DECISION ANALYSIS
Vol. 2, No. 3, September 2005, pp. 127–143
ISSN 1545-8490 | EISSN 1545-8504 | 05 | 0203 | 0127

DOI 10.1287/deca.1050.0020

Influence Diagrams

Ronald A. Howard and James E. Matheson
Strategic Decisions Group

Key words: influence diagram; expansion; expansion order; decision tree; decision-tree order; Bayes; arrow reversal; decision network; decision-tree network; value of clairvoyance; Bayesian network; belief network; knowledge map

1. Introduction

The rapid growth of electronic computation continues to challenge our ability to conceptualize and describe the world around us. Mathematical tools and formal descriptions serve poorly as a communication device with the majority of people not trained in nor used to mathematical means of expression. Yet virtually everyone has information useful in the solution of his own problems or the problems of others, if only it could be tapped.

The subject of this paper is a new form of description, the influence diagram, that is at once both a formal description of the problem that can be treated by computers and a representation easily understood by people in all walks of life and degrees of technical proficiency. It thus forms a bridge between qualitative description and quantitative specification.

The reason for the power of this representation is that it can serve at the three levels of specification of relation, function, and number, and in both deterministic and probabilistic cases. In the deterministic case, relation means that one variable can depend in a general way on several others; for example, profit is a function of revenue and cost. At the level of function, we specify the relationship; namely, that profit equals revenue minus cost. Finally, at the level of number, we specify the numerical values of revenue and cost and hence determine the numerical value of profit.

In the probabilistic case, at the level of relation we mean that given the information available, one variable is probabilistically dependent on certain other variables and probabilistically independent of still other variables. At the level of function, the probability distribution of each variable is assigned conditioned on values of the variables on which it depends. Finally, at the level of number, unconditional distributions are assigned on all variables that do not depend on any other variable and hence determine all joint and marginal probability distributions.

As an example of the probabilistic case, we might assert at the level of relation that income depends on age and education and that education depends on age. Next, at the level of function we would assign the conditional distribution of income given age and education and the distribution of education given age. Finally, at the level of number, we would assign the unconditional distribution on age.

The successive degrees of specification can be made by different individuals. Thus, an executive may know that sales depend in some way on price, but he may leave to others the probabilistic description of the relationship.

Because of its generality, the influence diagram is an important tool not only for decision analysis, but for any formal description of relationship and thus for all modeling work.

In the present paper, we shall focus on the probabilistic use of influence diagrams since the deterministic use is a special, but important, case of the probabilistic. We now proceed to development of the influence diagram concept, to examination of its properties, and to illustration of its use.

2. Probabilistic Dependence[1]

One of the most perplexing aspects of making decisions under uncertainty is the problem of representing

[1] This entire discussion applies as well to events as to variables.

and encoding probabilistic dependencies. A probabilistic dependency is one that arises as a result of uncertainty. For example, if a and b are known variables and $c = a + b$, then it is clear that c depends on both a and b, both in a vernacular sense and in a mathematical sense. However, suppose a is known and b is uncertain. Then c is probabilistically dependent on b but not on a. The reason is that knowing the specific value of b tells us something new about c, but there is no such possibility with respect to a.

3. Probabilistic Independence

Probabilistic independence, like the assigning of probability itself, depends on the state of information possessed by the assessor. Let x, y, and z be aleatory state variables of interest, which can be either continuous or discrete. Then $\{x \mid S\}$ is the probability distribution assigned to x given the state of information S. Two variables x and y are probabilistically independent given the state of information S if

$$\{x, y \mid S\} = \{x \mid S\}\{y \mid S\}$$

or equivalently, if

$$\{x \mid y, S\} = \{x \mid S\}.$$

4. Expansion

Regardless of whether x and y are probabilistically independent, we can write

$$\begin{aligned}\{x, y \mid S\} &= \{x \mid y, S\}\{y \mid S\}\\ &= \{y \mid x, S\}\{x \mid S\}.\end{aligned}$$

We call this the "chain rule of probabilities." Note that for three events there are six possible representations:

$$\begin{aligned}\{x, y, z \mid S\} &= \{x \mid y, z, S\}\{y \mid z, S\}\{z \mid S\}\\ &= \{x \mid y, z, S\}\{z \mid y, S\}\{y \mid S\}\\ &= \{y \mid x, z, S\}\{x \mid z, S\}\{z \mid S\}\\ &= \{y \mid x, z, S\}\{z \mid x, S\}\{x \mid S\}\\ &= \{z \mid x, y, S\}\{x \mid y, S\}\{y \mid S\}\\ &= \{z \mid x, y, S\}\{y \mid x, S\}\{x \mid S\}.\end{aligned}$$

For n variables there are $n!$ possible expansions, each requiring the assignment of a different set of probabilities and each logically equivalent to the rest. However, while the assessments are logically equivalent, there may be considerable differences in the ease with which the decision maker can provide them. Thus, the question of which expansion to use in a problem is far from trivial.

5. Probability Trees

Associated with each expansion is a probability tree. The expansion

$$\{x, y, z \mid S\} = \{x \mid y, z, S\}\{y \mid z, S\}\{z \mid S\}$$

implies the tree shown in Figure 5.1. The tree is a succession of nodes with branches emanating from each node to represent different possible values of a variable. The first assignment made is the probability of various values of z. The probability of each value of y is assigned conditioned on a particular value of z, and placed on the portion of the tree indicated by that value. Finally, the probabilities of various levels of x are assessed given particular values of z and y and placed on the portion of the tree specified by those values. When this has been done for all possible values of x, y, and z the tree is complete. The probability of any particular path through the tree is obtained by multiplying the values along the branches and is $\{x, y, z \mid S\}$. Notice that the tree convention uses small circles to represent chance nodes. If we wish to focus

Figure 5.1 A Probability Tree

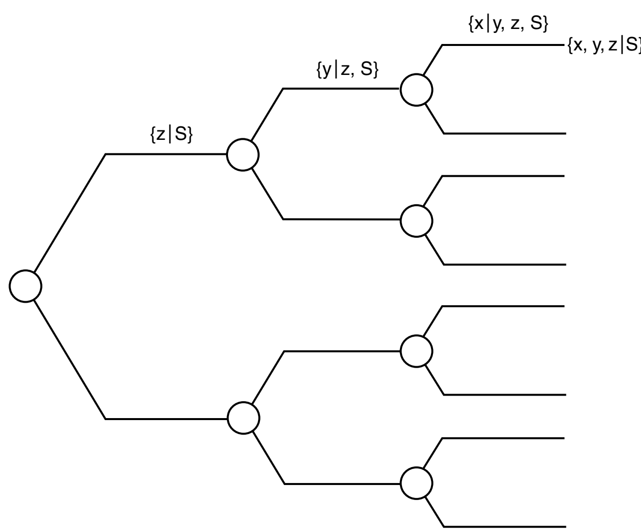

Figure 5.2 **A Generic Probability Tree**

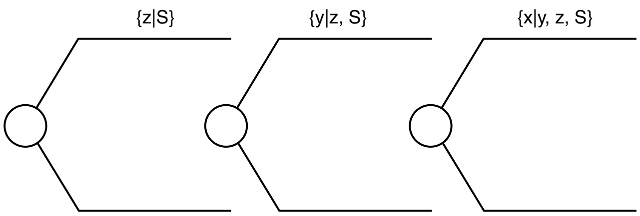

on the succession in the tree rather than the detailed connections, we can draw the tree in the generic form shown in Figure 5.2.

6. Decision Trees

If a variable is controlled by a decision maker, it is represented in a tree by a decision node. Thus, if y were a decision variable, Figure 5.2 could be redrawn as Figure 6.1. This tree states that the decision maker is initially uncertain about z and has assigned a probability distribution $\{z \mid S\}$ to it. However, he will know z at the time he must set y, the decision variable. This node is represented, like all decision nodes, by a small square box. Once z and y are given, the decision maker will still be uncertain about x; he has represented this uncertainty by $\{x \mid y, z, S\}$, Notice that a decision tree implies both a particular expansion of the probability assessments and a statement of the information available when a decision is made.

7. Probability Assignment for Decision Trees

The major problem with decision trees arises from the first of these characteristics. The order of expansion required by the decision tree is rarely the natural order in which to assess the decision maker's information. The decision-tree order is the simplest form for assessment only when each variable is probabilistically dependent on all preceding aleatory and

Figure 6.1 **A Generic Decision Tree**

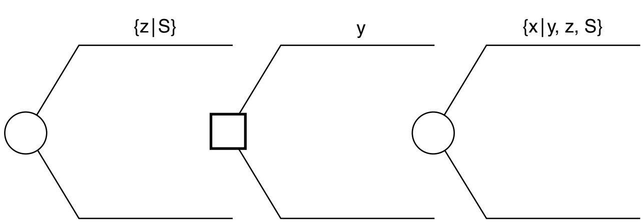

Figure 7.1 **A Four-Node Decision Tree**

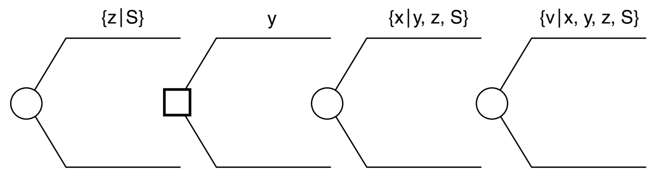

decision variables. If, as is usually the case, many independence assertions can be made, assessments are best done in a different order from that used in the decision tree. This means that we first draw a probability tree in an expansion form convenient to the decision maker and have him use this tree for assignment; it is called a probability assignment tree. Later the information is processed into the form required by the decision tree by representing it in one of the alternative expansion orders (Howard 1965). This is often called "using Bayes' Rule" or "flipping the tree." It is a fundamental operation permitted by the arbitrariness in the expansion order.

