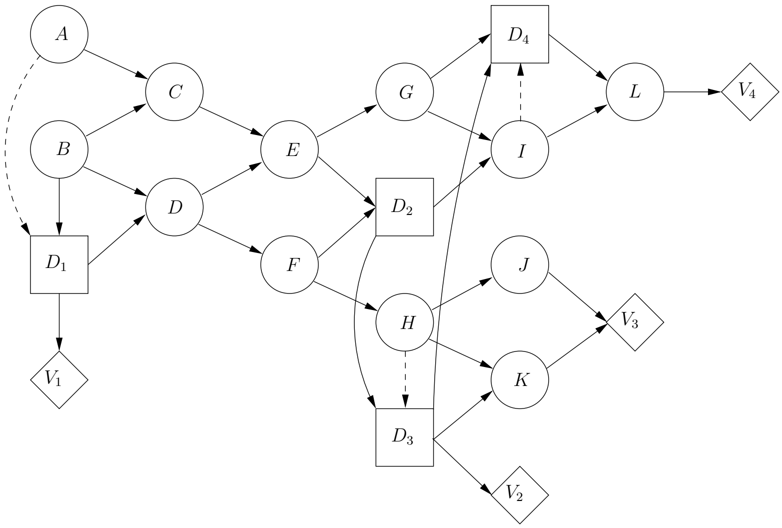

Fig. 11.4. An influence diagram with the option of not observing A, H, and I.

11.1.3 Non-Utility-Based Value Functions

If there is no proper model for actions and utilities, the reason for acquiring more information is to decrease the uncertainty of the hypothesis. This means that you will give high values to probabilities close to zero and one, while probabilities in the middle area should have low values. A classical function with this property is entropy (see Section 8.4).

The formula for the entropy of a distribution over H is in Section 8.4 defined as

$$\mathrm{Ent}(P(H)) = -\sum_{h \in H} P(h) \log_2(P(h)),$$

where $p \log_2 p = 0$ if $p = 0$.

Because we want the value function to increase with preference, we let an entropy-based value function be

$$V(P(H)) = -\,\mathrm{Ent}(P(H)) = \sum_{h \in H} P(h) \log_2(P(h)).$$

Variance

If the states of H are numeric, another classical measure can be used, namely the variance. Again, since small variances are preferred, the value function becomes

$$V(P(H)) = -\sum_{h \in H} (h - \mu)^2 P(h),$$

where $\mu = \sum_{h \in H} hP(h)$.

It is up to the modeler to specify the value function. If decisions with known utilities are attached to the hypothesis variable, then the utility value function should be preferred. If this is not the case, the user will mainly be interested in the precision of a diagnosis.

In the case of a Boolean hypothesis with states 0 and 1, the entropy function is $\log p^p (1-p)^{1-p}$, and the variance function is $-p(1-p)$. These two functions reflect that the value of p increases as it approaches its bounds 0 and 1. The entropy function is rather drastic in the way that the slope is infinite for 0 and 1. Therefore, small changes of p close to 0 and 1 will be highly valued. On the other hand, the variance is of polynomial degree 2, and the slope close to the bounds is 1 and -1, giving changes almost even value no matter how close they are to the bound.

Other Value Functions

In principle, any value function may be used. However, a particular class of functions called *convex functions* are best suited for the purpose.

Definition 11.1. *A function $f : R^n \rightarrow R$ is* convex *if for any two points P_1, P_2 on the graph of f, the line segment P_1P_2 lies above the graph (see Figure 11.5). Mathematically, the property is expressed as follows:*

$$\forall t \in [0,1], \forall \mathbf{x}, \mathbf{y} \in R^n : tf(\mathbf{x}) + (1-t)f(\mathbf{y}) \geq f(t\mathbf{x} + (1-t)\mathbf{y}).$$

The reason why a convex function is well suited is due to the following theorem, which we will not prove.

Theorem 11.1. *If the value function is a convex function, then the expected benefit of performing a test is never negative.*

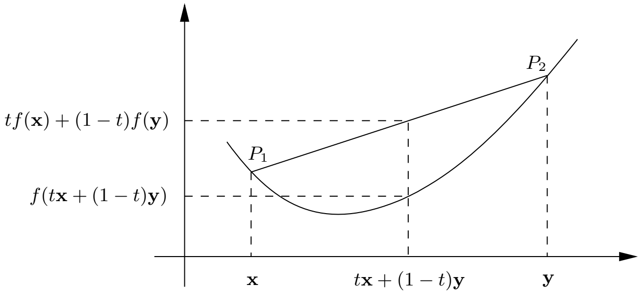

Fig. 11.5. A convex function. The line segment between two points of the graph lies above the graph.

Utility based functions are convex and so are entropy and variance.

11.2 Finding the Relevant Past and Future of a Decision Problem

When solving a decision problem we look for an optimal policy for each of the decisions. The optimal policy for a decision is in principle a function that for each possible configuration of the past, prescribes how to act in order to maximize the expected utility. Thus, for the poker domain modeled in Figure 11.6, a policy for the decision node D is a function over the entire past of D:

$$\delta_D : \mathrm{sp}(\mathit{MH0}, \mathit{MFC}, \mathit{MH1}, \mathit{OFC}, \mathit{MSC}, \mathit{OSC}) \rightarrow \mathrm{sp}(D)\,.$$

In general, if we represent such a policy function as a table, then the size of the policy increases exponentially in the number of variables in the past, and the policy can therefore quickly become intractable to handle.

