



ELSEVIER

European Journal of Operational Research 128 (2001) 479–498

EUROPEAN  
JOURNAL  
OF OPERATIONAL  
RESEARCH

www.elsevier.com/locate/dsw

## Theory and Methodology

# A Bayesian network approach to making inferences in causal maps

Sucheta Nadkarni, Prakash P. Shenoy \*

*School of Business, University of Kansas, Summerfield Hall, Lawrence, KS 66045-2003, USA*

Received 8 September 1998; accepted 19 May 1999

---

### Abstract

The main goal of this paper is to describe a new graphical structure called ‘Bayesian causal maps’ to represent and analyze domain knowledge of experts. A Bayesian causal map is a causal map, i.e., a network-based representation of an expert’s cognition. It is also a Bayesian network, i.e., a graphical representation of an expert’s knowledge based on probability theory. Bayesian causal maps enhance the capabilities of causal maps in many ways. We describe how the textual analysis procedure for constructing causal maps can be modified to construct Bayesian causal maps, and we illustrate it using a causal map of a marketing expert in the context of a product development decision. © 2001 Elsevier Science B.V. All rights reserved.

*Keywords:* Causal maps; Cognitive maps; Bayesian networks; Bayesian causal maps

---

### 1. Introduction

Recently, there has been a growing interest in the use of causal maps to represent domain knowledge of decision-makers (Huff, 1990; Eden et al., 1992; Laukkanen, 1996). Causal maps are cognitive maps that represent the causal knowledge of subjects in a specific domain. Causal maps (also called cognitive maps, cause maps, etc.) have been used extensively in the areas of policy analysis (Axelrod, 1976) and management sciences (Klein and Cooper, 1982; Ross and Hall, 1980) to represent salient factors, knowledge, and conditions that influence decision making.

In the last few years, researchers have emphasized that causal maps may be used as tools to facilitate decision making and problem solving within the context of organizational intervention (Eden, 1992; Laukkanen, 1996). Studies in artificial intelligence have also emphasized the importance of domain knowledge in constructing decision tools, especially in situations where domain knowledge is crucial and availability of data is scarce (Heckerman, 1996). Since causal maps represent domain knowledge more descriptively than other models such as regression or structural equations, they are more useful decision tools. An important aspect of a decision model is the ability to make inferences. Inference refers to drawing conclusions based on a premise. To make an inference is to come to believe a new fact based on other information. Inference is important in

---

\* Corresponding author. Tel.: +1-785-864-7551.

decision analysis for two reasons. First, it allows one to make predictions in case of interventions. For example, a marketing manager may want to know how a new marketing strategy will affect the company profits. Second, it provides a prescriptive framework for decision making which can be useful to individuals other than the expert in a specific decision context.

Current techniques used to analyze causal maps provide a qualitative interpretation of the variables representing a decision problem by focusing on the structure of a causal map (Axelrod, 1976; Bougon, 1992; Eden et al., 1979; Ross and Hall, 1980). This qualitative approach is useful in describing a decision problem or in making inferences in case of very simple maps with few variables. However, drawing inferences about variables in complex causal maps may be a non-trivial task because of two reasons. First, causal maps do not model uncertainty associated with the decision variables (Huff, 1990; Laukkanen, 1996). All variables in the maps have the same level of certainty. Identifying the level of uncertainty is important in making inferences because observations of variables may be uncertain, information may be incomplete, or the variables involved may be vague. Second, causal maps provide a static representation of the decision variables. They do not depict how beliefs of decision-makers about some target variables change when decision-makers learn additional information about relevant situational factors or decision options represented in the map. Such a dynamic approach is important in not only drawing inferences but also in learning about causal relations representing complex and uncertain decisions (Heckerman, 1996).

Recent advances in artificial intelligence such as Bayesian networks allow us to use causal maps to make inferences for decision making. This study uses the theory of Bayesian networks to suggest a quantitative way to make inferences in causal maps. A Bayesian network is a specification of a joint probability distribution of several variables in terms of conditional distributions for each variable. In general, the relations represented in a Bayesian network do not have to be causal relations. However, a Bayesian network representation is sparse and efficient when the relations are

causal (Shachter and Heckerman, 1987), and sparse Bayesian networks allow us to make inferences efficiently even when we have many variables in the network (Pearl, 1986; Lauritzen and Spiegelhalter, 1988; Shenoy and Shafer, 1990). Bayesian networks in which the dependence relations are causal are also called causal belief networks, causal probabilistic networks, etc. We simply refer to such representations as *Bayesian causal maps*. Using the framework of Bayesian networks, Bayesian causal maps can account for uncertainty associated with the variables in the map. Bayesian causal maps can use the evidence propagation algorithms of Bayesian networks to analyze the sensitivity of variables of interest (e.g. outcome variables) to additional findings about other variables (e.g. situational variables).

An outline of the remainder of the paper is as follows. In Section 2, we discuss the definition and analysis of causal maps. In Section 3, we discuss Bayesian networks, their semantics, and the process of making inferences. In Section 4, we propose a procedure for constructing a Bayesian causal map starting from a causal map. In Section 5, we discuss a case study of product development decision in a technology organization. Finally, in Section 6, we conclude with a summary and a statement of future research.

## 2. Causal maps

Causal maps are directed graphs that represent the cause–effect relations embedded in managers' thinking (Fiol and Huff, 1992). Eden et al. (1992) defines a causal map as a “directed graph characterized by a hierarchical structure which is most often in the form of a means/end graph.” Causal maps express the judgment that certain events or actions will lead to particular outcomes. The three components of a causal map are node representing *causal concept*, link representing *causal connection* among causal concepts (can be positive or negative), and strength representing *causal value* of a causal connection. Fig. 1 depicts a causal map of a marketing expert in the context of a new product decision.

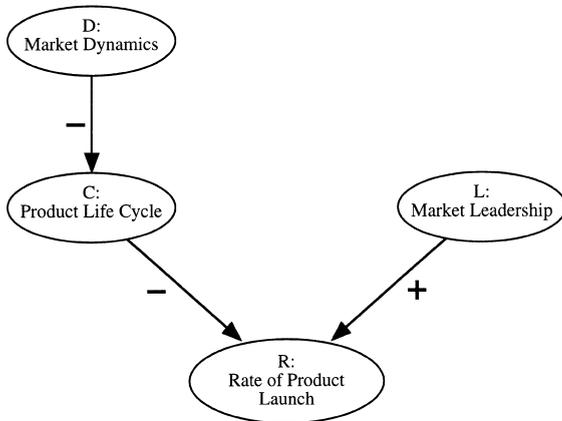


Fig. 1. A causal map of a marketing expert for a New Product Decision.

This causal map consists of four causal concepts: Market Dynamics, Product Life Cycle, Market Leadership and Rate of Product Launch. The causal concepts are linked to each other through causal connections represented by unidirectional arrows (either positive or negative). The concept at the tail of an arrow is the ‘cause’ of the concept at the head of the arrow. For example, high market dynamics leads to a short product life cycle (negative relation). A short product life cycle causes the rate of product launch of a company to be high (negative relation). Finally, market leadership causes a company to launch products at a high rate (i.e. a positive relation).