Consider, for example, the decision tree of Figure 6.1 with one additional aleatory variable v added, as shown in Figure 7.1. We interpret z as a test result that will become known, y as our decision, x as the outcome variable to which the test is relevant, and v as the value we shall receive if the test indicates z, we decide y, and x is the value of the outcome variable. Often y will not affect x in any way, even though y affects v; we write

$$\{x \mid y, z, S\} = \{x \mid z, S\}$$

to represent this assertion.

With this independence assertion we have the tree shown in Figure 7.2. This tree requires the specification of $\{z \mid S\}$ and $\{x \mid z, S\}$: the probability of various test results and the probability of various outcomes given test results. But typically in situations of this kind, the decision maker would prefer to assign

Figure 7.2 **A Four-Node Decision Tree Given the Assertion that y Will Not Affect x**

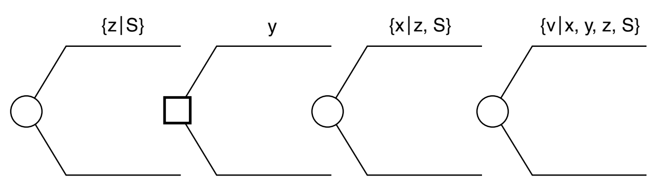

Figure 7.3 The Probability Assignment Tree

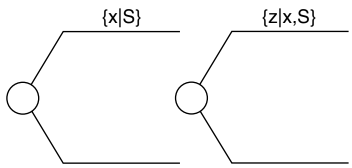

directly the probabilities of different outcomes $\{x \mid S\}$ and then the probabilities of various test results given the outcome, $\{z \mid x, S\}$. In other words, he would prefer to make his assessments in the probability tree of Figure 7.3 and then have them processed to fit the decision tree of Figure 7.2. Because

$$\{x \mid S\}\{z \mid xS\} = \{z \mid S\}\{x \mid zS\} = \{x, z \mid S\},$$

this is no more than choosing one expansion over the other. The exact processing required for the decision tree is then summation,

$$\{z \mid S\} = \int_x \{z \mid x, S\}\{x \mid S\},$$

and division,

$$\{x \mid z, S\} = \frac{\{z \mid x, S\}\{x \mid S\}}{\{z \mid S\}}.$$

Recall, however, that this whole procedure was possible only because variable x did not depend on the decision variable y.

8. Influence Diagrams

An influence diagram is a way of describing the dependencies among aleatory variables and decisions. An influence diagram can be used to visualize the probabilistic dependencies in a decision analysis and to specify the states of information for which independencies can be assumed to exist.

Figure 8.1 shows how influence diagrams represent the dependencies among aleatory variables and decisions. An aleatory variable is represented by a circle containing its name or number. An arrow pointing from aleatory variable A to aleatory variable B means that the outcome of A can influence the probabilities associated with B. An arrow pointing to a decision from either another decision or an aleatory variable means that the decision is made with the knowledge of the outcome of the other decision or aleatory variable. A connected set of squares and circles is called an influence diagram because it shows how aleatory variables and decisions influence each other.

Figure 8.1 Definitions Used in Influence Diagrams

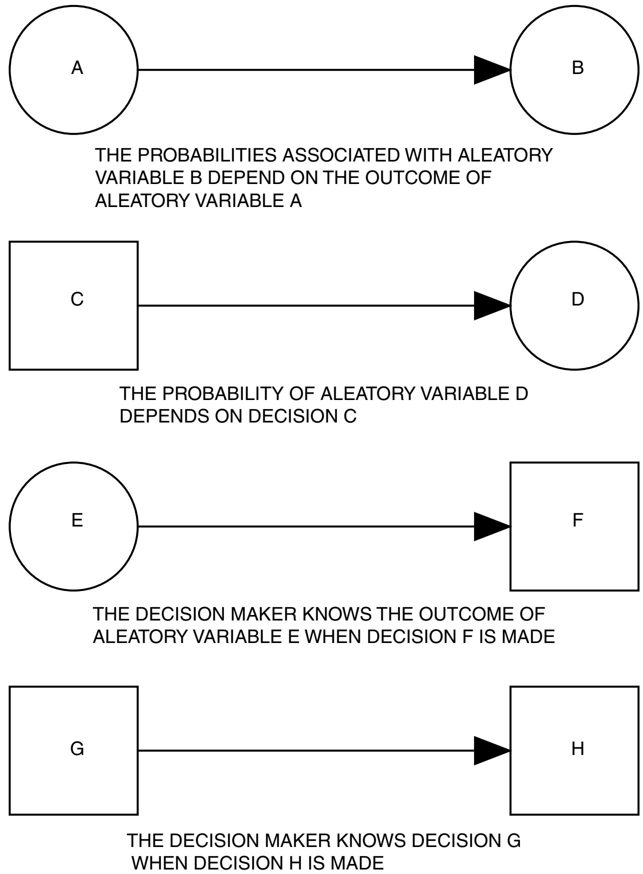

The influence diagram in Figure 8.2(a) states that the probability distribution assigned to x may depend on the value of y, whereas the influence diagram in Figure 8.2(b) asserts that x and y are probabilistically independent for the state of information with which the diagram was drawn. Note that the diagram of Figure 8.2(a) really makes no assertion about the probabilistic relationship of x and y since, as we know, any joint probability $\{x, y \mid S\}$ can be represented in the form

$$\{x, y \mid S\} = \{x \mid y, S\}\{y \mid S\}.$$

However, because

$$\{x, y \mid S\} = \{y \mid x, S\}\{x \mid S\},$$

the influence diagram of Figure 8.2(a) can be redrawn as shown in Figure 8.2(c); both are completely general

Figure 8.2 Two-Node Influence Diagrams

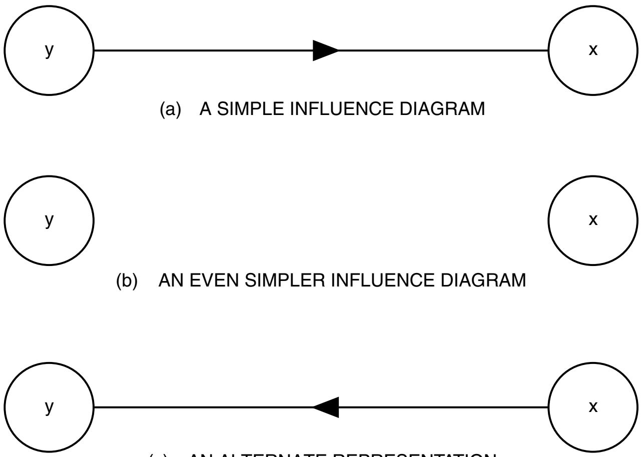

(a) A SIMPLE INFLUENCE DIAGRAM

(b) AN EVEN SIMPLER INFLUENCE DIAGRAM

(c) AN ALTERNATE REPRESENTATION

representations requiring no independence assertions. While the direction of the arrow is irrelevant for this simple example, it is used in more complicated problems to specify the states of information upon which independence assertions are made.

Figure 8.3 Alternate Influence Diagrams for $\{x, y, z \mid S\}$

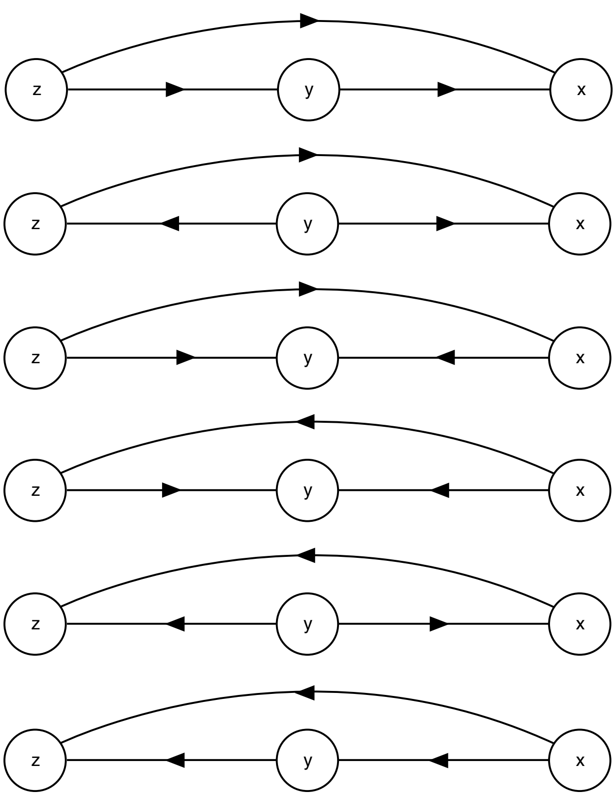

Similarly, with three variables x, y, z there are six possible influence diagrams of complete generality, one corresponding to each of the possible expansions we developed earlier. They are shown in Figure 8.3. While all of these representations are logically equivalent, they again differ in their suitability for assessment purposes. In large decision problems, the influence diagrams can display the needed assessments in a very useful way.

