4

Influence Diagrams and Decision Modelling

Richard E. Barlow & Carlos Pereira

IEOR Department, University of California, Berkeley, CA 94720, USA and Universidade de São Paulo, São Paulo, Brazil

1. MAKING DECISIONS USING INFLUENCE DIAGRAMS

In real life we are continually required to make decisions. Often these decisions are made in the face of a great deal of uncertainty. However, time and resources (usually financial) are the forcing functions for decision. That is, decisions must be made even though there may be a great deal of uncertainty regarding the unknown quantities related to our decision problem.

In considering a decision problem, we must first of all consider those things which are known as well as those things which are unknown but relevant to our decision problem. It is very important to restrict our analysis to those things which are relevant, since we cannot possibly make use of all that we know in considering a decision problem. So, the first step in formulating a decision problem is to *limit the universe of discourse* for the problem.

A decision problem begins with a list of the possible alternative decisions which may be taken. We must seriously consider all the exclusive decision alternatives which are allowed. That is, the set of decisions should be exhaustive as well as exclusive. We then attempt to list the advantages and disadvantages of taking the various decisions. This requires consideration of the possible uncertain events related to the decision alternatives. From these considerations we determine the *consequences* corresponding to decisions and possible events. At this point, in most instances, the decisions are 'weighed' and that decision is taken which is deemed to have the most 'weight'. It is this process of 'weighing' alternative decisions which concerns us.

An important distinction needs to be made between *decision* and *outcome*. A *good outcome* is a future state of the world that we prefer relative to other possibilities. A *good decision* is an action we take that is logically consistent with the alternatives we perceive, the information we have, and the preferences we feel at the time of decision. In an uncertain world, good decisions could lead to bad outcomes and bad decisions can conceivably lead to good outcomes. Making the distinction allows us to separate action from consequence. The same distinction needs to be made between *prior* and *posterior*. In retrospect, a prior distinction may appear to be very bad. However, based on prior knowledge alone, it may be the most logical assessment. The question "Suppose you have a 'bad' prior?" is essentially meaningless unless 'bad' means that a careful judgement was not used in prior assessment.

Our purpose is to introduce a 'rational method' for making decisions. By a 'rational method', we mean a method which is internally consistent, that is, it could never lead to a logical contradiction. The method we will use for making decisions can be described in terms of *influence diagrams*. Probabilistic influence diagrams need only probabilistic nodes, deterministic nodes and directed arcs. For decision making we also need *decision nodes* and *value nodes*. The following example demonstrates the need for these additional nodes.

1.1. Example (Two-headed coins). Suppose your friend tells you that he has a coin which is either a 'fair' coin or a coin with two heads. He will toss the coin and you will see which side comes face up. If you correctly decide which kind of a coin it is, he will give you $1.Otherwise you will give him $1. If you accept his offer, what *decision rule* should you choose? That is, based on the outcome of the toss, what should your decision be? In terms of probabilistic influence diagrams we only have:

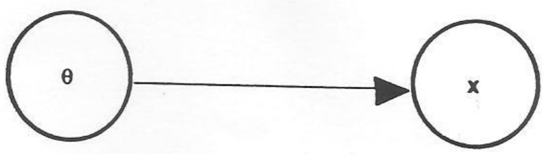

Fig. 1.

where

$$\theta = \begin{cases} \text{fair if the coin is fair} \\ \text{2-headed otherwise} \end{cases}$$

and

$$x = \begin{cases} T \text{ if the toss results in a tail} \\ H \text{ otherwise.} \end{cases}$$

Clearly, if $x = T$, you know the coin does not have 2 heads. The question is, what should you decide if $x = H$? To solve your problem, we introduce a decision node which is represented by a box or rectangle. The decision node represents the set of decisions which can be taken, namely

d_1: decide 'fair'

d_2: decide 2-headed.