Causal maps make the following three assumptions about cognition in the context of decision making (Huff, 1990).

1. Causal associations are a major way in which decision problems can be described and understood;
2. Causality is the primary form of post-hoc explanation of decision outcomes; and
3. Choice among alternative decision actions involves causal relations.

The primary emphasis in causal maps is in identifying dependence between variables in terms of explanation–consequence or means–end relationships. Different techniques have been employed to construct causal maps including structured methods such as the self-Q method (Bougon, 1983), open method based on personal

construct theory (Eden et al., 1979), and textual analysis approach (Axelrod, 1976). We use the textual analysis method for constructing Bayesian causal map. The textual analysis method is deemed appropriate because it does not require the researcher to identify causal concepts prior to or during the interview with subjects. The advantage of this method is that it enables the researcher to let the data drive the concepts rather than eliciting an individual’s cognition from a predefined set of concepts (Carley and Palmquist, 1992). The unintrusive nature of this method is especially important in representing an expert’s description and understanding of an unstructured decision problem.

Causal maps have been useful in practice. For example, Axelrod (1976) describes a causal map derived from text to represent a decision-maker’s beliefs concerning the relationships between factors in the public health system. Similarly, Swan (1995) describes textually derived causal maps of key managers to identify important factors affecting the decision of implementing computer-aided production management technologies in manufacturing firms. Such causal maps have also been used to represent cause–effect beliefs of decision-makers in strategic alliance decisions (Fiol, 1990), and marketing distribution decisions (Laukkanen, 1996).

Eden and Jones (1980) point out that there is no general approach to analyze causal maps. In general, causal maps can be analyzed qualitatively or quantitatively. Eden’s causal mapping technique, most widely applied to problem construction, is an aid to defining the nature of policy problem. It involves conflicting and qualitative views of the problem. This has been widely applied in strategic management for organizations and is constantly being developed (Eden, 1991; Eden, 1992; Eden et al., 1992; Eden and Ackermann, 1993). Causal maps can also be analyzed quantitatively using different analytical tools. For example, quantitative analysis could be used to explore the numerical implications of possible policy options based on constructed causal maps. Quantitative analysis provides an objective basis for assessments of causal maps so that subjective biases in qualitative analysis could be reduced.

Quantitative tools used to analyze causal maps largely differ depending on the context of application. Matrix algebra and network analytic methods have been used to explore the structural properties of causal maps (Axelrod, 1976; Bougon et al., 1977; Eden et al., 1979). System dynamics has been used to analyze causal maps from a dynamic perspective (Forrester, 1961; Wolstenholme and Coyle, 1983). Chen et al. (1979) uses an unstructured variation of causal maps to construct decision trees. They use these decision trees for decision analysis using probability assessments and utility functions. More recently, Wang (1996) describes a neural network model to compare the dynamic properties of causal maps. In this study, we use the probabilistic inference procedures of Bayesian networks to make inferences about variables in causal maps.

### 3. Bayesian networks

In this section, we briefly describe the semantics of Bayesian networks and their use in making inferences. A procedure for constructing a Bayesian

network starting from a causal map is described in Section 4.

Bayesian networks have their roots in attempts to represent expert knowledge in domains where expert knowledge is uncertain, ambiguous, and/or incomplete. Bayesian networks are based on probability theory. A primer on Bayesian networks is found in Spiegelhalter et al. (1993).

A Bayesian network model is represented at two levels, qualitative and quantitative. At the qualitative level, we have a directed acyclic graph in which nodes represent variables and directed arcs describe the conditional independence relations embedded in the model. Fig. 2 shows a Bayesian network consisting of four discrete variables: Market Dynamics (*D*), Product Life Cycle (*C*), Market Leadership (*L*), and Rate of Product Launch (*R*). At the quantitative level, the dependence relations are expressed in terms of conditional probability distributions for each variable in the network. Each variable *X* has a set of possible values called its *state space* that consists of mutually exclusive and exhaustive values of the variable. In Fig. 2, e.g., Market Dynamics has two states: ‘high’ and ‘low’; Market Leadership has two

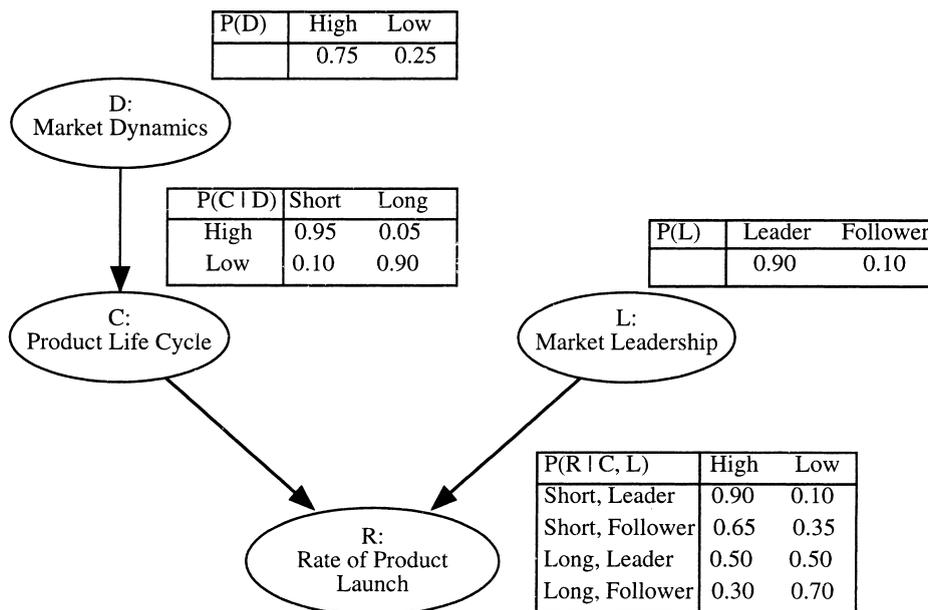


Fig. 2. A Bayesian network with conditional probability tables.

states: ‘leader’ and ‘follower’; Product Life Cycle has two states: ‘short’ and ‘long’; and Rate of Product Launch has two states: ‘high’ and ‘low’. If there is an arc pointing from  $X$  to  $Y$ , we say  $X$  is a *parent* of  $Y$ . For each variable, we need to specify a table of conditional probability distributions, one for each configuration of states of its parents. Fig. 2 shows these tables of conditional distributions –  $P(D)$ ,  $P(C|D)$ ,  $P(L)$ , and  $P(R|C, L)$ .