However, when analyzing the decision problem above, we find that not all variables can provide information influencing decision D. For example, if I know my current hand *MH2*, then knowledge about how many cards I discarded in the second round, *MSC*, will not affect my decision at D: at D I will try to maximize my profit represented by the utility function U. This utility function depends only on D and *BH*, and with knowledge of the state of *MH2*, the decision *MSC* becomes d-separated from *BH*. Hence *MSC* cannot tell me anything about *BH* (and therefore U), and it can therefore not affect my decision at D. By performing this type of analysis for the remainder of the variables in the past of D, we find that the only variables that can have

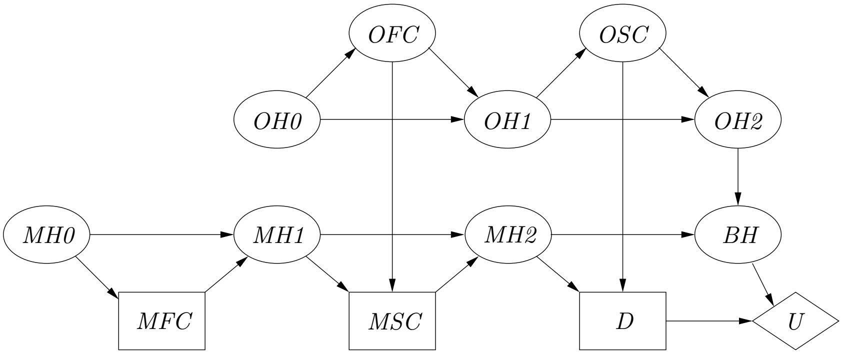

Fig. 11.6. An influence diagram representation of the poker domain described in Section 9.4.1. An optimal policy for decision D is a function over the past of D, namely *MH0*, *MFC*, *MH1*, *OFC*, *MSC*, and *OSC*.

an impact on D are *OFC*, *OSC*, and *MH2*. Hence, the optimal policy for D reduces to

$$\delta_D : \mathrm{sp}(\mathit{OFC}, \mathit{OSC}, \mathit{MH2}) \rightarrow \mathrm{sp}(D)\,.$$

This policy contains only 96 configurations, as opposed to the full policy function containing 165888 configurations. By doing the same exercise for the two remaining decisions we find that only *MH1* and *OFC* are relevant for *MSC*, and *MH0* is relevant for *MFC*.

Definition 11.2 (Required variables). *Let I be an influence diagram and let D be a decision variable in I. The variable $X \in \mathrm{past}(D)$ is said to be* required *for D if there exist a realization $\mathcal{R}$ of I, a configuration $\bar{y}$ over $\mathrm{dom}\,(\delta_D) \setminus \{X\}$, and states x_1 and x_2 of X such that $\delta_D(x_1, \bar{y}) \neq \delta_D(x_2, \bar{y})$, where δ_D is an optimal policy for D with respect to $\mathcal{R}$. The set of variables required for D is denoted by $\mathrm{req}(D)$.*

To take another example, consider the influence diagram in Figure 11.7, which specifies the partial ordering

$$\{B\} \prec D_1 \prec \{E, F\} \prec D_2 \prec D_3 \prec \{G\} \prec D_4 \prec \mathcal{C}_4;$$

$\mathcal{C}_4$ denotes the variables not observed before the last decision.

When looking for an optimal policy for D_4 we should in principle consider all the variables in the past of D_4, i.e., B, D_1, E, F, D_2, D_3, and G. However, when analyzing the influence diagram, we see that deciding on D_4 has an impact only on V_4, and from the d-separation properties of the model we have that by conditioning on G and D_2, all the other variables in the past of D_4 become d-separated from V_4. Hence, only G and D_2 are required for D_4.

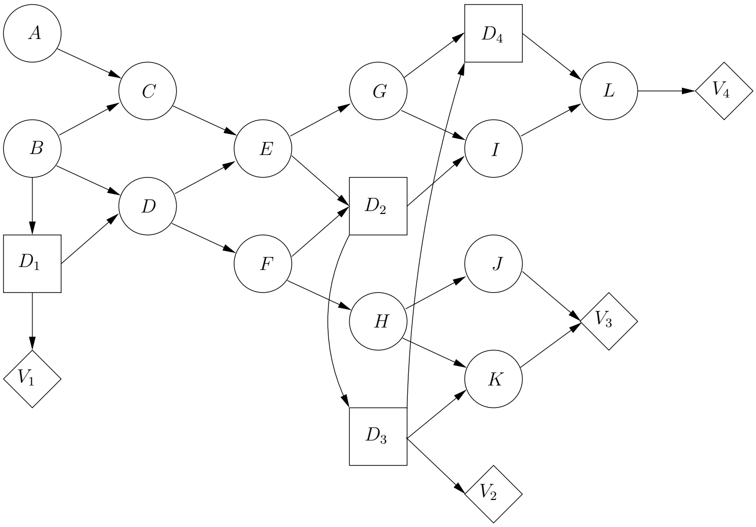

Fig. 11.7. The figure illustrates an influence diagram that specifies the partial order $\{B\} \prec D_1 \prec \{E, F\} \prec D_2 \prec D_3 \prec \{G\} \prec D_4 \prec \mathcal{C}_4$ ($\mathcal{C}_4$ denotes the chance variables observed after deciding on all the decisions.