Attached to the decision node is a set of allowed decision rules, which depend on the outcome of the toss, x. For example, one decision rule might be

$$\delta(x) = \begin{cases} d_1: & \text{if } x = T \\ d_2: & \text{if } x = H \end{cases}$$

Another decision rule might be

$$\delta(x) = d_1 \quad \text{for all } x.$$

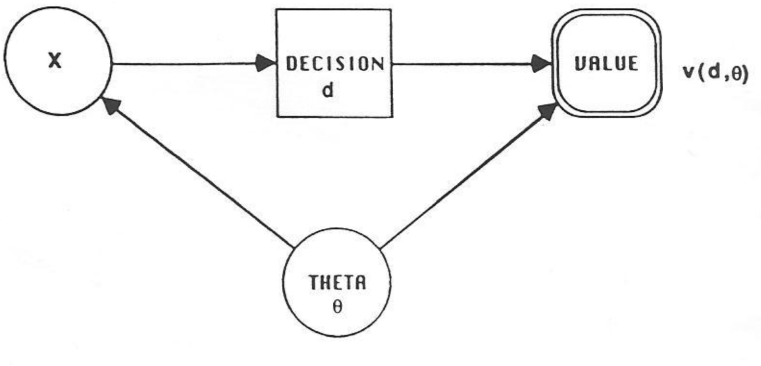

Fig. 2. Two-headed coins.

Figure 2 shows the influence diagram helpful for solving our problem. In this diagram, the double bordered node represents a value or utility node. The value node represents the set of consequences corresponding to possible pairs of decisions and states, (d, θ). Attached to the value node is the value function, $v(d, \theta)$, a deterministic function of the states of adjacent predecessor nodes. In this example

$$v(d,\theta)=\begin{cases} \$1 & \text{if } d=d_1 \text{ and } \theta=\text{fair} \\ & \text{or } d=d_2 \text{ and } \theta=2 \text{ heads} \\ -\$1 & \text{otherwise.} \end{cases}$$

The reason for initially drawing the arc $[\theta, x]$ rather than the arc $[x, \theta]$ is that, in general, it is easier to first assess $p(\theta)$ and $p(x|\theta)$ rather than to directly assess $p(x)$ and $p(\theta|x)$.

The optimal decision will depend on our initial information concerning θ, namely $p(\theta)$. However, since θ is unknown at the time of decision, there is no arc $[\theta, d]$. At the time of decision, we know x but not θ. Input arcs to a decision node indicate the information available at the time of decision. In general, there can be more than one decision node as the next example illustrates.

1.2. Example (Sequential Decision Making). Consider an urn containing white and black balls. Suppose we know that the proportion of white balls, θ, is either $\theta=\frac{1}{3}$ or $\theta=\frac{2}{3}$ but we do not know which. Our problem is to choose between two actions. One action, say a_1, would be appropriate were $\theta=\frac{2}{3}$, while a_2 would be appropriate were $\theta=\frac{1}{3}$. If we are wrong, we lose \$1. Otherwise, we lose nothing. We can, if we choose, first draw a ball from the urn at cost \$$c$ so as to learn more about θ. After observing the color of the ball drawn, say x, then we must choose either action a_1 or a_2 at cost \$$(1+c)$ if we are wrong and only cost \$$c$ if we are right.

The first decision can be either: (1) take action a_1; (2) take action a_2; or (3) draw a ball from the urn. If we draw a ball from the urn, then our second decision after drawing must be either: (1) take action a_1; or (2) take action a_2. In this problem there are two decision points and a second decision is needed only if the first decision is to continue sampling.

The following example is a decision problem of some practical importance.

1.3. Example (Inspection Sampling). Periodically, lots of size N of similar units arrive and are put into assemblies in a production line. The decision problem is whether or not to inspect units before they are put into assemblies. If we opt for inspection, what sample size, n, of the lot size, N, should be inspected? In any event, haphazard sampling to check on the proportion defective in particular lots is prudent.

Let π be the percent defective over many lots obtained from the same vendor. Suppose we believe that the vendor's production of units is in *statistical control*. That is, each unit, in our judgement, has the same chance, π, of being defective or good regardless of how many units we examine. Let $p(\pi)$ be our probability assessment for the parameter, π, based on previous experience. It could, for example, be degenerate at, say π_0.