### 3.1. Semantics

A fundamental assumption of a Bayesian network is that when we multiply the conditionals for each variable, we get the joint probability distribution for all variables in the network. In Fig. 2, e.g., we are assuming that

$$P(D, C, L, R) = P(D) \otimes P(C|D) \otimes P(L) \\ \otimes P(R|C, L),$$

where  $\otimes$  denotes pointwise multiplication of tables. The rule of total probability tells us that

$$P(D, C, L, R) = P(D) \otimes P(C|D) \otimes P(L|D, C) \\ \otimes P(R|D, C, L).$$

Comparing the two, we notice that we are making the following assumptions:  $P(L|D, C) = P(L)$ , i.e.,  $L$  is independent of  $D$  and  $C$ ; and  $P(R|D, C, L) = P(R|C, L)$ , i.e.,  $R$  is conditionally independent of  $D$  given  $C$  and  $L$ .

Notice that we can read these conditional independence assumptions directly from the Bayesian network graph as follows. Suppose we pick a sequence of the variables such that for all directed arcs in the network, the variable at the tail of each arc precedes the variable at the head of the arc in the sequence. Since the directed graph is acyclic, there always exists such a sequence. In Fig. 2, e.g., one such sequence is  $D C L R$ . Then, the conditional independence assumptions can be stated as follows. For each variable in the sequence, we are assuming that it is conditionally independent of its predecessors in the sequence given its parents. The essential point here is that missing arcs (from a node to its successors in the sequence) signify

conditional independence assumptions. Thus the lack of an arc from  $D$  to  $L$  signifies that  $L$  is independent of  $D$ ; the lack of an arc from  $C$  to  $L$  signifies that  $L$  is independent of  $C$ ; and the lack of an arc from  $D$  to  $R$  signifies that  $R$  is conditionally independent of  $D$  given  $C$  and  $L$ .

In general, there may be several sequences consistent with the arcs in a Bayesian network. In such cases, the list of conditional independence assumptions associated with each sequence can be shown to be equivalent using the laws of conditional independence (Dawid, 1979). Pearl (1988) and Lauritzen et al. (1990) describe other equivalent graphical methods for identifying conditional independence assumptions embedded in a Bayesian network graph.

Unlike a causal map, the arcs in a Bayesian network do not necessarily imply causality. The (lack of) arcs represent conditional independence assumptions. How are conditional independence and causality related? Conditional independence can be understood in terms of relevance. If  $Z$  is conditionally independent of  $X$  given  $Y$ , then this statement can be interpreted as follows. If the true state of  $Y$  is known, then in assigning probabilities to states of  $Z$ , the states of  $X$  are irrelevant. In practice, the notion of direct causality is often used to make judgments of conditional independence. Consider a situation where  $X$  directly causes  $Y$  and  $Y$  in turn directly causes  $Z$ , i.e., the causal effect of  $X$  on  $Z$  is completely mediated by  $Y$ . Then it is clear that although  $X$  is relevant to  $Z$ , if we know the true state of  $Y$ , further knowledge of  $X$  is irrelevant (for assigning probabilities) to  $Z$ , i.e.,  $Z$  is conditionally independent of  $X$  given  $Y$ . This situation is represented by the Bayesian network  $X \rightarrow Y \rightarrow Z$  in which there is no arc from  $X$  to  $Z$ . As another example, consider the situation where  $X$  directly causes  $Y$  and  $X$  also directly causes  $Z$ . Although knowledge of  $Y$  is relevant to  $Z$  (if  $Y$  is true then it is more likely that  $X$  is true which in turn means that it is more likely that  $Z$  is true), once we know the true state of  $X$ , then further knowledge of  $Y$  is irrelevant to  $Z$ , i.e.,  $Y$  is conditionally independent of  $Z$  given  $X$ . This situation is represented by the Bayesian network  $Z \leftarrow X \rightarrow Y$  in which there is no arc from  $Y$  to  $Z$  or vice-versa. Finally as a third example, consider the situation

where  $X$  and  $Y$  are two independent direct causes of  $Z$ , i.e.,  $X$  and  $Y$  are unconditionally independent. But if we learn something about the true state of  $Z$ , then  $X$  and  $Y$  are no longer irrelevant to each other (if  $Z$  is believed to be true and  $X$  is false, then it is more likely that  $Y$  is true), i.e.,  $Y$  is not conditionally independent of  $X$  given  $Z$ . This situation is represented by the Bayesian net  $X \rightarrow Z \leftarrow Y$  in which there is no arc from  $X$  to  $Y$  or vice-versa.

### 3.2. Making inferences

Inference (also called probabilistic inference) in a Bayesian network is based on the notion of *evidence propagation*. Evidence propagation refers to an efficient computation of marginal probabilities of variables of interest, conditional on arbitrary configurations of other variables, which constitute the observed evidence (Speigelhalter et al., 1993). Probabilistic inference refers to the process of computing the posterior marginal probability distributions of a set of variables of interest after obtaining some observations of other variables in the model.

Once a Bayesian network is constructed, it can be used to make inferences about the variables in the model. The conditionals given in a Bayesian network representation specify the *prior* joint distribution of the variables. If we observe (or learn about) the values of some variables, then such observations can be represented by tables where we assign 1 for the observed values and 0 for the unobserved values. Then the product of all tables (conditionals and observations) gives the (unnormalized) *posterior* joint distribution of the variables. Thus, the joint distribution of variables changes each time we learn new information about the variables.

In theory, the posterior marginal probability of a variable  $X$ , say  $P(X)$ , can be computed from the joint probability by summing out all other variables except  $X$  one by one. In practice, such a naive approach is not computationally tractable when we have a large number of variables because the joint distribution has an exponential number of states and values. The key to efficient inference lies

in the concept of *local computation* where we compute the marginal of the joint without actually computing the joint distribution. A crucial feature of a Bayesian network is that it describes a joint distribution as built out of local relationships within groups of variables – such as a node and its parents. Instead of tackling the whole collection of variables simultaneously, we can use the concept of factorization. Factorization involves breaking down the joint probability distributions into subgroups called factors or belief universes in such a way that the naive computations described above need only be performed within each belief universe. Since the state space of such belief universes is much smaller than that of the joint probability distribution, the calculations become manageable.

To ensure that we obtain the correct answers when considering all the variables together, we need to develop ways for the belief universes to communicate with each other, so that (for example) the effect of conditioning on a variable in one belief universe can be felt by those in another. This is called the procedure of evidence propagation, i.e., how information propagates between the different variables in the Bayesian network. Several propagation architectures have been developed which provide different algorithms to pass messages between the belief universes. Pearl (1986) developed a message-passing scheme that updates the probability distributions for each node in response to observations of one or more variables. Lauritzen and Speigelhalter (1988), Jensen et al. (1990) and Shenoy and Shafer (1990) devised propagation algorithms that first transforms the Bayesian network into a tree where each node in the tree corresponds to a belief universe. The algorithm then exploits several mathematical properties of this tree to perform probabilistic inference.