11.2.1 Identifying the Required Past

In the examples above we informally characterized a variable as being required for D if it can provide information about the utility functions that we are trying to maximize when deciding on D. To test whether a variable X can provide information about these utility functions, we used the d-separation criterion. The question is then how to identify the utility functions that can influence D. To be on the safe side you might simply include all utility functions, but this may result in variables that are falsely identified as required for D. So we would like to identify the minimal set of utility functions to take into account when deciding on a particular decision.

Definition 11.3 (Relevant utility nodes). *The utility function U is* relevant *for decision D if there exists two realizations $\mathcal{R}_1$ and $\mathcal{R}_2$ of I that differ only on U such that the optimal policies for D are different in $\mathcal{R}_1$ and $\mathcal{R}_2$.*

Luckily, it turns out that this semantic definition also supports a simple syntactic characterization. For the last decision we have the following specification:

Proposition 11.1. *Let D_n be the last decision variable in the influence diagram I, and let U be a utility node in I. Then U is relevant for D_n if and only if there is a directed path from D_n to U.*

Proof. For the last decision D_n we know that the optimal policy is

$$\begin{aligned}\delta_{D_n}(\text{past}(D_n)) &= \arg\max_{D_n} \sum_{\mathcal{C}_n} P(\mathcal{C}_n \mid \text{past}(D_n), D_n)\Big[U(\text{pa}(U)) \\ &\qquad\qquad + \sum_{i=1}^{m} U_i(\text{pa}(U_i))\Big] \\ &= \arg\max_{D_n}\Big[\sum_{\mathcal{C}_n} P(\mathcal{C}_n \mid \text{past}(D_n), D_n) U(\text{pa}(U)) \\ &\qquad + \sum_{\mathcal{C}_n} P(\mathcal{C}_n \mid \text{past}(D_n), D_n) \sum_{i=1}^{m} U_i(\text{pa}(U_i))\Big].\end{aligned}$$

Since

$$\begin{aligned}&\sum_{\mathcal{C}_n} P(\mathcal{C}_n \mid \text{past}(D_n), D_n) U(\text{pa}(U)) \\ &\qquad = \sum_{\mathcal{C}_n \cap \text{pa}(U)} P(\mathcal{C}_n \cap \text{pa}(U) \mid \text{past}(D_n), D_n) U(\text{pa}(U)),\end{aligned}$$

we have that U is relevant for D_n if and only if D_n is either a parent of U or D_n is d-connected to a variable in $\mathcal{C}_n \cap \text{pa}(U)$ given $\text{past}(D_n)$; otherwise, the above expression would be independent of D_n. In order for D_n to be d-connected to a variable $X \in \mathcal{C}_n \cap \text{pa}(U)$ given $\text{past}(D_n)$, there must be an active path between $\text{pa}(U)$ and D_n. Since a converging connection on such a path cannot be opened by evidence (a descendant of D_n cannot be observed), the path must be directed from D_n to a node in $\text{pa}(U)$. □

Based on this proposition, we now have a full syntactic characterization of the variables required for the last decision.

Proposition 11.2. *Let D be the last decision variable in the influence diagram I and let X be a variable in* $\text{past}(D)$*. Then X is required for D if and only if X is d-connected to a utility node relevant for D given* $\text{past}(D) \setminus \{X\}$.

Proof. Follows the proof above. □

For example, if we go back to the influence diagram shown in Figure 11.7, we see that V_4 is the only utility node to which there exists a directed path from D_4; hence V_4 is the only utility node relevant for D_4. Moreover, using Proposition 11.2 we find that only G and D_2 are required for D_4, $\text{req}(D_4) = \{G, D_2\}$.

Suppose now that we also want to identify the required variables for D_3. This can be done by substituting D_4 with its chance-variable representation (actually, we need not calculate the policy). This is done in Figure 11.8.

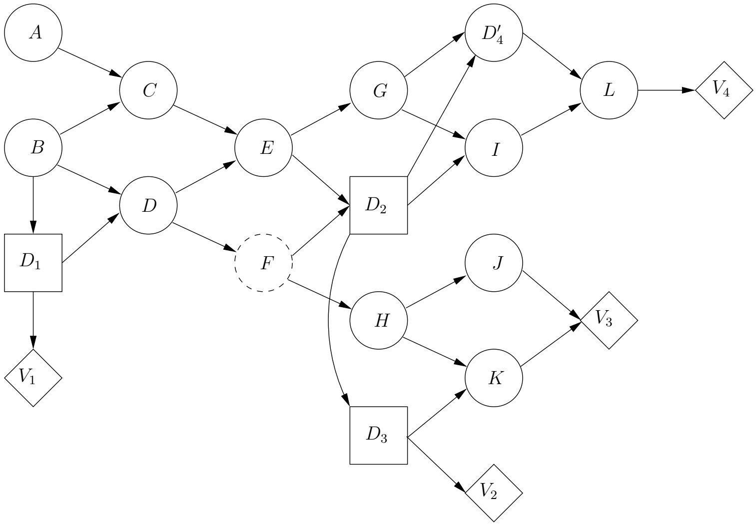

Fig. 11.8. The influence diagram obtained from the influence diagram in Figure 11.7 by substituting D_4 with its chance-variable representation. Since D_3 is the last decision, we see from Proposition 11.2 that F is the only variable required for D_3.

In this transformed influence diagram, D_3 appears as the last decision, and by applying the propositions we find that V_2 and V_3 are relevant for D_3 and that F is the only variable required for D_3.