Let k_1 be the cost to inspect one unit before installation. Let k_2 be the cost of a defective unit that gets into the production line. This cost will include the cost to tear down the assembly at the point where the defect is discovered. If a unit is inspected and found defective, additional units from another lot are inspected until a good unit is found. (We make this model assumption since all defective units which are found will be replaced at vendor's expense.) Figure 3 illustrates our production line problem. We assume the inspection process is perfect; i.e. if a unit is inspected and found good then it will also be found good in the assembly test.

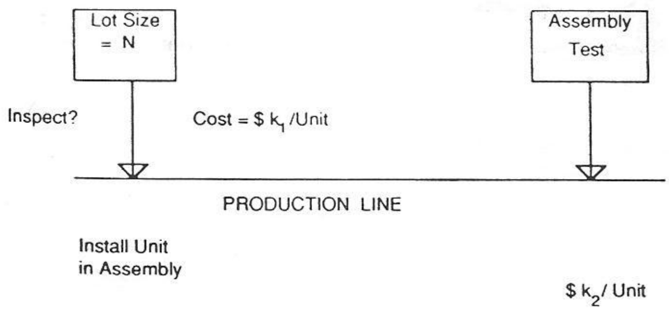

Fig. 3. Deming's inspection problem.

1.3.1 The All or None Rule

It has been suggested (cf. Deming[2]) that the optimal decision rule is always to either choose $n=0$ or $n=N$, depending on the costs involved and π, the chance that a unit is defective. This is not always valid if π is unknown.

In this example we consider the problem under the restriction that the initial inspection sample size is either $n=0$ or $n=N$. The decisions are: $n=0$ and $n=N$. Figure 4 is an influence diagram for this problem, where x = number of defectives in the lot, and y = number of additional inspections required to find good replacement units for bad units.

The value (loss) function is:

$$v[d,(x,y)] = \begin{cases} k_2 x + k_1 y & \text{if } n=0 \\ k_1 N + k_1 y & \text{if } n=N. \end{cases}$$

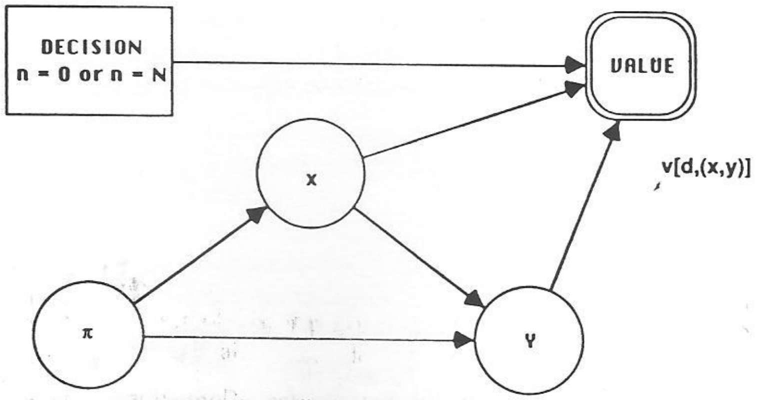

Fig. 4. Influence diagram for Deming's inspection problem.

2. DECISION INFLUENCE DIAGRAMS: DEFINITIONS AND BASIC RESULTS

A decision influence diagram (or influence diagram for short) is a diagram helpful in solving decision problems. These ideas were discussed in Howard and Matheson[3] and also in Shachter.[4] Probabilistic influence diagrams are discussed in detail in Barlow and Pereira.[5]

In Section 1 we introduced decision nodes and value nodes. Figure 5 is a typical decision node with input and output arcs.

2.1. Definition. A decision node:

(1) represents the possible decisions which may be taken at a given point in time.

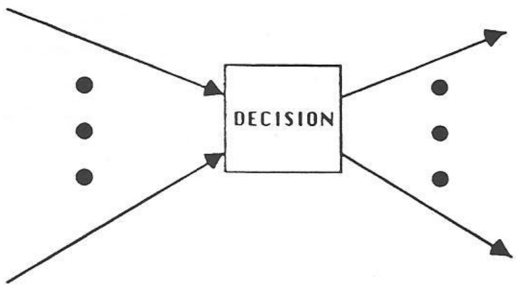

Fig. 5. Decision node.

(2) Attached to each decision node is a set of *allowed decision rules* or mappings from possible states of adjacent predecessor nodes to the set of possible alternative decisions represented by the node itself.