These propagation algorithms can be used to make two types of inferences. Often we are interested in the values of some target variables. In this case, we make inference by computing the marginal of the posterior joint distribution for the variables of interest. Consider the situation described by the Bayesian network in Fig. 2. Suppose we are interested in the true state of Rate of

Product Launch ( $R$ ). Given the prior model (as per the probability tables shown in Fig. 2), the marginal distribution of  $R$  is 0.71 for high and 0.29 for low. Now suppose we learn that Market Dynamics is low. The posterior marginal distribution of  $R$  changes to 0.52 for high and 0.48 for low. Suppose we further learn that the state of Market Leadership is follower. Then the marginal distribution of  $R$  changes to 0.34 for high and 0.66 for low. This type of inference is referred to as ‘*sum propagation*’.

Sometimes we are more interested in the configuration of all variables (“the big picture”) rather than the values of individual variables. In this case, we can make inferences by computing the mode of the posterior joint distribution, i.e., a configuration of variables that has the maximum probability. Consider again the situation described by the Bayesian network in Fig. 2. Given the prior model (as per the probability tables shown in Fig. 2), the mode of the prior joint distribution is (high market dynamics, leader, short product life cycle, high rate of product launch). Now suppose we learn that Market Dynamics is low. The mode of the posterior joint distribution changes to (low market dynamics, leader, long product life cycle, high/low rate of product launch) (both configurations have the same maximum probability). Suppose we further learn that the state of Market Leadership is follower. Then the mode of the joint distribution changes to (low market dynamics, leader, long product life cycle, low rate of product launch). This type of inference is referred to as ‘*max propagation*’.

The results of inference are more sensitive to the qualitative structure of the Bayesian network than the numerical probabilities (Darwiche and Goldszmidt, 1994). For decision making, the inference results are robust with respect to the numerical probabilities (Henrion et al., 1994).

There are several commercial software tools such as Hugin [www.hugin.com] and Netica [www.norsys.com] that automate the process of inference. These tools allow the user to enter the Bayesian network structure graphically, enter the numerical details, and then do inference of either type. The results of the inference are then shown graphically using bar charts.

## 4. Constructing Bayesian causal maps

In this section, we sketch a procedure for constructing Bayesian causal maps starting from a causal map based on concepts from Bayesian networks and causal maps. The procedure for constructing Bayesian causal maps proceeds in two stages: qualitative and probabilistic.

### 4.1. Qualitative stage

In this stage, the structure of the causal maps is modified for two reasons. First, to eliminate some of the limitations of the modeling procedure used in the derivation of causal maps. Second, to make the causal maps compatible with the Bayesian network approach. In the qualitative stage, we address four major modeling issues. These issues are discussed in the following paragraphs.

#### 4.1.1. Conditional independencies

A network model can be either a dependence map (D-map) or an independence map (I-map) (Pearl, 1988). A D-map guarantees that vertices (concepts) found to be connected are indeed dependent; however, it may display a pair of dependent concepts as a pair of separated vertices (concepts). An I-map, on the other hand, guarantees that vertices (concepts) found to be separated are indeed conditionally independent, given other variables. However, it may display a pair of independent concepts as connected vertices (concepts). A model that is both a D-map and an I-map is called a *perfect* map.

A causal map is a directed graph that depicts causality between variables as perceived by individuals. Since an arrow between two variables implies dependence, it is a D-map. The absence of an arrow between two variables does not imply a lack of dependence. In other words, a causal map does not guarantee that variables found to be separated correspond to independent concepts, i.e., it is not an I-map.

Bayesian networks, on the other hand, are I-maps. Given a sequence of variables, an absence of arrow from a variable to its successors in the sequence implies conditional independence between

the variables. Conditional independence is an important issue in making inferences since it specifies the relevance of information on one variable in making inference on another. Thus if we are to regard a causal map as a Bayesian network, it is important to first convert the causal map model from a D-map to an I-map. This can be done in consultation with the expert (whose causal map is at issue) by ensuring that all dependencies are depicted in the causal map. An adjacency matrix can be used to achieve this. An adjacency matrix is the matrix of all possible direct cause–effect relations among the variables where the  $(i, j)$  entry indicates the existence of a direct causality relation from variable  $i$  to variable  $j$  (Axelrod, 1976). The adjacency matrix is represented as an  $n$  by  $n$  matrix in which the columns and rows are labeled with the names of the variables. The rows represent causes and columns represent effects. For each pair of variables, the subject (whose causal map is at issue) is asked to specify if there is a causal relation (scored initially as 1 for yes, and 0 for no) and if yes, to specify which is the cause and which is the effect. Subject should be instructed that if (s)he considers the two variables to have reciprocal influences, (s)he should indicate which one has the more dominant causal influence (to eliminate the occurrence of direct reciprocal causality wherein a variable is both a direct cause and an effect of another variable). Once the subject has identified the existence of a causal relation, (s)he can then be asked to indicate nature of the relation (positive or negative). Thus when the subject is finished, each cell in the adjacency matrix will have one of the three values:  $-1$  for a negative causal relation,  $0$  for no causal relation, or  $+1$  for a positive causal relation.

Fig. 3 indicates a part of a Bayesian causal map of a marketing expert relating to a product development decision. The solid arrow (from Product Life Cycle to Rate of Product Launch) was identified in the original causal map. The dashed arrows (from Market Dynamics to Product Life Cycle, from Market Dynamics to Sales Uncertainty, from Market Share Distribution to Sales Uncertainty, from Sales Uncertainty to Product Risk, and from Competition Strategy to Price Level) were identified in the process of making the

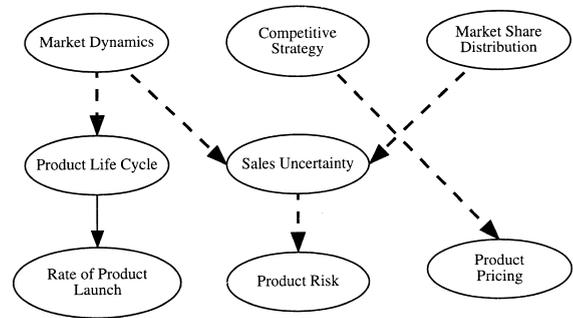


Fig. 3. Making a causal map a perfect map.

original causal map an I-map. These unmodeled dependencies affect the inferences made on variables – Sales Uncertainty, Product Risk and Price Level. In the original causal map, inference on Sales Uncertainty is not conditional on any other variables in the map. For example, Market dynamics and Market Share Distribution are not relevant in making inference on Sales Uncertainty. However, in the modified causal map, inference on Sales Uncertainty is conditional on Market Share Distribution as well as Market Dynamics. This example illustrates how unmodeled dependencies can change the inferences that can be made on the various variables in the causal maps.

After incorporating all possible dependencies between variables, the revised causal map is a perfect map – a D-map as well as an I-map.<sup>1</sup> An arrow between two variables in the map implies a causal relation between them; whereas missing arrows imply conditional independence.