9. Graphical Manipulation

Because there are many alternative representations of an influence diagram, we might ask what manipulations can be performed on an influence diagram to change it into another form that is logically equivalent.

The first observation we should make is that an arrow can always be added between two nodes without making an additional assertion about the independence of the two corresponding variables (as long as no loops are created). That is, saying that x *may* depend on variable y is not equivalent to saying that x

Figure 9.1 Graphical Manipulation of Influence Diagrams

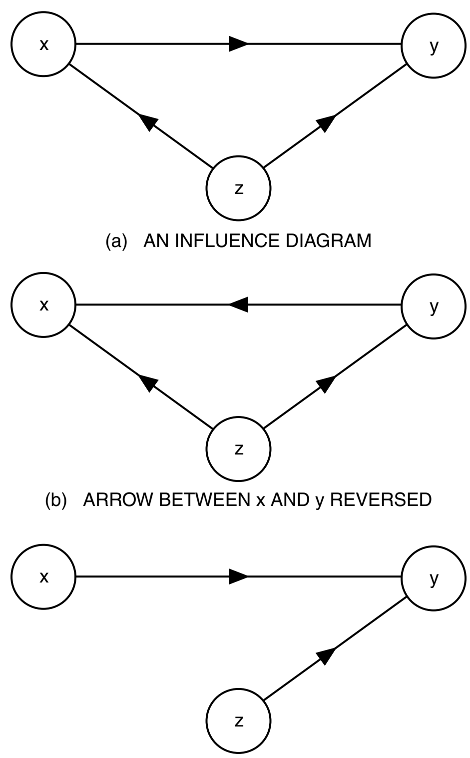

(a) AN INFLUENCE DIAGRAM

(b) ARROW BETWEEN x AND y REVERSED

(c) ANOTHER INFLUENCE DIAGRAM

Figure 9.2 Graphical Manipulations Producing Non-Unique Results

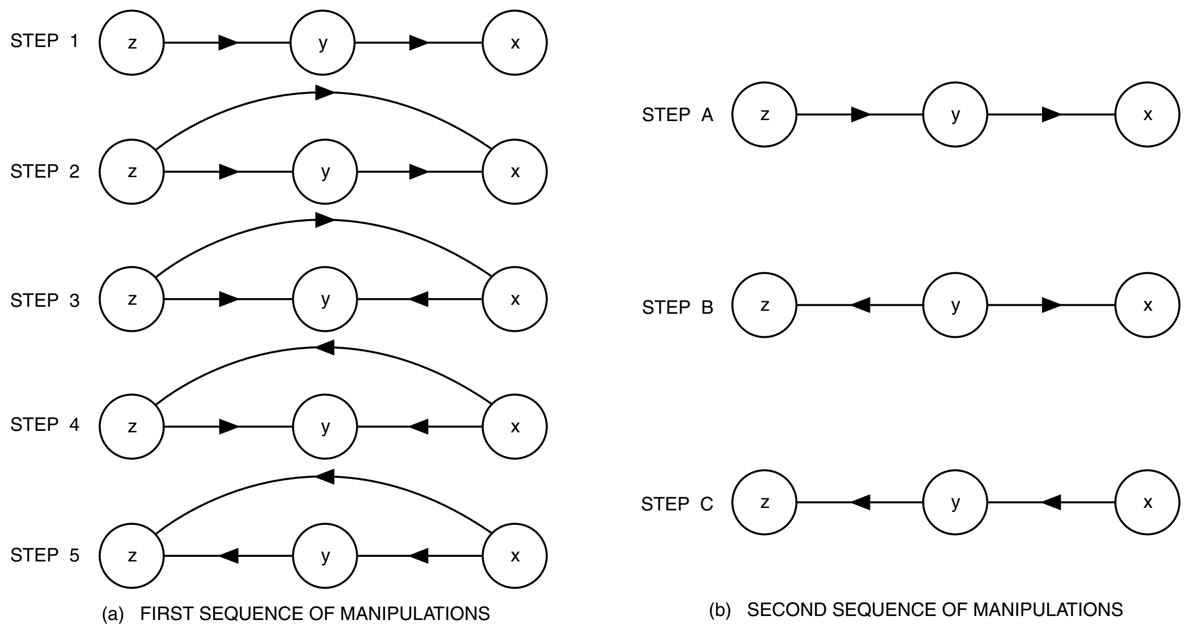

must depend on y. Thus, the diagram of Figure 8.2(b) can be changed into either of the diagrams shown in Figures 8.2(a) and 8.2(c) without making an erroneous assertion. However, the reverse procedure *could* lead to an erroneous assertion. Creating additional influence arrows will not change any probability assessment, but may destroy explicit recognition of independencies in the influence diagram.

Thus, Figures 8.2(a) and 8.2(c) are two equivalent influence diagrams. They are equivalent in that they imply the same *possibility* of dependencies between x and y given the state of information on which the diagram was based.

An arrow joining two nodes in an influence diagram may be reversed provided that all probability assignments are based on the same set of information. For example, consider the influence diagram of Figure 9.1(a). Because the probability assignment to both x and y are made given knowledge of z, the arrow joining them can be reversed as shown in Figure 9.1(b) without making any incorrect or additional assertions about the possible independence of x and y. Figure 9.1(c) shows another example where the assignment of probability to x does not depend on the value of z, and so it might appear that no reversal of the arrow from x to y is possible. However, recall that we can always add an arrow to a diagram without making an incorrect assertion. Thus, we can change the diagram of Figure 9.1(c) to that of Figure 9.1(a), and then that of Figure 9.1(a) to that of Figure 9.1(b). The influence arrow between x and y can be reversed after an influence arrow is inserted between z and x.

The graphical manipulation procedure may yield more than one result. For example, consider the reversal of the three-node influence diagram shown in Step 1 of Figure 9.2(a). Suppose we first attempt to reverse the y to x arrow. For x and y to have only common influences, we must provide x with an influence from z (Step 2), before performing the reversal (Step 3). Since both x and z now are based on the same state of information (there are no impinging influences, i.e., arrows into x or z from any other node), the influence joining them may be reversed (Step 4). Finally, since both z and y are assigned probabilities after x is known, the influence joining them can be reversed (Step 5) .

Suppose however, that the same diagram (Step A, Figure 9.2(b)) was transformed by first reversing the arrow joining z and y (Step B), which is possible since y and z are based on the same state of information (i.e., there are no impinging influences). Then the arrow joining x and y can be reversed (Step C) because neither x nor y now have impinging influences. Both this transformation and the one in Figure 9.2(a) are correct. However, Step C of Figure 9.2(b) shows that there is no need to indicate conditioning of z on x. Step 5 of Figure 9.2(a) contains this unnecessary, but not incorrect, influence.

10. Influence Diagrams with Decision Variables

We shall now extend the concept of influence diagrams to include decision variables. We begin with a formal definition of influence diagrams.

An influence diagram is a directed graph having no loops. It contains two types of nodes:

— *Decision nodes* represented by boxes (□)

— *Chance nodes* represented by circles (◯)

Arrows between node pairs indicate influences of two types:

— *Informational influences*, represented by arrows leading into a decision node. These show exactly which variables will be known by the decision maker at the time that the decision is made.

— *Conditioning influences*, represented by arrows leading into a chance node. These show the variables on which the probability assignment to the chance node variable will be conditioned.

The informational influence on a decision node represent a basic cause/effect ordering, whereas the conditional influence into a chance node represent, as we have seen, a somewhat arbitrary order of conditioning that may not correspond to any cause/effect notion and that may be changed by application of the laws of probability (e.g., Bayes' Rule).

Figure 10.1 is an example of an influence diagram. Chance node variables $a, b, c, e, f, g, h, I, j, k, l, m$, and o all indicate aleatory variables whose probabilities must be assigned given their respective conditioning influences. Decision node variables d and n represent decision variables that must be set as a function of their respective informational influences. For example, the probability assignment to variable i is conditioned upon variables f, g, and l, and only these variables. In inferential notation, this assignment is

$$\{i \mid f, g, l, E\},$$

where E represents a special S, the initial state of information upon which the construction of the entire diagram is based. As another example, the decision variable d is set with knowledge of variables a and c, and only these variables. Thus, d is a function of a and c.

11. Node Terminology

One of the most important, but most subtle, aspects of an influence diagram is the set of possible additional influences that are not shown on the diagram. An influence diagram asserts that these missing influences do not exist.

To illustrate this characteristic of influence diagram more clearly we must make a few more definitions.

• A *path* from one node to another node is a set of influence arrows connected head to tail that forms a directed line from one node to another.