For example, there could be only *one* allowed decision rule corresponding to a given decision node. In this case the decision node can be replaced by a deterministic node.

2.2. Definition. A *decision rule*, δ, corresponding to a decision node is a mapping from possible states of adjacent predecessor nodes (deterministic and probabilistic as well as previous decisions) to a set of possible alternative decisions.

Decision nodes will, in general, have both directed input and directed output arcs. What makes decision nodes very different from probabilistic and deterministic nodes is that these arcs may never be reversed. Adjacent predecessor nodes to a decision node indicate information available to the decision maker at the time of that particular decision. *The value of an adjacent predecessor node to a decision node is known at the time of decision — it represents certain knowledge — not possible knowledge*. In this sense this arc is different from arcs between probability nodes which can indicate only possible dependence. Since decision nodes imply a time ordering, the corresponding directed arcs can never be reversed.

Directed arcs emanating from a decision node denote possible dependence of adjacent successor nodes on the decision taken. These arcs can likewise never be reversed.

2.3. Example. (Selling a car). Suppose you plan to sell your car tomorrow, but the finish on your car is bad. Your decision problem is whether or not to have your car painted today in order to increase the value of your car tomorrow. Note that tomorrow's selling price depends on today's decision, **d**. Let x be yesterday's blue book value for your car and c the cost of a paint job. Let θ be the price you will be able to obtain for your car tomorrow. Let $v(d, \theta)$ be the value to you of today's decision d and tomorrow's selling price, θ. The influence diagram associated with your decision problem is as follows:

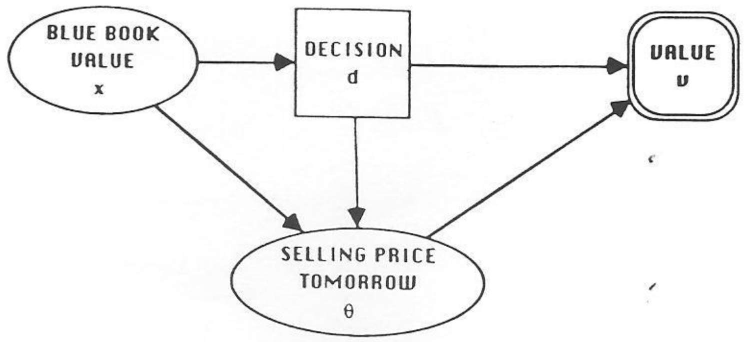

Fig. 6.

Obviously you cannot reverse arc $[x, d]$, since today's decision cannot alter yesterday's blue book value. Likewise, you cannot reverse arc $[d, \theta]$, since we cannot know today, for sure, what tomorrow will bring.

The value node is similar to a deterministic node. What makes it different is that it has no successor nodes.

2.4. Definition. A *value node* is a *sink node* that:

(1) represents possible consequences corresponding to the states of adjacent predecessor nodes; (A consequence can, for example, be a monetary loss or gain.)

(2) has an attached utility or loss function which is a deterministic function of consequences represented by the value node itself.

A value node shares with a decision node the property that input arcs can not be reversed. However if a value node, v, has a probabilistic adjacent predecessor node, say θ, and v is the only adjacent successor of θ (as in Fig. 6), then node θ can be eliminated.

We now give a formal definition of a Decision Influence Diagram.

2.5. Definition. A *Decision Influence Diagram* is an acyclic directed graph in which:

(i) nodes represent random quantities, deterministic functions and decisions;
(ii) directed arcs into probabilistic and deterministic nodes indicate *possible* dependence while directed arcs into decision nodes indicate information available at the time of decision;
(iii) attached to each probabilistic (deterministic) node is a conditional probability (deterministic) function while attached to each decision node is a set of allowed decision rules;
(iv) decision nodes are *totally ordered* and there is a directed arc (perhaps implicit) to each decision node from all predecessor decision nodes;
(v) there is *exactly one* deterministic sink node called the value node.

2.5.1. Using Decision Influence Diagrams to Model Problems

In Example 1.1 (Two-Headed Coins) you were offered the opportunity to make a dollar with also the risk of losing a dollar. Your allowable decisions were:

d_1: decide the coin is fair
d_2: decide the coin is two-headed.