By replacing D_3 in Figure 11.8 with its chance-variable representation we obtain the influence diagram in Figure 11.9, where D_2 is the last decision. From this model we find that V_4 is the only utility function relevant for D_2, and E is therefore the only variable required for D_2.

Finally, we can find the required variables for D_1 by substituting D_2 with its chance-variable representation . The resulting model is shown in Figure 11.10, where we see that all four utility functions are relevant for D_1, and since B is d-connected to V_2, V_3, V_4 we have that B is required for D_1.

More generally, we can specify an algorithm for finding the required variables for the decisions in an influence diagram as follows.

Algorithm 11.1 [Identify required variables] *Let I be an influence diagram and let $D_1, D_2, \ldots, D_n$ be the decision variables in I ordered by index. To determine* $\mathrm{req}(D_i)$*, the variables required for D_i ($\forall 1 \leq i \leq n$), do:*

1. *Set $i := n$.*
2. **For** *each decision variable D_i not considered ($i > 0$)*
   a) *Let $\mathcal{V}_i$ be the set of utility nodes to which there exists a directed path from D_i in I.*
   b) *Let* $\mathrm{req}(D_i)$ *be the set of nodes X such that $X \in \mathrm{past}(D_i)$ and X is d-connected to a node in $\mathcal{V}_i$ given* $\mathrm{past}(D_i) \setminus \{X\}$.

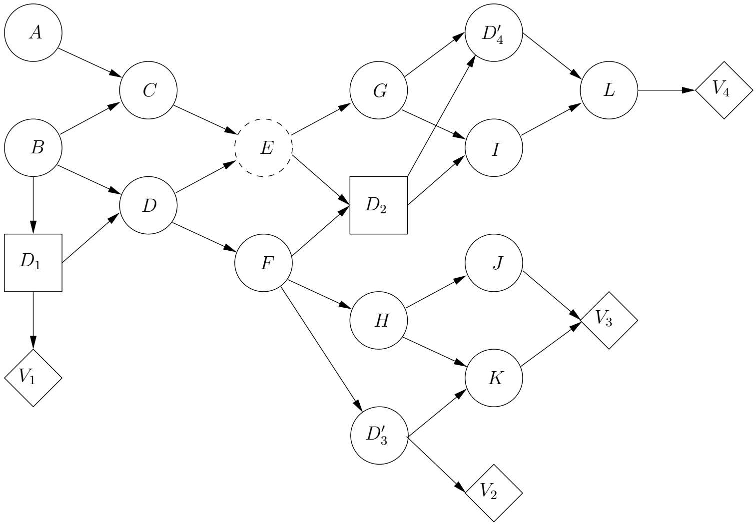

Fig. 11.9. The influence diagram obtained form the influence diagram in Figure 11.8 by substituting D_3 with its chance-variable representation.

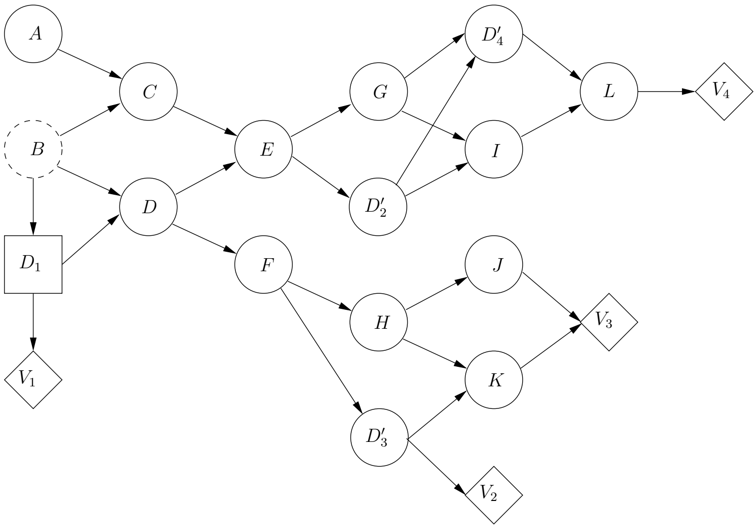

Fig. 11.10. The influence diagram obtained from Figure 11.9 by replacing D_2 with its chance variable policy.

c) *Replace D_i with a chance-variable representation of the policy for D_i, and let I be the resulting model.*
d) *Set $i := i - 1$.* □

Identifying the Relevant Future

Analogously to the idea of identifying the required variables in the past of a decision, we can also identify the future variables that are relevant for that decision. By relevant variables we mean the variables whose probability distributions (or policies) should be taken into account when deciding on D. Having such a characterization will not reduce the complexity of the policies, but it may provide insight into the overall structure of the decision problem. For example, if some decision variable is of particular interest, then the relevant variables may pinpoint the part of the model that we should focus on when specifying the probabilities.

Definition 11.4. *Let I be an ID and let D be a decision variable in I. The future variable X is said to be* relevant *for D if either:*

- *X is a chance node and there exist two realizations $\mathcal{R}_1$ and $\mathcal{R}_2$ of I that differ only on the probability distribution associated with X such that the optimal policies for D are different in $\mathcal{R}_1$ and $\mathcal{R}_2$, or*
- *X is a decision variable and there exist a realization of I and two different policies δ_X^1 and δ_X^2 for X such that the optimal policies for D are different with respect to δ_X^1 and δ_X^2.*

Together with the required past, the relevant variables describe the part of a decision problem that is sufficient to take into account when one is focusing on a particular decision.