Figure 10.1 An Influence Diagram with Decision Nodes

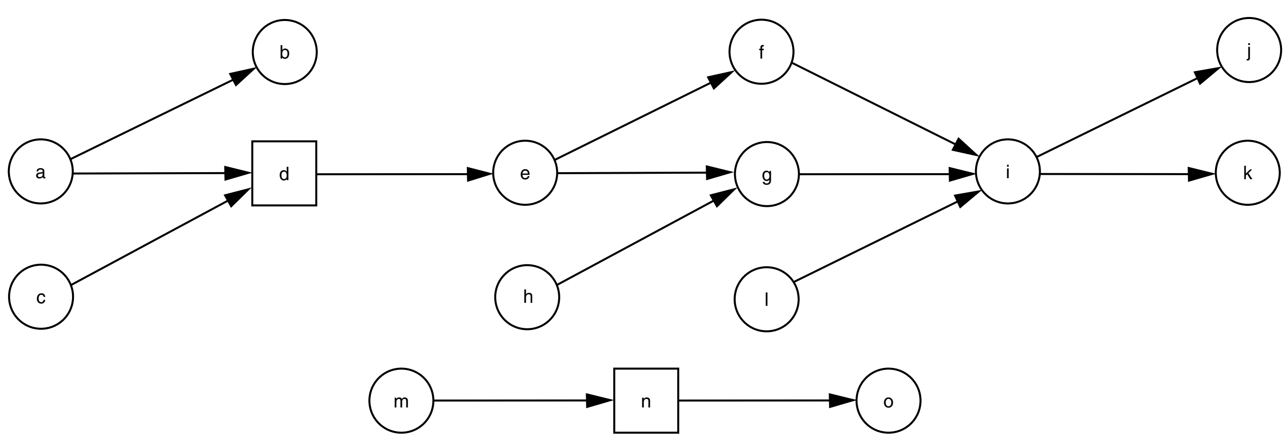

Figure 11.1 **Some Sets Defined by the Node** *g*

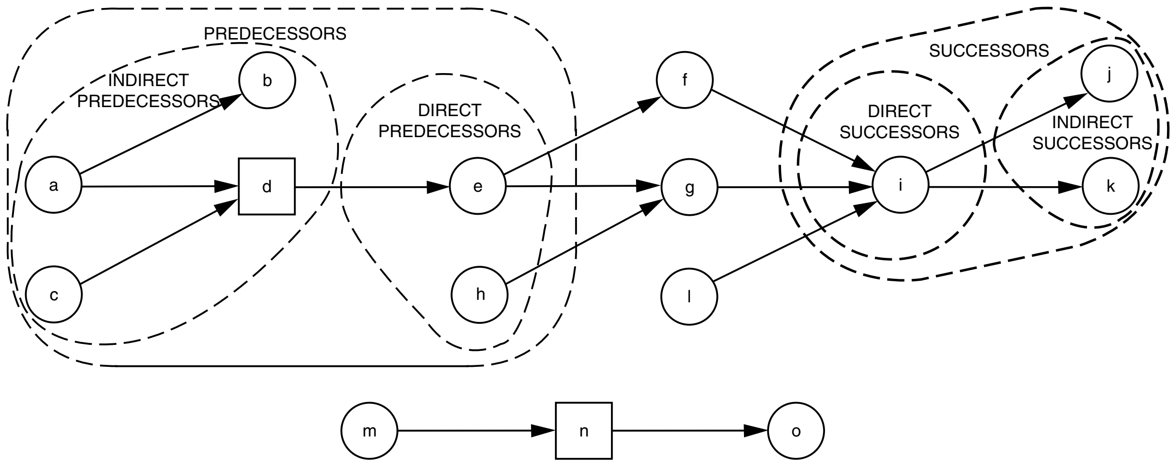

With respect to any given node we make the following definitions:

- The *predecessor set* of a node is the set of all nodes having a path leading *to* the given node.
- The *direct predecessor set* of a node is the set of nodes having an influence arrow connected *directly to* the given node.
- The *indirect predecessor set* of a node is the set formed by removing from its predecessor set all elements of its direct predecessor set.
- The *successor set* of a node is the set of all nodes having a path leading *from* the given node.
- The *direct successor set* of a node is the set of nodes having an influence arrow connected *directly from* the given node.
- The *indirect successor set* of a node is the set formed by removing from its successor set all elements of its direct successor set.

We refer to members of these sets as predecessors, direct predecessors, indirect predecessors, successors, direct successors, and indirect successors. Figure 11.1 shows the composition of each of these sets in relation to node g.

12. Missing Influences

We now are prepared to investigate the implications of influences not shown in a diagram. A given node could not have any arrows coming into it from successor nodes because this addition would form a loop in the diagram. A loop is prohibited since it could not represent any possible expansion order. However, the given node could conceivably have an additional arrow coming from any predecessor node.

The situation for decision nodes is relatively simple. The diagram asserts that the only information available when any decision is made is that represented by the direct predecessors of the decision. The addition of a new arrow, or informational influence, would usually add to the information available for decision making, and destroy the original logic of the diagram. The influence diagram asserts that this information is not directly available; however, all or part of it might be inferred indirectly from the direct predecessor set.

The situation for chance nodes is more complex. The diagram partially constrains the probabilistic conditioning (expansion) order for chance nodes. In general, the probability assignment for a given chance node, x, might be conditioned on all nonsuccessors (except for x itself). Let us call this set N_x, and let D_x be the set of direct predecessors of x. The set D_x is, of course, contained in N_x. The diagram asserts that the probability assignment to x given N_x is same as to x given D_x; that is,

$$\{x \mid N_x, E\} = \{x \mid D_x, E\}.$$

The addition of a new arrow or conditioning influence from an element of N_x to x would increase the set of direct predecessors and seem to increase the dimensionality of the conditional probability assignment. While this addition would not violate the logic of the diagram, it would cause a loss of information regarding independence of the added conditioning influence. The original diagram asserts that all information in the set N_x that is relevant to the probability assignment to x is summarized by the direct predecessors D_x. In classical terms, with respect to x, D_x is a sufficient statistic for N_x.

Returning to Figure 10.1 as an example, the probability assignment to variable g is in principle conditioned on all variables except g, i, j, and k. However, the diagram asserts that the variables on which g depends are sufficiently summarized by only e and h. This means

$$\{g \mid a, b, c, d, e, f, h, l, m, n, o, E\} = \{g \mid e, h, E\}.$$

This strong and useful assertion is based as much on the lack of arrows as on the ones that are present.

We have seen that an influence diagram indicates a specific, but possibly nonunique, order for conditioning probability assignments as well as the information available as the basis for each decision. When decision rules are specified for each decision node and probability assignments are made for each chance node, the influence diagram relationship can be used to develop the joint probability distribution for all variables.

13. Relationship of Influence Diagrams to Decision Trees

Some influence diagrams do not have corresponding decision trees. As in a decision tree, all probability assignments in an influence diagram—including the assignment limitations represented by its structure—must be founded on a base state of information, E. Unlike the nodes in a decision tree, the nodes in an influence diagram do not have to be totally ordered nor do they have to depend directly on all predecessors. The freedom from total ordering allows convenient probabilistic assessment and computation. The freedom from dependence on all predecessors allows the possibility of decisions in the diagram being made by decision makers who agree on the common base state of information E, but who differ in their ability to observe certain variables in the diagram. If the diagram represents a single decision maker who does not forget information, then the direct predecessor set of one decision must be a subset of the direct predecessor set of any subsequent decision. In the influence diagram of Figure 11.1, decisions d and n have mutually exclusive direct predecessor sets, (a, c) and (m). This situation could not be represented by a conventional decision tree.

If the informational arrows shown as dashed lines in Figure 13.1 are added to Figure 11.1, then the influence diagram can be represented by a decision tree. Many different valid decision trees can be constructed from this new influence diagram. The only conditions are that they must (1) preserve the ordering of the influence diagram and (2) not allow a chance node to be a predecessor of a decision node for which it is not a direct predecessor. For example, the chance node m must not appear ahead of decision node d in a

Figure 13.1 **An Influence Diagram Representable by a Decision Tree**

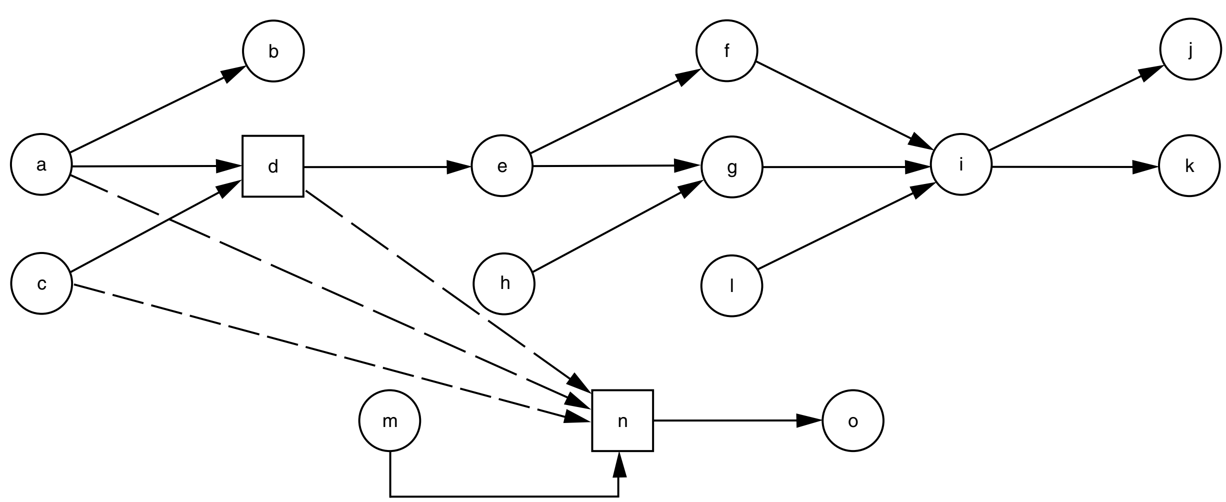

Figure 13.2 **Influence Diagram Requiring Probabilistic Manipulation Before Decision-Tree Construction**