#### 4.1.2. Reasoning underlying cause–effect relations

Causal maps identify individuals' perceptions of cause–effect relationships between variables based on language rather than the reasoning processes (Carley and Palmquist, 1992). Studies in managerial cognition indicate that individuals reason by accumulating possibly significant pieces of infor-

<sup>1</sup> If we start with an arbitrary probability distribution, it is not always possible to find a Bayesian network representation that is a perfect map (Pearl, 1988). Our starting point is an expert's causal model and we are assuming that we can always find a Bayesian network representation that is a perfect map.

mation and organizing them in relation to each other so as to be able to combine them into a conclusion and decision (Jaques and Clement, 1996). Individuals use such reasoning processes to put information together as a cause–effect series of events leading to predicted future courses of events. These reasoning processes are important in decision making and in making inferences about future decision outcomes.

Literature on logic suggests that individuals perceive cause–effect relationships based on two types of reasoning: *deductive* and *abductive* (Charniak and McDermott, 1985; Winston, 1984). A reasoning process is called deductive when we reason *from* causes *to* effects, i.e., in the direction of causation. For example, in the medical domain, risk factors (e.g., smoking) are regarded as causes, and the diseases (e.g., lung cancer) as effects. When a physician, confronted with a patient who has been a smoker, reasons that the patient is at risk for lung cancer, (s)he is reasoning deductively.

A reasoning process is called abductive when we reason from effects to causes, i.e., in the direction opposite to causation. For example, diseases (e.g., lung cancer) are regarded as causes of symptoms (e.g., positive X-ray). When a physician, after observing a patient’s positive X-ray result, concludes that the patient is probably suffering from lung cancer, (s)he is reasoning abductively.

The difference between deductive and abductive reasoning underlying causal statements and their effect on representation of causal linkages are illustrated in Fig. 4. Causal statement 1 involves the use of logical deduction and the reasoning is in the direction of causation. This is correctly reflected in

the arc from Product Risk to Sales Projection. Causal statement 2 involves abductive reasoning. Since Product Risk is a latent variable, and Sales Projection is one of the observable measures of Product Risk, the decision maker is making inference about the latent variable Product Risk based on the observable variable Sales Projection. This does not imply that Sales Projection causes Product Risk. The reasoning in this causal statement is in the direction opposite of causation. Causal statements involving abductive reasoning are misrepresented in a causal map by an arc from effect to cause. Such misrepresentation can also lead to redundant circular relations between variables in the causal map. For instance, both the arrows in Fig. 4 may be represented in a causal map creating a loop.

A distinction between deductive and abductive reasoning behind the causal linkages is essential to establish accurate directions of linkages in causal maps. The emphasis in deriving causal maps should be on the causal theory underlying the causal statements rather than the language used.

#### 4.1.3. Distinguishing between direct and indirect relationships

The procedure for deriving causal maps does not provide for a distinction between ‘direct’ and ‘indirect’ relationships between concepts (Eden et al., 1992; Laukkanen, 1996). This distinction is important to identify conditional independencies in the causal maps. Fig. 5 depicts how a lack of distinction between direct and indirect relationship affects conditional independence assumptions in a

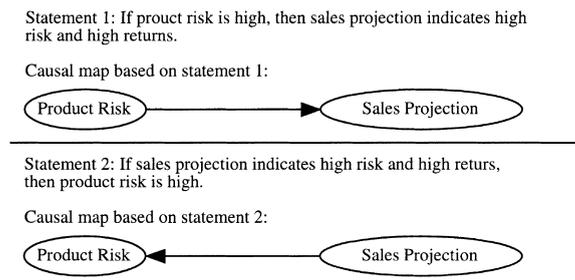


Fig. 4. Distinguishing between deductive and abductive reasoning.

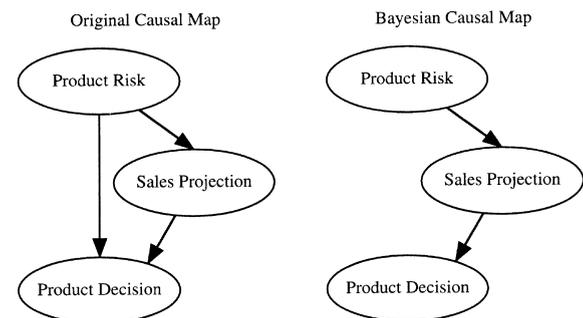


Fig. 5. Distinguishing between direct and indirect relations.

causal map. In Fig. 5, both Product Risk and Sales Projection directly influence Product Decision. In the modified Bayesian causal map, there is no linkage between product risk and product decision implying that product risk impacts product decision strictly through Sales Projection. If we have complete information on Sales Projection, any additional information on Product Risk would be irrelevant in making inferences about Product Decision.

A clear distinction between direct and indirect cause–effect relations is important for three reasons. First, it helps us understand the nature of relations between variables. It tells us whether the effect of a variable on another is completely modeled by the effect of the first on a third mediating variable (which in turn is a cause of the second). Second, if Product Risk (in Fig. 5) affects Product Decision only through Sales Projection Figures, then an arrow from product risk to product decision is redundant and increases the complexity of the representation. Finally, distinction between direct and indirect cause–effect relations allows incorporation of conditional independencies in causal maps. As we have seen earlier, conditional independencies are critical in making inferences on the variables in large causal maps.

#### 4.1.4. Eliminating circular relations

Causal maps are directed graphs and are characterized by a hierarchical (or acyclic) structure. However, circular relations or causal loops destroy the hierarchical form of a graph. Circular relations in the causal maps violate the acyclic graphical structure required in a Bayesian network. It is therefore essential to eliminate circular relations to make causal maps compatible with Bayesian networks. Causal loops can exist for two reasons (Eden et al., 1992; Huff, 1990; Laukkanen, 1996). First, they may be coding mistakes that need to be corrected. Second, they may represent dynamic relations between variables across multiple time frames.

Coding mistakes can be rectified by clarifying causal linkages between variables in terms of deductive versus abductive reasoning or direct versus indirect linkage; issues already discussed

in previous paragraphs. However, if causal loops exist despite these clarifications, then they can be eliminated by aggregating the variables into a single variable. In causal maps, causal assertions of individuals are aggregated (or clustered) into broader concepts. All variables or nodes in circular relations are of the same hierarchical status and so if aggregated to a single node, the general form of the causal map can still be hierarchical.

In addition to coding mistakes, feedback loops may indicate dynamic relations between variables overtime. In such cases, part of the linkages in the loop pertains to a current time frame and some linkages pertain to a future time frame. In such cases, deaggregating the variables into two time frames can often solve the problem of circularity. For example, Fig. 6 shows a reciprocal causal relation between Market Leadership and Rate of Product Launch and reasoning underlying this circular relation. Arrow  $t_1$  implies that companies having market leadership tend to introduce more products that are new. Arrow  $t_2$  implies that introduction of new products affects the Market Leadership of the firm at some future point in time. The circular relation has resulted from aggregation of the variable Market Leadership across two time frames:  $t_1$  and  $t_2$ . After de-aggregating Market Leadership into two time frames, we get an acyclic relation between the three variables. To make the causal map acyclic, we arbitrarily retain one of the two relations and exclude the other from the causal map. An acyclic structure of the causal map is essential to the inference process and to make causal maps compatible with Bayesian networks. Bayesian networks are unable to represent reciprocal causal relations.