Hence you might start by drawing the decision node, d, with the two allowable choices $d_1$and d_2. You might then draw the value node, v, since the value node denotes the objective of your decision analysis. Your objective is to calculate $v(d)$, the *unconditional* value function, as a deterministic function of the decision taken, d. However, it may be easier to first determine the *conditional* value function as a deterministic function of relevant unknown quantities as well as perhaps prior decisions relevant to your problem.

2.6. Definition. A value function is called *unconditional* if it only depends on the decision taken and not on relevant unknown quantities. It is called *conditional* when it also depends on relevant *unknown* quantities.

Since it is easier to think of the value function as a conditional deterministic function of your decision and the property of the coin, say θ, you also need to draw a node for θ. Since θ is unknown to you, node θ is a probabilistic node. Since the value node, v, depends on both the decision taken, d, as well as the property of the coin, θ, draw arcs $[d, v]$ and $[\theta, v]$. Attach a deterministic function, $v(d, \theta)$, to node v. Figure 7 shows the influence diagram at this stage of the analysis.

You are allowed to see the result of one coin toss and this information is available at the time you make your decision. Hence draw a node, x, for the outcome of the toss. Since the outcome of the toss (before you see it) is an unknown random quantity for you, draw a probabilistic node for x. Since the probability function for x depends on θ, draw arc $[\theta, x]$ and assess $p(x \mid \theta)$ and $p(\theta)$. Draw arc $[x, d]$ since you will know x at the time you make your decision, d. The diagram now looks like Fig. 2.

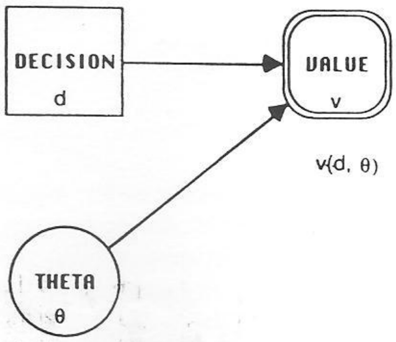

Fig. 7.

2.6.1. Decision Influence Diagram Operations

There are essentially *two* influence diagram graph operations which are used to solve decision problems expressed as influence diagrams. They are:

(1) the elimination of probabilistic nodes; and
(2) the elimination of decision nodes.

In the process of eliminating probability nodes you will obtain the *unconditional* value function, $v(d)$, as a function of decisions allowed. Having obtained $v(d)$ you can determine the *optimal* decision with respect to the unconditional value function and in this sense eliminate the decision node. The justification for both graph operations is based on the idea that you should be self-consistent in making decisions.

2.6.2. Solution of the Two-Headed Coin Problem

Having described how to construct the influence diagram (Fig. 2) for the two-headed coin Example 1.1, we now discuss its solution. Our objective is to calculate $v[d(x)]$ where $d(x)$ depends on x and is either

d_1: decide the coin is fair; or
d_2: decide the coin is two-headed.

You must eliminate θ since, as it stands, the *optimal* decision would depend on θ which is unknown. If the arc $[\theta, v]$ were the only arc emanating from θ you would be able to do this immediately by summation or by integration. Since there is also another arc, namely $[\theta, x]$, emanating out of θ, this is not possible.

2.6.3 Elimination of Node θ.

Since nodes x and θ are probabilistic nodes, you can reverse arc $[\theta, x]$ using the arc reversal operation. If you do this, $p(x|\theta)$ attached to node x is changed to $p(x)$ using the theorem of total probability while $p(\theta)$ is changed to $p(\theta|x)$ using Bayes' formula. After this arc reversal, arc $[\theta, v]$ is the only arc emanating from node θ and node θ can now be eliminated by summing $v(d, \theta)$ with respect to $p(\theta|x)$. This results in the influence diagram of Fig. 8. The deterministic function attached to node v is now

$$v[d|x] = v(d, \theta = \text{fair})\ p(\theta = \text{fair}|x) + v(d, \theta = \text{2-headed})\ p(\theta = \text{2-headed}|x).$$

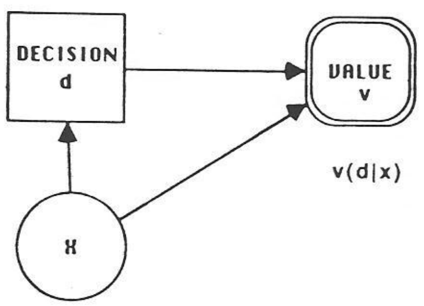

Fig. 8.