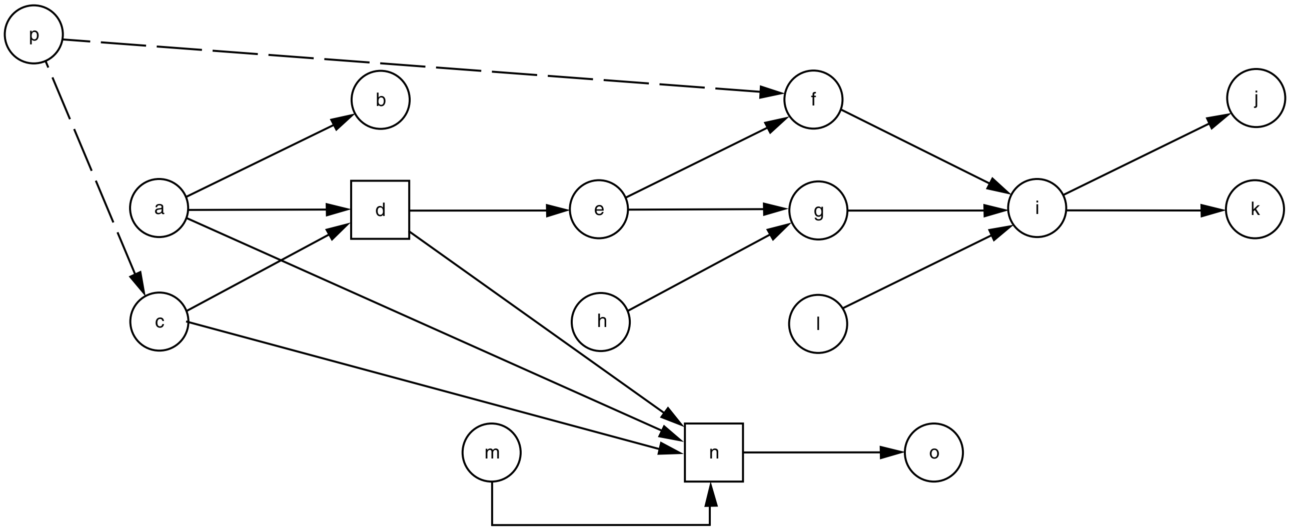

decision tree because this would imply that the decision rule for d could depend on m, which is not the case.

The situation becomes more complex when we add a node such as p in Figure 13.2. If we were to construct a decision tree beginning with chance node p, it would imply that the decision rules at nodes d and n could depend on p, which is not the actual case. Node p represents a variable that is used in the probability assignment model, but that is not observable by the decision maker at the time that he makes his decisions. In this situation, we would normally use the laws of probability (e.g., Bayes' Rule) to eliminate the conditioning of c on p. This process would lead to a new influence diagram reflecting a change in the sequence of conditioning, and could result in the inclusion of additional influence.

In Figure 13.3, the dashed arrow represents the influence as "turned around" by Bayes' Rule. The resulting diagram can be developed into a decision tree without further processing of probabilities. Also note that the change in the influence diagram required

Figure 13.3 **Influence Diagram Ready for Development into a Decision Tree**

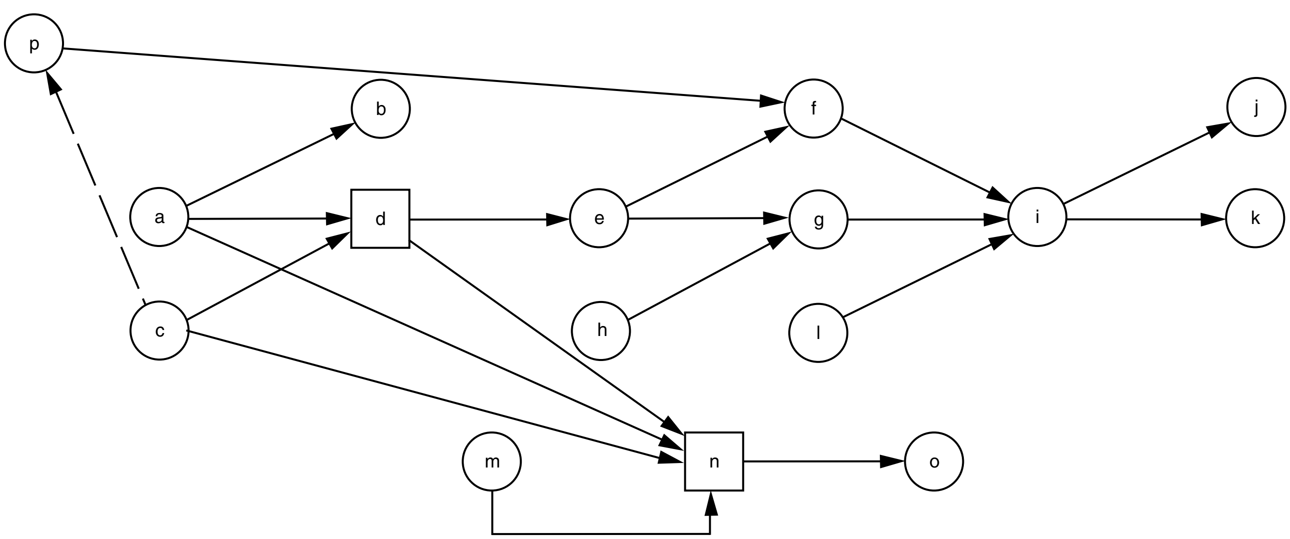

only information already specified by the original influence diagram (Figure 13.2) and its associated numerical probability assignments. Thus, it can be carried out by a routine procedure.

The foregoing considerations motivate two new definitions.

- A *decision network* is an influence diagram:

(i) that implies a total ordering among decision nodes,

(ii) where each decision node and its direct predecessors directly influence all successor decision nodes.

- A *decision-tree network* is a decision network:

(iii) where all predecessors of each decision node are direct predecessors.

Requirement (i) is the "single decision maker" condition and requirement (ii) is the "no-forgetting" condition. These two conditions guarantee that a decision tree can be constructed, possibly after some probabilistic processing. Requirement (iii) assures that no probabilistic processing is needed so that a decision tree can be constructed in direct correspondence with the influence diagram.

As an example, consider the standard inferential decision problem represented by the decision network of Figure 13.4(a). This influence diagram cannot be used to generate a decision tree directly because the decision node c has a nondirect predecessor that represents an unobservable chance variable. To convert

Figure 13.4 The Process of Converting a Decision Network to a Decision-Tree Network

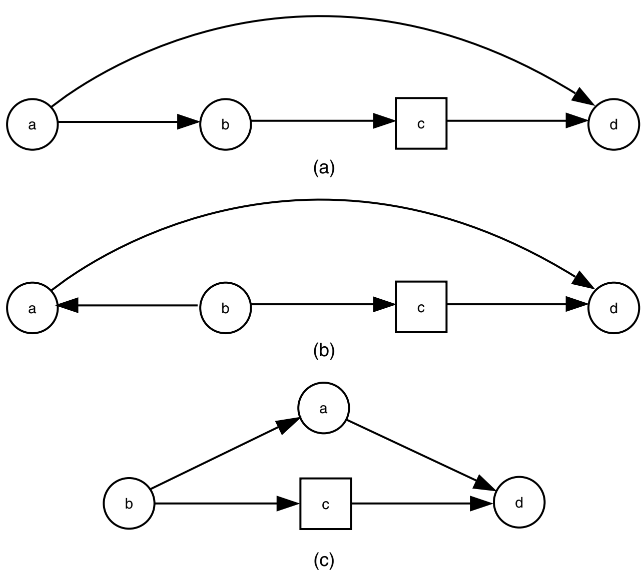

this decision network to a suitable decision-tree network, we simply reverse the arrow from a to b, which is permissible because they have only common predecessors, namely none. We thus achieve the decision tree network of Figure 13.4(b), and with redrawing we arrive at Figure 13.4(c).

Specifying the limitations on possible conditioning by drawing the influence diagram may be the most significant step in probability assignment. The remaining task is to specify the numerical probability of each chance node variable conditioned on its direct predecessor variable by a probability assessment procedure.

14. Example: The Used-Car Buyer

As an illustration of the use of influence diagrams, we consider a problem known as "The Used-Car Buyer" (Howard 1977) presented elsewhere in detail. For our purposes, we need only specify that the buyer of a used car can select among various tests T at different costs, observe their results R, choose a purchase alternative A, and then receive some value V that depends on the state of the car he bought, the outcome O. Figure 14(a) shows the influence diagram. The arrows show that the test results R depends on the test selected T and the state of the car O. The buying alternative A is chosen knowing the test selected T and its results R. The value V depends on the buying alternative chosen A, on the test selected T (as a result of the cost of the test), on the outcome O, and on the test results R. This last influence allows for the possibility that the value may depend directly on the results of the test; for example, if the testing is destructive. The outcome O does not depend on any other variable, and in particular, not on the test T,

$$\{O \mid E\} = \{O \mid T, E\}.$$

This assumption is based on the belief that the seller of the car will not switch the car to be tested as a result of the test selected.