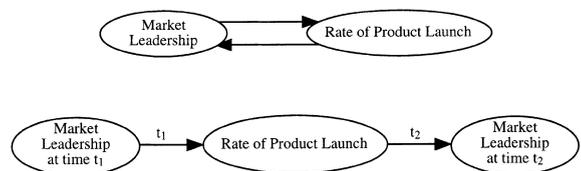


Fig. 6. Deaggregating variables over time.

## 4.2. Probabilistic stage

In this stage, the numerical parameters of the Bayesian causal map are assessed. Also, the causal Bayesian map is validated using validation procedures.

### 4.2.1. Assessment of numerical parameters

The parameters of the causal map consist of the strength of the linkages between the different variables. However, in the context of inference, we are interested in the uncertainty associated with every variable in the map and the interactive effects of multiple causal variables on effect variables. The Bayesian network approach provides tools to address both of these issues in defining the parameters of the map. Major issues in the derivation of the parameters of the causal map are discussed in the following sections.

One of the most challenging problems in generating a causal map is specifying the strength associated with the causal connections in the map. Traditional analytical methods to represent the parameters of the causal maps are based on social network methods (Knoke and Kuklinski, 1982). The parameters of a causal map constitute the strength of each pairwise linkage in the map and are measured in terms of the frequency of the linkage in the narrative. This measure has been criticized by researchers for several reasons (Huff, 1990; Laukkanen, 1996). First, it does not explicitly incorporate an individual's perception of the strength associated with a linkage in the causal map. A number of factors such as the type of interview (structured versus unstructured), length of the narrative (either interview or archival data), etc., can affect the frequency with which concepts and linkages are mentioned in the map. Second, the strength of the causal connection is based on pairwise comparison of each causal variable independently on the effect variable. It does not capture the interactive effects of causal variables on the effect variables. Finally, the uncertainty associated with each variable in the map is not represented in the frequency measure. All variables are assumed to have the same level of uncertainty.

Treatment of uncertainty is necessary in causal maps because there are different sources of uncertainty in a decision problem. The observations may be uncertain, the information about variables in the map may be incomplete, the relations in the domain represented by the map may be of a non-deterministic type, or the variables involved may be vague. One common way of capturing uncertainty of the variables in a Bayesian network is to measure a person's 'degree of belief' for that variable conditional on the states of its parents. This uncertainty associated with the variables in a decision model is sensitive to the context in which the certainties have been established. The process of measuring degrees of belief is commonly referred to as *probability assessment* or *probability encoding procedure*.

Many different probability encoding techniques are available (for a detailed review see (Spetzler and Staël von Holstein, 1975)) wherein a subject responds to a set of questions either directly by providing numbers or indirectly by choosing between simple alternatives or bets. The choice of response mode (direct or indirect) as well as the choice of a method within each mode depends on the preferences of the subject. Spetzler and Staël von Holstein (1975) describe three direct response encoding methods – cumulative probability, fractiles and verbal encoding – to elicit probabilities. In the cumulative probability method, the subject is asked to assign the cumulative probability associated with a variable conditioned on the states of its parent variables. The probability response can be expressed either as an absolute number (0.30), as a discrete scale ("three on a scale from zero to ten"), or as a fraction using a discrete scale ("three in ten"). Verbal encoding uses verbal descriptions to characterize events in the first phase of the encoding procedure. The descriptors used are those to which the subject is accustomed to such as "high", "medium" or "low". The quantitative interpretation of the descriptors is then encoded in a second phase. The form chosen to express the probability (absolute number, percentage, fraction or verbal) should be the one most familiar to the subject.

When a variable has many parents, the number of probability assessments can be reduced by

assessing the nature of the relationship between the variable and its parents such as noisy-OR, noisy-AND, etc. (Henrion, 1989; Pradhan et al., 1994). Once the parameters of the causal map are identified, propagation algorithms can be used to make inferences about the variables in the causal maps.

#### 4.2.2. Validation

We propose validation procedures for both qualitative and quantitative stages of constructing Bayesian causal maps. In the qualitative stage, we extend the reliability and validity techniques suggested by causal mapping studies (Axelrod, 1976; Carley, 1997; Eden et al., 1992; Huff, 1990) to the modification procedures proposed in our study. A typical method used to establish reliability of the cognitive mapping procedure is to achieve a consensus among multiple raters for all the stages of the cognitive mapping procedure. Accordingly, the modifications to the original causal map (described in Section 4.1) can be made through a consensus between multiple raters. For example, modifications agreed to by two or more raters can be retained, whereas those for which agreement cannot be reached can be thrown out. Such inter-rater reliability also reduces the possibility that researchers' own inter-

pretations have contaminated the modification process.

Validation of Bayesian causal maps involve establishing that the modifications made to the original causal map measure what they are intended to measure. This validation has to be established at the qualitative level as well as the quantitative level. Typical practice suggests that researchers check the interpretation of variables in a causal map with the subjects involved. This can be extended to the modification procedure in the qualitative phase wherein modifications can be made to the original causal maps after consultation with the subjects involved.

At the quantitative level, sensitivity analysis approach from operations research can be used to validate the map (Churchman et al., 1957; Howard and Matheson, 1981). Sensitivity analysis consists of examining the posterior marginals of decision variables under different scenarios and then corroborating these marginals with the domain expert. The sensitivity analysis procedure is illustrated in the next section in the context of a case study.

Bayesian causal maps combine the strengths of causal maps and Bayesian networks and reduce the limitations of either. A comparison of the three network models is shown in Table 1.

Table 1  
A comparison of causal maps, Bayesian networks and Bayesian causal maps

Features	Causal map	Bayesian network	Bayesian causal map
<i>Graphical structure:</i>			
1. Type of network model	Dependence map	Independence Map	Perfect map
2. Type of graph	Directed graph	Directed acyclic graph	Directed acyclic graph
3. Definition of relations	Based on causality	Based on conditional independence	Based on causality and conditional independence
<i>Model parameters:</i>			
4. Uncertainty of variables	Not represented	Represented	Represented
5. Background noise or unmodeled factors	Not represented	Represented	Represented
6. Interactive effects of multiple causal variables on effect variables	Not represented	Represented through conditional probabilities	Represented through conditional probabilities
7. Validation	Qualitative	Quantitative using sensitivity analysis	Qualitative and quantitative
<i>Inference procedure:</i>			
8. Inference procedure	Qualitative based on the graphical structure	Quantitative based on evidence propagation	Qualitative and quantitative

## 5. A case study: Product development decision

This section describes a case study of a construction of a Bayesian causal map. The main goal of the case study is to provide a test bed for the construction of a Bayesian causal map using the technique described in the previous section. The case study is also useful for illustration purposes. We describe how starting from a causal map, we construct the qualitative structure of a Bayesian causal map. We show how additional information can be collected from a subject to address the modeling issues discussed in Section 4.1 as well as to derive the numerical parameters of the Bayesian causal map. In addition, we describe how Bayesian network software can be used to draw probabilistic inferences in a Bayesian causal map.