2.6.4. Elimination of the Decision Node

To solve your problem you need only eliminate the decision node. This is accomplished by maximizing $v[d(x)]$ over decision alternatives since you want to maximize your winnings.

2.6.4.1. Exercise. Let $p(\theta = \text{fair}) = p$ and determine your optimal decision rule as a function of x and p.

The following example is a decision problem of some practical importance.

2.6.4.2. Example. (Example 1.3 Inspection Sampling). We can now discuss the solution of Deming's inspection problem when the only decisions allowed are that the sample size $n=0$ or $n=N$, the lot size. Figure 4 is repeated below in Fig. 9.

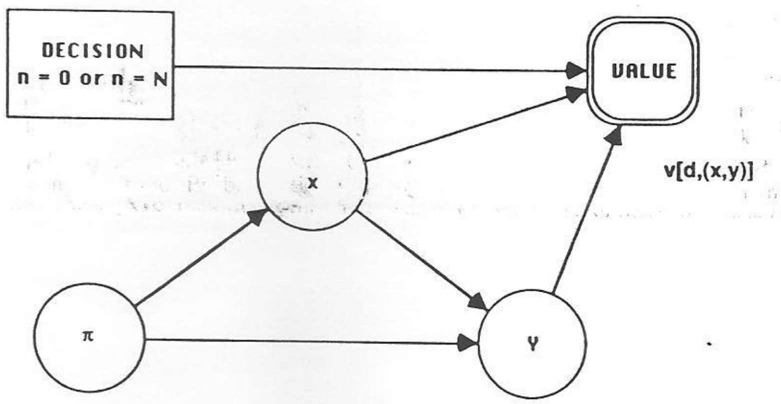

Fig. 9.

To determine whether $n=0$ or $n=N$ is best, we eliminate nodes x and y by calculating the expected value of the value function, first conditioning on π. The expected value given π is

$$E\{v[d,(x, y)]|\pi\} = \begin{cases} k_2 N\pi + k_1 E[y|\pi] & \text{if } n=0 \\ k_1 N + k_1 E[y|\pi] & \text{if } n=N \end{cases}$$

Hence $n=0$ is best if

$$\pi < k_1/k_2$$

and $n=N$ is best if

$$\pi \geqslant k_1/k_2.$$

The solution is also valid if we replace π by $E(\pi)$. If π or $E(\pi)=k_1/k_2$, then we may as well let $n=N$ since we may obtain additional information without additional expected cost.

Suppose now that we allow $0 \leqslant n \leqslant N$. If we are certain that $\pi < k_1/k_2$, i.e. $p(\pi)=0$ for $\pi > k_1/k_2$, then $n=0$ is always best. On the other hand, if we are certain that $\pi > k_1/k_2$ then $n=N$ is always best. In the intermediate case, when $p(\pi)$ straddles $\pi = k_1/k_2$, the optimal sample inspection size may be neither 0 nor N.

REFERENCES

1. Barlow, R. E., Using Influence Diagrams. In *Accelerated Life Testing and Experts' Opinions in Reliability*, eds Clarotti & Lindley, North-Holland, 1988, pp. 145–57.
2. Deming, W. E., *Out of the Crisis*, MIT, Center for Advanced Engineering Study, Cambridge, MA, 1986.
3. Howard, R. A. & Matheson, J. E., Influence diagrams. In *Readings in the Principles and Applications of Decisions Analysis*, Two volumes. eds Howard & Matheson, Strategic Decision Group, Menlo Park, CA, 1984.
4. Shachter, R., Evaluating influence diagrams. *Oper. Res.*, **34** (1986) 871–82.
5. Barlow, R. E. & Pereira, C., Conditional Independence and Probabilistic Influence Diagrams. In *Topics in Statistical Dependence*, IMS Lecture Note Series, 1990, pp. 1–14.