This influence diagram is a decision network, but not a decision-tree network because node O is a predecessor of node A, but not a direct predecessor. To create a decision-tree network, we must reverse the arrow connecting node O to node R. The first step in this reversal is to assure that these nodes have a common information state. We accomplish this by adding

Figure 14 The Influence Diagram for the Used-Car Buyer Converted to a Generic Decision Tree

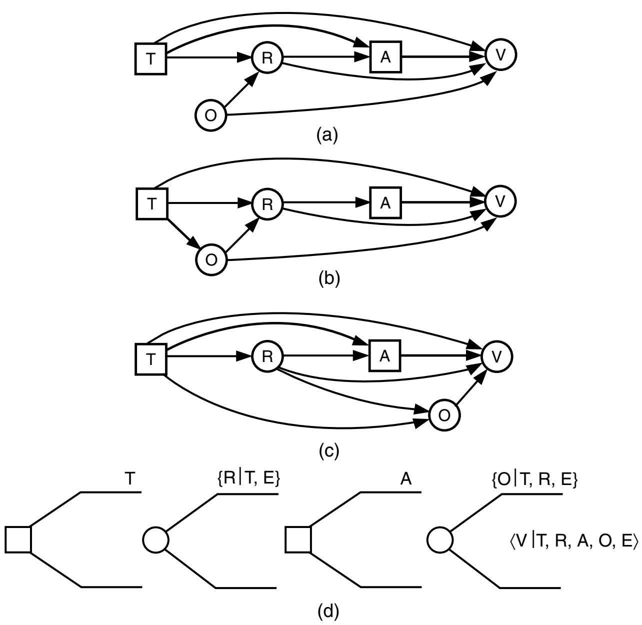

an influence from node T to node O as shown in Figure 14(b). Then we reverse the arrow from node O to node R and redraw the diagram as a decision-tree network in Figure 14(c).

This reversal means, of course, that the original probability assessments $\{R \mid T, O, E\}$ and $\{O \mid E\} = \{O \mid T, E\}$ must be changed to the probability distributions $\{R \mid T, E\}$ and $\{O \mid T, R, E\}$ according to the equation

$$\{R \mid T, E\} = \int_O \{R \mid T, O, E\}\{O \mid T, E\}$$

and Bayes' equation

$$\{O \mid T, R, E\} = \frac{\{R \mid T, O, E\}\{O \mid T, E\}}{\{R \mid T, E\}}.$$

The resulting *generic* decision tree appears in Figure 14(d), where the value assigned to each path through the tree, $\langle V \mid T, R, A, O, E\rangle$, is recorded at the endpoint of the path. The detailed calculations are shown in Howard (1977).

15. Toxic-Chemical Testing Example

To illustrate the power of influence diagrams to solve complex problems of decision making and information acquisition, we shall apply this method to a problem of toxic-chemical testing. We shall carry out the analysis under the assumption that an automated influence diagram system is available to provide the flavor of its use.

Let us suppose that a chemical having some benefits is also suspected of possible carcinogenicity. We wish to determine whether to ban, restrict, or permit its use, and also whether to undertake any information gathering regarding cancer-producing activity of the chemical or its degree of exposure to humans.

The primary decision problem can be formulated by drawing the influence diagram on the input screen of the system as in Figure 15.1. This figure shows that the system has been told that the economic value of the product and the cancer cost attributed to it both depend on the decision regarding usage of the chemical. The (probability assignment on) economic value given the usage decision is independent of the human exposure, carcinogenic activity, and the cancer cost. However, the cancer cost is dependent on the usage decision as well as on both the carcinogenic activity and human exposure levels of the chemical. The net value of the chemical given the economic value and the cancer cost is independent of the other variables. Also, human exposure and carcinogenic activity are independent.

These relationships are not necessarily obvious ones; they depend on knowledge of the problem at hand. For

Figure 15.1 Influence Diagram for Primary Decision

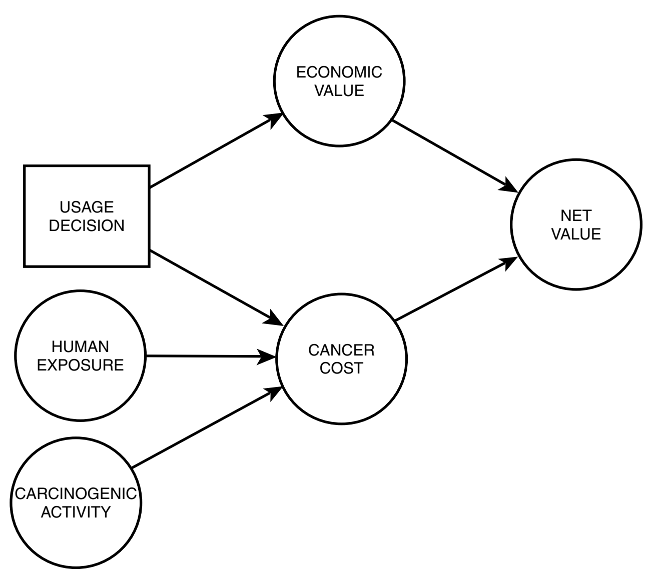

Figure 15.2 **Initial Probability Assignments**

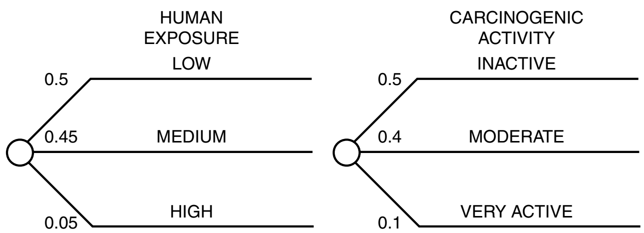

example, the economic value of a particular chemical might well depend on its chemical activity, which in turn might be closely related to its carcinogenic activity. In such a case an arrow might have to be added from "carcinogenic activity" to "economic value."

The next step is to obtain probability and value assessments corresponding to the influence diagram. The automated influence diagram system asks for a list of usage decision alternatives. In this case they are BAN, RESTRICT, and PERMIT. Next it asks for the economic value given each of these alternatives. In this case the permit alternative is considered to have a reference value of zero, the restrict alternative a substitute process cost of $1 million, and the ban alternative a substitute process cost of $5 million.

The next request is to assess possible outcomes for human exposure and carcinogenic activity along with their corresponding (unconditional) probabilities. The probability trees of Figure 15.2 illustrate these assignments. Then we are asked for the cancer cost given human exposure and carcinogenic activity levels as well as the usage decision. We assess the expected values of this cost as given in Table 15.1. Finally we state that the net value is simply the sum of the economic value and cancer cost.

All of this information is based on detailed modeling and expert judgment regarding the decision situation. Once it has been captured with the influence diagram, analysis can proceed. The automated

Table 15.1 **Cancer Cost ($millions)**

	PERMIT ALTERNATIVE			RESTRICT ALTERNATIVE			BAN ALTERNATIVE		
	Exposure			Exposure			Exposure		
Activity	Low	Med	High	Low	Med	High	Low	Med	High
Inactive	0	0	0	0	0	0	0	0	0
Moderate	0.5	5	50	0.05	0.5	5	0	0	0
Very active	10	100	1000	1	10	100	0	0	0

Figure 15.3 **Partial Decision Tree for Primary Decision**

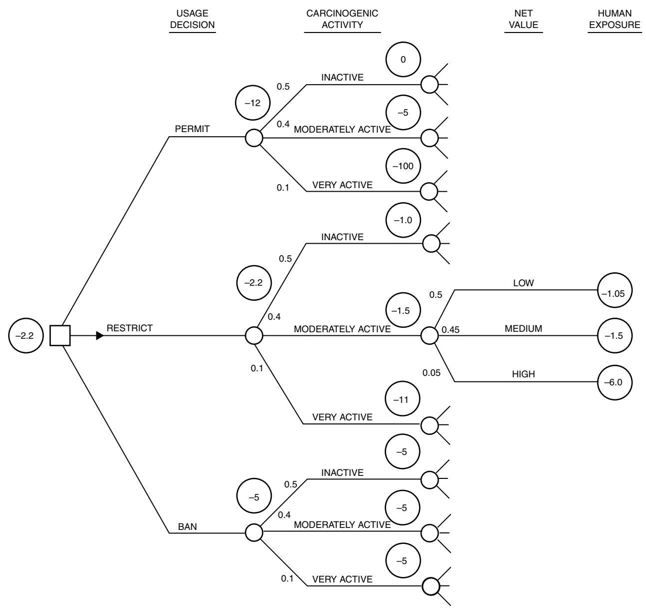

influence diagram procedure generates the appropriate decision tree, displays it if desired by the user, and determines that the best decision is to restrict usage. The expected value given this decision is a cost of $2.2 million. An example display containing this information is shown in Figure 15.3. In the example we consider only the expected value or risk-neutral case, although the case of risk aversion can be treated without difficulty.