*Decision context.* We used a modified version of the *New Product Decision* case (Clemen, 1996) from decision analysis literature as a decision context. We chose this case because it was sufficiently unstructured to allow an expert to develop his/her own framework in diagnosis, analyses, and recommendations of decision options. Also, the decision alternatives, outcomes, and situational factors all involved uncertainty.

The decision scenario presented in the case is as follows. A marketing executive of a computer-manufacturing firm is faced with a decision of introducing a new computer model. There are three different options. The first option is to continue with the existing product line that is essentially a risk-free proposition with stable expected profits. The second option, which is riskier than the first, involves introducing a modified version of the existing model. This option involves significant initial investment and very little information is available on sales predictions. However, a working prototype of the modified product has already been developed and pricing decisions have been made. In other words, it is ready to be produced and marketed. Moreover, this product, if successfully marketed, has the potential of yielding profits much higher than the first product. The third option, which is riskier than the second, consists of introducing a product that has not yet been fully developed. The prototype developed has some problems, and the engineering division has indi-

cated its confidence in solving the problems, but there is a slight chance that devising a workable solution may take a much longer than predicted resulting in a significant delay in introducing this modified product. The best case profit scenario for the third option is much higher than the second.

*Subject.* The subject was a marketing professor whose prior work experience and current line of research was relevant to the decision context.

*Data collection.* The subject was given the decision case a week before the interview. An intensive interview (two to three hours) was conducted with the subject to elicit his/her analysis of the case and decision recommendations. The subject was asked to identify factors that he/she would consider as the marketing executive responsible for the new product decision. The subject was encouraged to identify factors not identified in the decision scenario that, in his/her opinion, might be relevant to the decision. The narrative yielded by this interview was used to construct the original causal map shown in Fig. 7 below using a textual analysis method (for details see Carley and Palmquist, 1992; Eden et al., 1979; Huff, 1990; Laukkanen, 1996). Follow-up interviews were conducted with the subject to modify the structure of the causal map to construct the qualitative structure of the Bayesian causal map. The next few paragraphs discuss this procedure in detail.

*Causal map.* The original causal map shown in Fig. 7 describes the subject's causal perceptions of the decision problem in the new product decision case. There are 18 variables in the map representing industry factors, firm factors, product characteristics, and decision options. The three major industry factors are Market Dynamics, Product Life Cycle, and Market Share Distribution. Firm factors represent the prior market experience of the firm and long term strategies of the firm and include Competitive Strategy, Brand Loyalty, Market Leadership, Cross-Functional Teams, R&D, Rate of Product Launch and Prior Presence (in the market). The product characteristics are represented by Sales Uncertainty, Production Delay, Initial Investment, Product Risk, and Sales Projection. The three decisions the manager has to make in this specific decision context are Willingness to Take Risk, Product Pricing, and Product



dependence between the two variables. This resulted in five additional links in the modified map. First is the link between Market Dynamics and Product Life Cycle that did not exist in the original causal map. Second is the link between Market Dynamics and Sales Uncertainty. These variables were shown as conditionally independent in the original map. Similarly, the modified map shows links between Market Share Distribution and Sales Uncertainty, Market Share Distribution and Sales Uncertainty, and Market Share Distribution and Prior Presence that did not exist in the original causal map.

3. *Deductive reasoning.* The subject was instructed that an arrow should be directed from cause to effect only (deductive reasoning). The original map shows a link from Sales Projection to Product Risk. However, as shown in Fig. 4, the direction of the arrow is based on abductive reasoning from the observable measure (Sales Projection) to the latent cause (Product Risk). The direction of this arrow changed in the modified causal map from Product Risk to Sales Projection.

4. *Similar time frames.* The subject was instructed that all variables should pertain to a specific time frame  $t_1$ . We defined  $t_1$  as the period until the final decision of marketing the new product was reached (to eliminate circular relations due to relations pertaining to different periods). This resulted in the elimination of a reciprocal relationship between Market Leadership and Rate of Product Launch shown in Fig. 6.

At the end of the first interview, the subject was asked to provide a brief rationale or explanation for every link added or deleted in the original causal map. This explanation provided a validity check for modifications made to the map. Based on the suggestions made and the explanation provided by the subject, two raters (one researcher and another marketing expert) modified the original causal map through consensus. Modifications agreed to by both the raters were made to the original causal map, and those not agreed to were excluded.

A second follow-up interview was conducted to validate the modification process. Accordingly, the subject was provided with a modified causal map (based on the information collected in the first

follow-up interview) to check if the modified map correctly reflected the modifications made to the original map. The final modifications were made to the map through close consultations with the subject. The resultant structure of the Bayesian causal map is shown in Fig. 8.

In the probabilistic stage, the parameters of the model were assessed. The parameters of a Bayesian causal map consist of marginal probabilities and conditional probabilities. To assess the marginal probabilities, the expert was asked to provide the following information.

5. To rate the marginal and conditional probabilities on a discrete scale (0 to 10); and

6. To identify the type of interactive effects of multiple causal variables on effect variables. For example, whether each causal variable affects the effect variable independently (noisy-OR model), or whether each causal variable affect the effect variables through interactions of two or more variables (noisy-AND), or some combination of the two (Henrion, 1989; Pradhan et al., 1994).

*Making inferences.* We used Netica [www.norsys.com] to make probabilistic inferences using both sum propagation and max propagation.

*Sum propagation.* The sum propagation computes the marginal probabilities of all the model variables and updates the marginals with all additional evidence received about other variables.

In our case study, we can evaluate each product decision option under different scenarios. The scenarios were defined in consultation with three marketing experts (including the subject), and they represent situations in which there are unambiguous prescriptions for product decisions in the marketing literature. We illustrate how predictions can be made about our subject's perceptions of product option 1 (risk free proposition) in terms of 'go' or 'no go' under different information conditions. The same procedure can be used to evaluate product options 2 and 3. Product option 1 is to continue with the current product line and is a risk-free proposition with stable expected profits. This option can be represented by specifying the states of two variables in the Bayesian causal map: Initial Investment = low, and Production Delay = no delay. Since the product is similar to the one already being marketed, the initial cost of