To complete the characterization, we need an algorithm for identifying the variables that are relevant for a decision D. The first thing to notice is that by using the chance-variable representation of a decision node, we again need to consider only the situation in which D is the last decision variable in the influence diagram. Hence we can identify the relevant future decisions as the decision variables whose chance-variable representations are relevant for D. This also means that in order to identify all the relevant variables we just need a method for identifying the relevant chance variables.

Theorem 11.2. *Let I be an ID and let D be the last decision variable in I. Then the future chance variable X is relevant for D if and only if*

- *X is not barren in the ID formed from I by removing all utility nodes that are not relevant for D, and*[1]
- *there exists a utility node U relevant for D such that X is d-connected to U in I given $\{D\} \cup \mathrm{past}(D)$.*

[1] If X is barren, then it does not affect any decisions and it can simply be removed.

By going back to the influence diagram in Figure 11.7, we see that I and L are the only future variables d-connected to the relevant utility function, V_4, for D_4. Hence, no other future chance variables are relevant, and the decision problem for D_4 can therefore be described by the utility node V_4, the required variables G and D_2, and the relevant chance variables I and L, see Figure 11.11(a). To determine the relevant variables for D_3 we substitute D_4 with a chance variable and apply the same procedure as above. That is, from Figure 11.8 we see that H, I, and K are relevant for D_3, and together with the relevant utility nodes and the required variables we can identify the part of the decision problem relevant for D_3. See Figure 11.11(b). By continuing to D_2, we use the influence digram in Figure 11.9. When performing the analysis, we identify the variable D_4' as relevant for D_2, which in turn means that the decision node D_4 is relevant for D_2 (the identification of the remaining variables is left as an exercise).

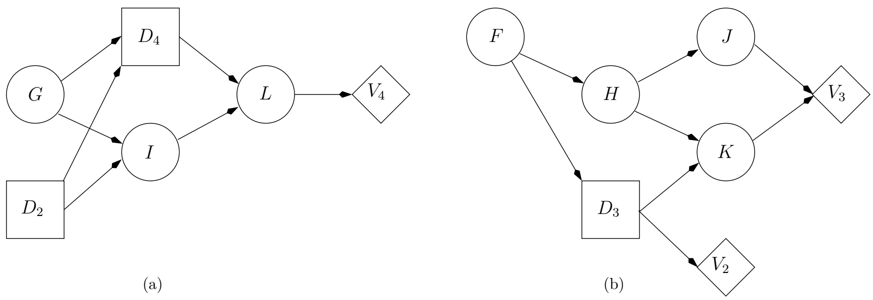

Fig. 11.11. The figures illustrate the parts of the influence diagram in Figure 11.7 relevant for D_4 and D_3, respectively.

11.3 Sensitivity Analysis

One of the main difficulties in modeling a decision problem is the elicitation of utilities and probabilities. This makes it desirable to be able to investigate how sensitive the solution is to variations in some utility or probability parameter, and how robust the solution is to joint variations over a set of parameters.

We distinguish between *value sensitivity* and *decision sensitivity*. Value sensitivity concerns variations in the maximum expected utility when a set of parameters is changed, and decision sensitivity refers to changes in the optimal strategy.

11.3.1 Example

Consider the following simplified binary version of the Oil Wildcatter Problem from Exercise 9.11. The influence diagram is shown in Figure 11.12. The hole can be *good* or *bad*. If the hole is good, the gain is \$260,000, and if the hole is bad, the gain is \$0. The test has no false negatives, and the probability of a false positive is 0.05. The prior probability for the hole being good is 0.2. The cost of drilling is \$60,000, and the cost of the test is \$5,000.

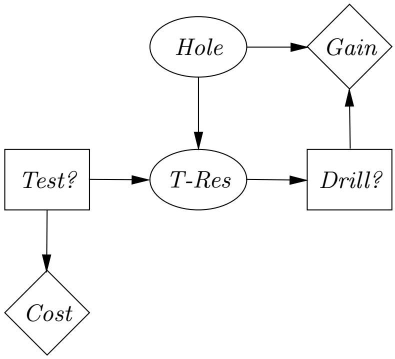

Fig. 11.12. An influence diagram for the Oil Wildcatter Problem.

The optimal strategy, Δ, is to test and then to drill if and only if the test is positive. However, although the oil wildcatter is quite certain of the specifics of the test, he is rather uncertain of the gain of a good hole as well as of the prior probability for this particular hole being good. If the gain and the prior for a good hole are large, he need not test, because he will drill regardless of the result of the test, and if the prior and the gain for a good hole are low, he will just leave the hole.

To be precise, the optimal strategy consists of two optimal policies, $\delta_{Test?} = y$ for *Test?*, and $\delta_{Drill?}(Test?, T')$ for *Drill?*, where $\delta_{Drill?}(y, pos) = y$, $\delta_{Drill?}(y, neg) = n$, $\delta_{Drill?}(n, no\text{-}test) = n$, and the values for other configurations are of no importance, since they will never be realized.