Figure 15.1.1 **Influence Diagram Modification to Determine the Value with Perfect Information on Carcinogenic Activity**

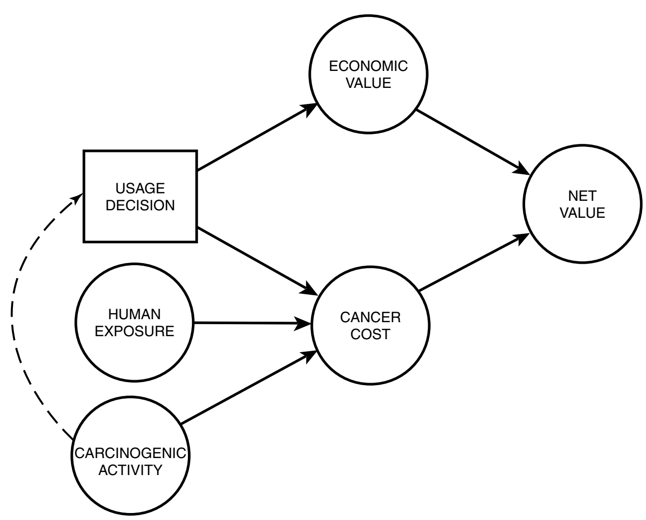

Figure 15.1.2 Partial Decision Tree for Perfect Activity Information

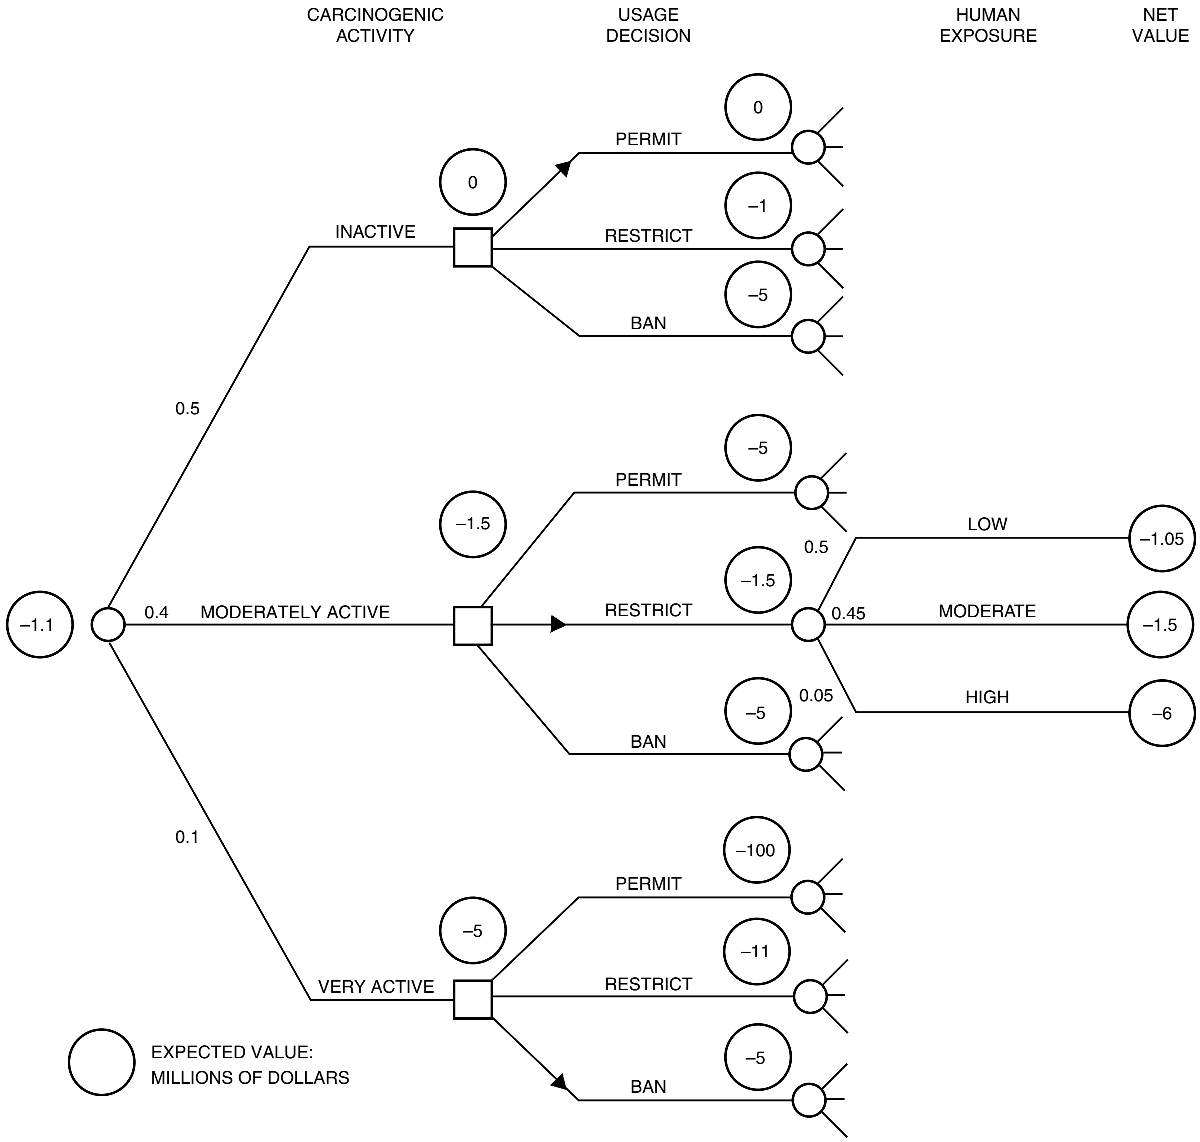

15.1. The Value of Clairvoyance (Perfect Information)

Before investigating actual information-gathering alternatives, the usual decision analysis practice is to determine the value of clairvoyance (perfect information) on the uncertain variables. The value of clairvoyance furnishes an upper limit on the value of real information gathering.

With the automatic influence diagram procedure, these calculations are trivial. For example, to calculate the value of the problem with clairvoyance on carcinogenic activity, we need only add the influence arrow indicated by a dotted line in Figure 15.1.1. This modification states that the decision maker knows the degree of carcinogenic activity when he makes the usage decision. The result is an expected cost of \$1.1 million and a decision rule to permit if inactive, restrict if moderate, and ban if very active. This means that the expected value of clairvoyance is the original \$2.2 million minus this \$1.1 million, which is \$1.1 million. Figure 15.1.2 shows a more complete display of the decision tree for this case than would be automatically generated on request of the user.

The value of clairvoyance on exposure can be calculated to be \$0.4 million by adding an influence arrow from the human exposure node to the usage decision node in Figure 15.1. The associated decision rule is to restrict if exposure is low or medium, and to ban if exposure is high.

Finally, by adding influence arrows from both the carcinogenic activity node and the human exposure node to the usage decision node, we find the value of clairvoyance on both activity and exposure to be \$1.38 million, which is less than the sum of the values of clairvoyance on each quantity separately. The decision rule is shown later in Table 15.2.1.

15.2. Value of Imperfect Information

To place a value on imperfect information we must model the information source. To be useful, the informational report must depend probabilistically on one or more of the uncertain variables in the problems. To incorporate this dependence we augment the influence diagram with a model of the information-gathering activity.

In the example at hand, it might be possible to carry out a laboratory test of the carcinogenic activity of the chemical. In this case we begin by adding a chance node to represent the report from the activity test. In Figure 15.2.1 we have added an activity test node, we have drawn an arrow to it from the carcinogenic activity node showing that the test result depends on the actual carcinogenic activity of the chemical,

Table 15.2.1 Perfect Information Summary

		EXPOSURE INFORMATION VALUE WITH = 1.8 VALUE OF = 0.40		
JOINT INFORMATION		DECISION RULE		
VALUE WITH −0.82 VALUE OF 1.38		LOW	MEDIUM	HIGH
DECISION RULE		RESTRICT	RESTRICT	BAN
ACTIVITY INFORMATION VALUE WITH = −1.10 VALUE OF + 1.10 DECISION RULE				
INACTIVE	PERMIT	PERMIT	PERMIT	PERMIT
MODERATELY ACTIVE	RESTRICT	PERMIT	RESTRICT	BAN
VERY ACTIVE	BAN	RESTRICT	BAN	BAN

Figure 15.2.1 Influence Diagram to Determine the Value with Imperfect Information on Carcinogenic Activity

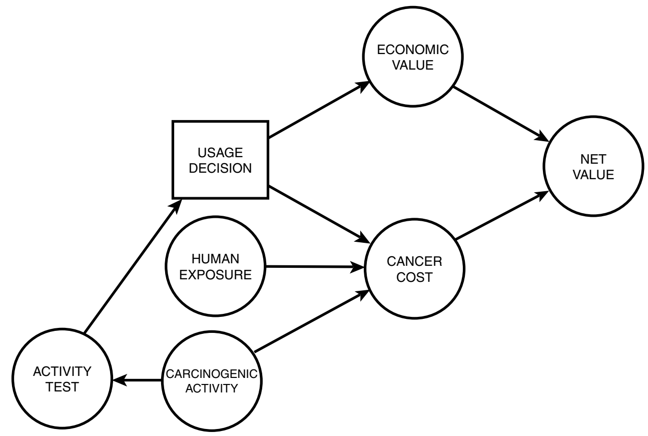

and we have drawn an arrow from the activity test to the usage decision, showing that the decision maker will know the test result when he makes the usage decision. We must also check the logic of each probabilistic statement represented in the diagram because additional knowledge, in principle, could change the probabilistic dependence elsewhere in the diagram.