Let t denote $P(Hole = good)$ and let s denote $Gain(Hole = good) - 60000$. Then $\delta_{Drill?}$ is optimal for $(t, s) = (0.2, 200000)$, and the wildcatter would like to know which parameter values support this policy. To determine the support, we calculate the expected utilities of the various options. The relevant utilities are only the utilities on which *Drill?* has an impact, namely *Gain*; the descendant of *Drill?*. We now get

$$\begin{aligned}
\mathrm{EU}(\mathit{Drill?} \,|\, n, \mathit{no\text{-}test}) &= (P(\mathit{good} \,|\, \mathit{no\text{-}test})s - P(\mathit{bad} \,|\, \mathit{no\text{-}test})60000, 0) \\
&= (ts - (1-t)6000, 0), \\
\mathrm{EU}(\mathit{Drill?} \,|\, y, \mathit{pos}) &= (P(\mathit{good} \,|\, \mathit{pos})s - P(\mathit{bad} \,|\, \mathit{pos})60000, 0) \\
&= \left(\frac{ts - 0.05(1-t)60000}{0.95t + 0.05}, 0\right), \\
\mathrm{EU}(\mathit{Drill?} \,|\, y, \mathit{neg}) &= (P(\mathit{good} \,|\, \mathit{neg})s - P(\mathit{bad} \,|\, \mathit{neg})60000, 0) \\
&= (-60000, 0).
\end{aligned}$$

The policy $\delta_{\mathit{Drill?}}$ is optimal if

$$\begin{aligned}
\mathrm{EU}(\mathit{Drill?} = n \,|\, n, \mathit{no\text{-}test}) &\geq \mathrm{EU}(\mathit{Drill?} = y \,|\, n, \mathit{no\text{-}test}), \\
\mathrm{EU}(\mathit{Drill?} = y \,|\, y, \mathit{pos}) &\geq \mathrm{EU}(\mathit{Drill?} = n \,|\, y, \mathit{pos}), \\
\mathrm{EU}(\mathit{Drill?} = n \,|\, y, \mathit{neg}) &\geq \mathrm{EU}(\mathit{Drill?} = y \,|\, y, \mathit{neg}).
\end{aligned}$$

This gives the following inequalities:

$$\begin{aligned}
&0 \geq ts - (1-t)60000, \\
&0 \leq ts - 0.05(1-t)60000, \\
&0 \geq -6000.
\end{aligned}$$

That is,

$$ts + 3000t - 3000 \geq 0 \geq ts + 60000t - 60000. \tag{11.1}$$

For $s = 200000$ we get that $\delta_{\mathit{Drill?}}$ is optimal for $\frac{3}{203} \leq t \leq \frac{3}{13}$, and for $t = 0.2$ it is optimal for $12000 \leq s \leq 240000$. These intervals are called the *admissible domains* for the parameters in $\delta_{\mathit{Drill?}}$.

Next we analyze the first decision. The decision node *Drill?* is substituted with the chance node D (Figure 11.13), and $P(D \,|\, T')$ reflects the optimal policy (see Section 10.2.3).

Using the model in Figure 11.13 we calculate

$$\begin{aligned}
\mathrm{EU}(\mathit{Test?} = y) &= -5000 + P(\mathit{pos})(\mathrm{EU}(\mathit{Drill?} = y \,|\, \mathit{pos}) \\
&= -5000 + (0.95t + 0.05)\frac{ts - 0.05(1-t)60000}{0.95t + 0.05} \\
&= -5000 + ts - 0.05(1-t)60000 \\
&= ts + 3000t - 8000, \\
\mathrm{EU}(\mathit{Test?} = n) &= 0.
\end{aligned}$$

This yields that testing is optimal if

$$ts + 3000t - 8000 \geq 0. \tag{11.2}$$

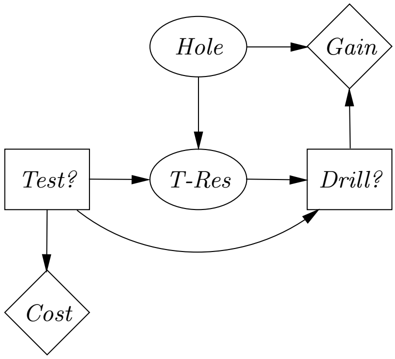

Fig. 11.13. The decision node *Drill?* is substituted by its chance-node representation.

For $s = 200000$ it holds for $t \geq \frac{8}{203}$ and for $t = 0.2$ it holds for $s \geq 37000$. The strategy is optimal in the intersection of the admissible domains of the two policies. That is, for $s = 200000$ the admissible domain for t is $[\frac{8}{203}, \frac{3}{13}]$. For $t = 0.2$, the admissible domain for s is $[37000, 240000]$.

11.3.2 One-Way Sensitivity Analysis in General

Let t be a parameter with initial value t_0 in an influence diagram, and let Δ be an optimal strategy for the value t_0. We wish to determine the admissible interval for t. The method starts determining the admissible interval for the policy $\delta_{Drill?}$ for the last decision D. Then D is substituted by its chance-variable representation, and the admissible interval for t is determined for the last decision in this influence diagram. The procedure is repeated until the first decision has been analyzed. The admissible interval for t in Δ is the intersection of the admissible intervals for all the policies. Since t_0 is a member of all intervals, we know that the intersection is nonempty.

In the example above it turned out that the expected utilities were simple expressions in the parameters. This holds in general.

Theorem 11.3. *Let s be a utility parameter in the influence diagram ID, let D be the last decision in ID, and let π be any configuration of the required past of D. Then for any d in D, the expected utility of d given π is a linear function in s.*

Let t be probability parameter in the influence diagram ID, let D be the last decision in ID, and let π be any configuration of the required past of D. Then for any d in D, the expected utility of d given π is a fraction of two linear functions in t.

Proof. [Sketch] The expected utility is calculated as

$$\sum_{Parents} P(\text{Parents} \,|\, \text{past}) U(\text{Parents}).$$

For utility parameters, this expression is linear. A probability parameter has an effect on $P(\text{Parents} \,|\, \text{past})$, and from Corollary 5.2, it can be expressed as a fraction of two linear functions. □

As for sensitivity analysis for Bayesian networks, this theorem can be exploited to establish a functional expression for the expected utilities. Assume that we analyze a utility parameter s with initial value s_0. We have a solution for ID with value s_0. That is, we have a value of the expected utility for the last decision D_n for each configuration of the required past. Next, substitute s_0 with s_1 and solve the influence digram. Again, we get the expected utility for each option and any configuration of the required past. Now, for each option and for each parent configuration we have two values of the expected utility, and the two coefficients in the linear expression can be determined.

The next step is to establish a new influence diagram, and do the same with D_{n-1} as the last decision. However, if the value s_1 lies in the admissible interval for the policy for D_n, the solution from before can be reused. The optimal policy for D_n is guaranteed, also for the value s_1, to be the same as the conditional probability for its chance-node representation. This holds for the next decisions too, so by careful choice of the new value, one extra solution of the influence diagram is sufficient for the calculation of all the expected utilities required for determining the admissible domain for the parameter. In the case of probability parameters, three extra solutions are sufficient.

We shall illustrate the method for the parameter s in the oil wildcatter example above.

Solving the influence diagram with $s = 200000$ we get the following expected utilities:

$$\begin{aligned} \mathrm{EU}(\mathit{Drill?} \,|\, \mathit{pos}) &= (156666, 0), \\ \mathrm{EU}(\mathit{Drill?} \,|\, \mathit{neg}) &= (-60000, 0), \\ \mathrm{EU}(\mathit{Drill?} \,|\, \mathit{no\text{-}test}) &= (-8000, 0), \\ \mathrm{EU}(\mathit{Test?}) &= (32600, 0). \end{aligned}$$

Changing s to 150000 we get

$$\begin{aligned} \mathrm{EU}(\mathit{Drill?} \,|\, \mathit{pos}) &= (115000, 0), \\ \mathrm{EU}(\mathit{Drill?} \,|\, \mathit{neg}) &= (-60000, 0), \\ \mathrm{EU}(\mathit{Drill?} \,|\, \mathit{no\text{-}test}) &= (-18000, 0), \\ \mathrm{EU}(\mathit{Test?}) &= (22600, 0). \end{aligned}$$

This yields the following expressions:

$$\mathrm{EU}(\mathit{Drill?} = y \mid \mathit{pos}) = 0.833s + 10000,$$
$$\mathrm{EU}(\mathit{Drill?} = y \mid \mathit{neg}) = -60000,$$
$$\mathrm{EU}(\mathit{Drill?} = y \mid \mathit{no\text{-}test}) = 0.2s - 48000,$$
$$\mathrm{EU}(\mathit{Test?} = y) = 0.2s - 7400,$$

which are the same as the result of the expressions in Section 11.3.1.

If you wish to find out how stable the strategy is to joint variations of several the parameters, one-way sensitivity analysis for each parameter may not provide the full picture and you may need to resort to n-way sensitivity analysis. However, the work becomes much harder. For example, in the case of a probability parameter t and a utility parameter s, the expected utilities have the form $\alpha s + \beta$, where α and β are fractions of linear expressions over t. This means that there are eight coefficients to determine. For illustration, the admissible area for (t, s) in the strategy from Section 11.3.1 is shown in Figure 11.14.

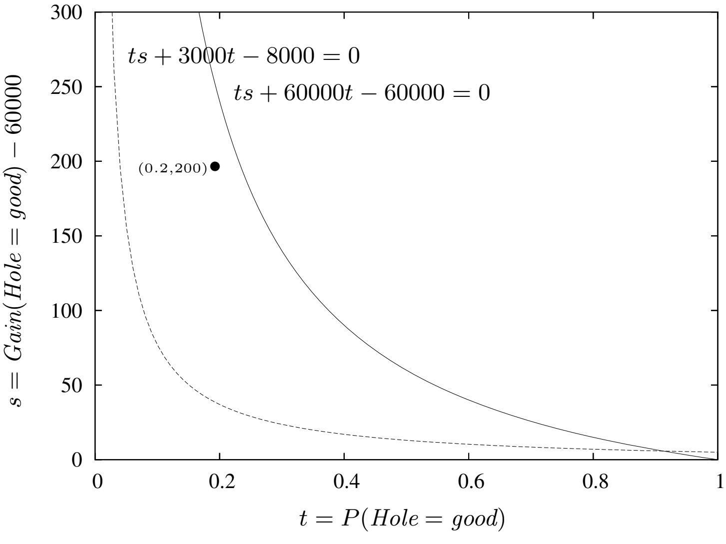

Fig. 11.14. The admissible area for (t, s) in the strategy for the the oil wildcatter. The y-axis is scaled by a factor of 1000.

If all parameters are utility parameters, $s_1, \ldots, s_n$, then the situation is much simpler. Since utilities are never multiplied, the expected utilities are linear expressions over $s_1, \ldots, s_n$. Therefore, there are only $n + 1$ coefficients to determine, and n extra solutions are sufficient.