The automated system would now ask us to define the test results. We reply that there are three test results called "INACTIVE," "MODERATELY ACTIVE," and "VERY ACTIVE" corresponding to the possibilities for the actual activity. However, unlike the case of perfect information, these test indications may be mislead-

Figure 15.2.2 Activity Test Probability Assignments

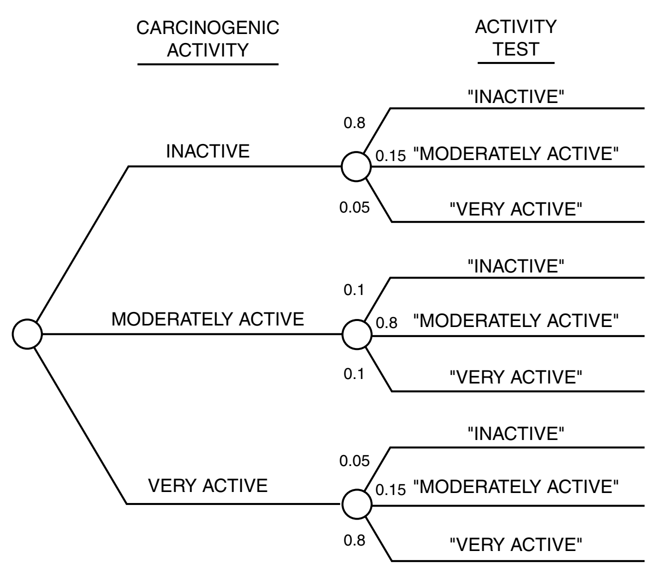

Figure 15.2.3 Exposure Test Probability Assignments

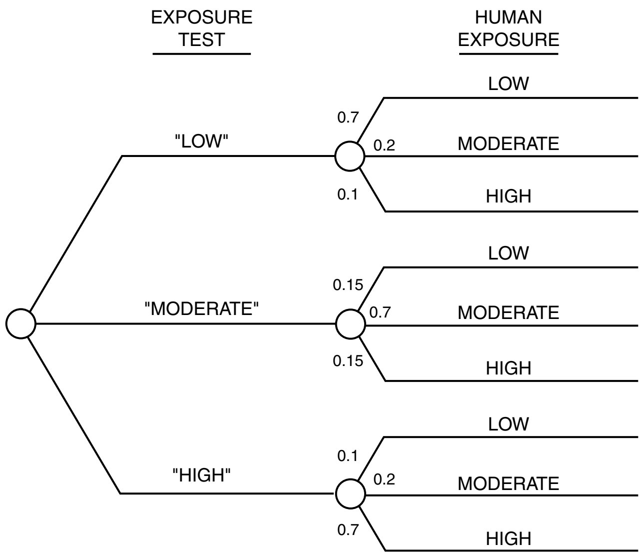

ing. The system now asks us to supply the probabilities of these test results for each state of carcinogenic activity (i.e., to supply the likelihood function). Figure 15.2.2 shows a possible display with the assigned probabilities.

All of the information needed to determine the value of the carcinogenic activity test has now been supplied. However, the influence diagram of Figure 15.2.1 is a decision network, rather than a decision-tree network, so it must be manipulated into decision-tree network form before a decision tree can be generated and evaluated. The problem is that the carcinogenic activity node precedes the usage decision node, but activity is unknown to the decision maker when he makes the usage decision. A decision tree beginning with resolution of carcinogenic activity would incorrectly give this information to the decision maker. The problem is resolved by turning around the influence arrow between carcinogenic activity and the activity test; the reversal is possible because both nodes have no impinging influences. This manipulation requires the application of Bayes' rule to determine from the original possibility assignments new assignments conditional in the opposite order. The procedure is straightforward for an automated system and results in the desired decision-tree network. In fact, a sophisticated system could determine that this manipulation was required and carry it out without being asked by the user.

Evaluation of this network yields an expected cost, given the activity test option, of \$1.96 million. Subtracting this cost from the original cost of \$2.20 million yields an expected value of \$0.24 million from a free activity test. This is the upper limit on the price the decision maker should pay for the actual test.

A test of the degree of human exposure also could be treated by adding an exposure test node to the influence diagram. The necessary probability assign-

Figure 15.2.4 Influence Diagram to Determine the Value of Imperfect Information on Both Carcinogenic Activity and Human Exposure

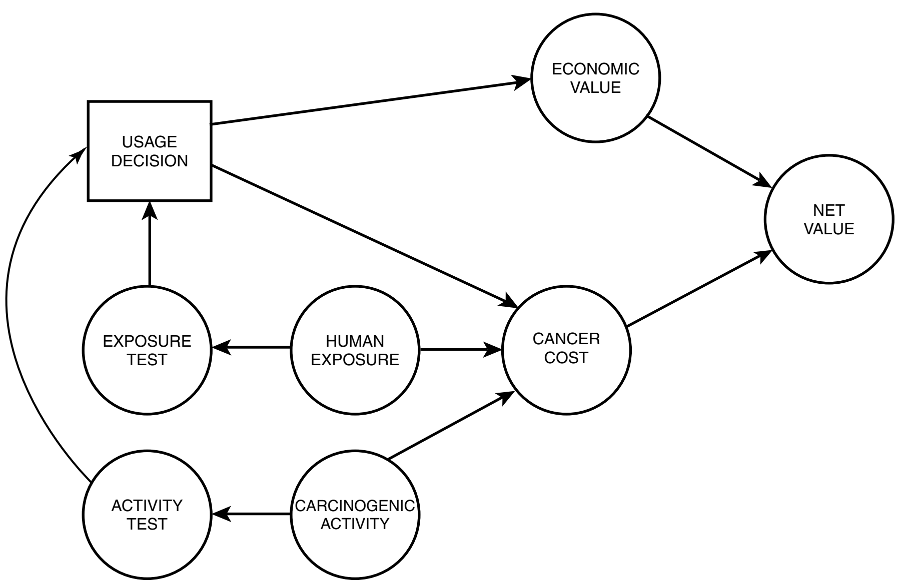

ments are shown in Figure 15.2.3. Finally, the value of testing both carcinogenic activity and human exposure could be determined by making both modifications as illustrated in Figure 15.2.4. This influence diagram indicates that given human exposure and carcinogenic activity, exposure test and activity test results are probabilistically independent.

In this example, we have shown how influence diagrams can be used to designate the initial structure of the problem. The automated system can then interact with the user to request and develop values for the probability assignments that are implicitly specified in the influence diagram. The automated system can then process the information to solve the decision problem. The automated system, not the user, develops the decision tree from the influence diagram specifications. This method allows the user to ask value of perfect information questions through simple modifications of the initial influence diagram, and to ask value of imperfect information questions by augmenting the influence diagram to model the information-gathering activities.

For this example, the decision rules and values for all information-gathering possibilities are displayed in Tables 15.2.1 and 15.2.2. The value of the situation with the specified information is given, as well as this value less the value of the primary decision (−2.2 in this case). This difference is the value of the specified information. The decision rules for joint information are given in matrix form and the ones for individual information are given along the edges of the matrix. These summaries allow the user to see easily which information is most useful. For example, Table 15.2.2 shows that imperfect exposure information is useless because the decision rule is to restrict usage regardless of the outcome of the test, even though as shown in Table 15.2.1, perfect information would be valuable. Examination of the two decision-rule matrices for the joint information cases shows four differences in choice of alternatives between the perfect and imperfect information cases. Perfect joint information is three times more valuable than imperfect information.

We have shown in this example how influence diagrams can be used to model the primary decision problem, to determine the value of perfect information on the uncertain variables, and finally to determine the value of actual, but imperfect, information. The latter calculation usually requires the application of Bayes' law. Decision-tree methods require the user to apply Bayes' law and supply the answers, or at least the formulas, for the appropriate probabilities on the decision tree. Because the influence diagram captures the logic of the problem in a more fundamental way, the user need only supply the initial probabilities that represent his model of the information-gathering activity, and an automated system can carry out the rest of the analysis. This example shows how influence diagrams can greatly simplify the probabilistic modeling and decision-making process.

Table 15.2.2 Imperfect Information Summary

		EXPOSURE INFORMATION VALUE WITH = − 2.2 VALUE OF = 0.00		
JOINT INFORMATION VALUE WITH = −1.74 VALUE OF = 0.46		DECISION RULE		
		"LOW"	"MEDIUM"	"HIGH"
		RESTRICT	RESTRICT	RESTRICT
ACTIVITY INFORMATION VALUE WITH = −1.96 VALUE OF = 0.25 DECISION RULE				
"INACTIVE"	RESTRICT	PERMIT	RESTRICT	RESTRICT
"MODERATELY ACTIVE"	RESTRICT	RESTRICT	RESTRICT	RESTRICT
"VERY ACTIVE"	BAN	RESTRICT	BAN	BAN

Acknowledgments

The authors acknowledge the contributions of Dr. Miley W. Merkhofer and Dr. Allen C. Miller, III to many of the early discussions in which the concept of influence diagrams was developed.

References

Howard, R. A. 1965. Bayesian decision models for system engineering. *IEEE Trans. Systems Sci. Cybernetics*, Vol. SSC-1, Vol. 1, November, 1965; reprinted in *Reading in Decision Analysis*, 2nd Edition, Stanford Research Institute, Menlo, CA.

Howard, R. A. 1977. The used car buyer. *Reading in Decision Analysis*, 2nd Ed. Stanford Research Institute, Menlo Park, CA.

Howard, Ronald A., James E. Matheson, Miley W. Merkhofer, Allen C. Miller III, Thomas R. Rice. 1976. *Development of Automated Aids for Decision Analysis.* DARPA Contract MDA 903-74-C-0240, SRI International, Menlo Park, CA.