11.4 Summary

Value of Information

Value function (one utility function U, one decision D):

$$V(P(H)) = \max_{d \in D} \sum_{h \in H} U(d,h) P(h \mid d).$$

Expected value of performing test T:

$$EV(T) = \sum_{t \in T} P(t) \max_{d \in D} \sum_{h \in H} U(d,h) P(h \mid t, d).$$

Expected profit:

$$EP(T) = EV(T) - V(P(H)) - C_T.$$

The value $EV(T)$ can be calculated for all tests T by entering the states of h as evidence and using Bayes' rule.

Myopic approach: Choose repeatedly a test with the highest positive expected profit, if any.

Nonutility value functions:

- Entropy: $V(P(H)) = \sum_{h \in H} P(h) \log_2(P(h))$;
- Variance: $V(P(H) = -\sum_{h \in H}(h-\mu)^2 P(h)$, where $\mu = \sum_{h \in H} hP(h)$.

The Required Past for a Decision

Let I be an influence diagram and let D be a decision variable in I. The variable $X \in \mathrm{past}(D)$ is said to be *required* for D if there exist a realization of I, a configuration $\bar{y}$ over $\mathrm{dom}\,(\delta_D) \setminus \{X\}$, and states x_1 and x_2 of X such that $\delta_D(x_1, \bar{y}) \neq \delta_D(x_2, \bar{y})$. The set of variables required for D is denoted by $\mathrm{req}(D)$.

To determine $\mathrm{req}(D_i)$ ($\forall 1 \leq i \leq n$) do:

1. Set $i := n$.
2. **For** each decision variable not considered ($i > 0$)
   a) Let $\mathcal{V}_i$ be the set of value nodes to which there exists a directed path from D_i in I.
   b) Let $\mathrm{req}(D_i)$ be the set of nodes X such that $X \in \mathrm{past}(D_i)$ and X is d-connected to a node in $\mathcal{V}_i$ given $\mathrm{past}(D_i) \setminus \{X\}$.
   c) Replace D_i with a chance-variable representation of the policy for D_i, and let I be the resulting model.
   d) Set $i := i - 1$.

Sensitivity Analysis

Value sensitivity: How much can the utility and probability parameters be varied without changing the optimal strategy? This question can be answered by performing an analysis of the expected utility as a function of these parameters.

Utility parameters: Let s be a utility parameter in the influence diagram ID, let D be the last decision in ID, and let π be any configuration of the required past of D. Then for any d in D, the expected utility of d given π is a linear function in s.

Probability parameters: Let t be probability parameter in the influence diagram ID, let D be the last decision in ID, and let π be any configuration of the required past of D. Then for any d in D, the expected utility of d given π is a fraction of two linear functions in t.

Calculating the coefficients: If there are only utility parameters to investigate, then all coefficients can be found by performing only one extra propagation for each parameter. This will also give all the information necessary for performing n-way sensitivity analysis (that is, sensitivity analysis in which you consider joint variations of the parameters).

11.5 Bibliographical Notes

Value of information is formally treated in (Howard, 1966) and (Lindley, 1971), where utilities are guiding the test selection. The myopic approximation was introduced by Gorry and Barnett (1968). In (Ben-Bassat, 1978), entropy and variance are used. Value of information for influence diagrams has been treated by Dittmer and Jensen (1997) and Shachter (1999). The required past of decisions in influence diagrams was introduced independently by Shachter (1999) and Nielsen and Jensen (1999). The relevant future of decisions was described in (Nielsen, 2002). Sensitivity analysis for multiple parameters in decision problems was investigated in (Felli and Hazen, 1999a). A method using value of information was given in (Felli and Hazen, 1999b). Sensitivity analysis for influence diagrams in particular was treated in (Nielsen and Jensen, 2003b).

The oil wildcatter's problem is due to Raiffa (1968). The used car buyer's problem is due to Howard (1962).

11.6 Exercises

Exercise 11.1. ^{E} Consider the insemination model from Exercise 3.8. Assume that you have the options to repeat the insemination or to wait another six-week period. The cost of repeating the insemination is 65 regardless of the pregnancy state of the cow. If the cow is pregnant and you wait, it will cost you nothing, but if the cow is not pregnant and you wait, it will cost you further 30 (that makes a total of 95 for waiting plus the eventual repeated insemination). The cost of *BT* is 1 and the cost of *UT* is 2. Perform a myopic value of information analysis.

Exercise 11.2. Solve the problem in Exercise 9.11 as a value of information problem.

Exercise 11.3. ^{E} Consider the influence diagram obtained by adding arcs from *FC*, *SC*, and *MH* to *D* in the network in Figure 9.3, using the probabilities found in Section 3.2.3 and the utilities found in Section 9.1.1. Assume that prior to the game, a shady-looking person at the table next to me offers to tell me the first hand of my opponent (*OH0*) for the price of \$0.1. Ignoring ethical issues, should I take the offer?

Exercise 11.4. Consider the influence diagrams in Figures 9.23 and 9.24. What is the required past of decision FV_4 in the two diagrams?

Exercise 11.5. What are the relevant futures of decisions D_1 and D_2 in the influence diagram in Figure 11.7?

Exercise 11.6. Consider again the influence diagram in Example 10.10 and the strategy Δ, stating that one should always choose d_1 and d_2 if and only if C_1 is in state c_1. Denoting by t the utility parameter $U(d_1, c_2, \neg d_2)$, what is the support of Δ?