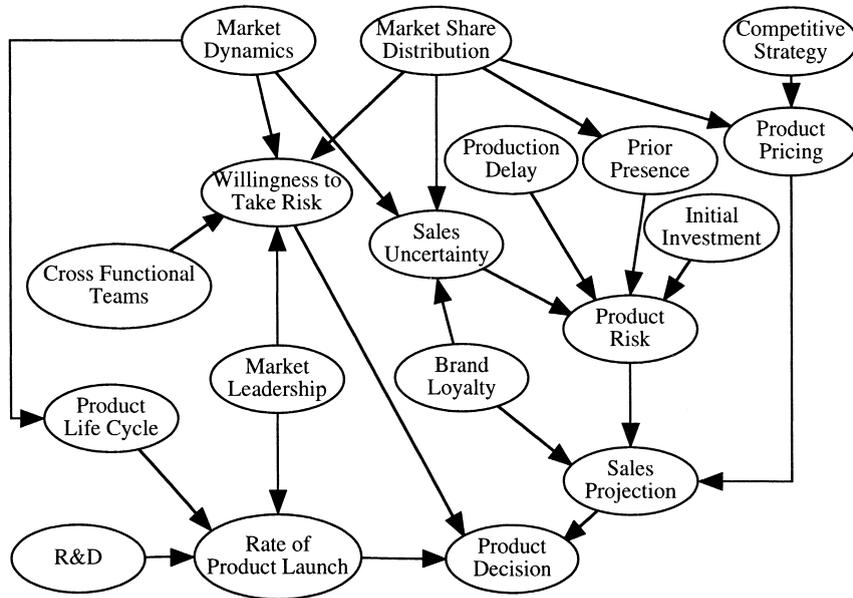


Fig. 8. The final bayesian causal map of the product development decision.

development and commercialization is minimal. Similarly, the chance of delay in production due to development problems is also very low. Table 2 shows the prior marginal probabilities of the three

variables directly affecting product decision when all other variables (except Initial Investment and Production Delay) in the map are unknown. The marginal probabilities of these variables are as

Table 2  
Prior and posterior marginal probabilities under two different scenarios

Variables	Prior marginals	Posterior marginals in scenario 1 <sup>a</sup>	Posterior marginals in scenario 2 <sup>b</sup>
1. Rate of Product Launch:			
High	0.70	0.15	0.80
Low	0.30	0.85	0.20
2. Willingness to Take Risk:			
High	0.66	0.13	0.83
Low	0.34	0.87	0.17
3. Sales Projection:			
High sale high risk	0.49	0.32	0.41
Low sale low risk	0.51	0.68	0.59
4. Product Option 1 <sup>c</sup> :			
Go	0.51	0.63	0.44
No go	0.49	0.37	0.56

<sup>a</sup> Scenario 1: Market Share Distribution=low share, Competitive Strategy=cost-efficiency, Market Leadership=follower and R&D=low.

<sup>b</sup> Scenario 2: Market Share Distribution=high share, Competitive Strategy=cost-efficiency, Market Leadership=leader and R&D=high.

<sup>c</sup> Product Option 1: Production Delay=no delay, and Initial Investment=low.

follows. Rate of Product Launch: high = 0.70, low = 0.30; Willingness to Take Risk: high = 0.66 and low = 0.34; Sales Projection: high sale high risk = 0.49, low sale low risk = 0.51; Product Decision: go = 0.51, no go = 0.49. Since the industry specified in the decision case is the computer industry in which market dynamics are high, the priors show that any company in the industry would be willing to take risk and have a high rate of product launch. However, if we have additional information about specific firm characteristics, then our inferences about Product Decision may change.

We specify two different scenarios and show how our inferences about product decision option one change depending on the firm factors. In the first scenario, we consider a struggling firm with a low market share that follows a price-based strategy. Accordingly we specify the states of the following four firm variables in the map: Market Share Distribution = low, Competitive Strategy = cost-efficiency, Market Leadership = follower, and R&D = low. Based on this information, we propagate the information to compute the posterior marginals of variables of interest. A comparison of prior and posterior marginals of Rate of Product Launch, Willingness to Take Risk, Sales Projection and Product Decision is shown in Table 2. When additional information is received about firm factors, the posterior probability of Rate of Product Launch = low increases from 0.30 to 0.85, that of Willingness to Take Risk = low increases from 0.37 to 0.87 and that of Sales Projection = low sale low risk increases from 0.51 to 0.68. These posterior marginals change our inference about the state of Product Decision. The posterior probability of 'go' for product option 1 is 0.63 in comparison to a prior of 0.51. Under the conditions described in scenario 1, our marketing expert is likely to select product option 1 since the firm factors make him/her risk averse.

In the second scenario, we consider a market-leading firm with a high market share and that follows a technology-based strategy. Accordingly, the states of the four firm variables in the map are specified as follows: Market Share Distribution = high share, Competitive Strategy = innovation, Market Leadership = leader and R&D = yes.

As shown in Table 2, the posterior marginal of Rate of Product Launch = high increases from 0.70 to 0.80, that of Willingness to Take Risk = high increases from 0.66 to 0.83, and that of Sales Projection = low sale low risk increases from 0.51 to 0.59. The posterior marginal probability of Product Decision = no go increases from 0.50 to 0.56. This implies that in scenario two, our marketing expert is more likely to reject Product Option 1.

A relevant question is which product option will our marketing expert choose in the second scenario. To find an answer to this question, we evaluated both product options two and three following the procedure described above. The results show that the highest probability of 'go' in scenario two is for Product Option 3. Product Option 3 can be described in the map by specifying Initial Investment = high and Delay in Production = high. This increases the probability of Sales Projection = high sale high risk to 0.56, whereas the probabilities of Willingness to Take Risk and Rate of Product Launch remain the same. The posterior marginal probability of Product Option 3 = go is 0.62. This implies that in scenario two, our subject is more likely to go for a high sale high risk proposition such as Product Option 3 since he/she has an organizational infrastructure which allows him/her to take risk in order to go for a new product with potentially very high sales.

*Max propagation.* The max propagation computes the most likely configuration of the states of model variables under different scenarios. This is useful in getting a macro picture of the situation. We illustrate max propagation for Product Option 1 (risk free proposition with stable profits) under the same two scenarios described in the section on sum propagation. Table 3 shows the most likely configuration of the variables when all variables (except Initial Investment = low and Production Delay = no delay) are unknown. However, when we specify Scenario 1, the most likely configuration of the variables changes. The most likely state of Rate of Product Launch changes from high to low, that of prior presence in the market changes from yes to no, that of Sales Uncertainty changes from high to low. The most likely states of all the three decision variables change: Willingness to

Table 3  
Most likely configuration of variables under two different scenarios

Variables	Prior model	Scenario 1	Scenario 2
<i>Industry variables:</i>			
1. Market Dynamics	High	High	High
2. Product Life Cycle	Short	Short	Short
3. Market Share Distribution	High share	Low share	High share
<i>Firm variables:</i>			
4. Competitive Strategy	Technology-based	Price based	Technology-based
5. Brand Loyalty	High	High	High
6. Market Leadership	Leader	Follower	Leader
7. Cross Functional Teams	Yes	No	Yes
8. R&D	High	Low	High
9. Rate of Product Launch	High	Low	High
10. Prior Presence in the Market	Yes	No	Yes
<i>Product variables:</i>			
11. Sales Uncertainty	High	Low	High
12. Production Delay	No delay	No delay	No delay
13. Initial Investment	Low	Low	Low
14. Product Risk	Low	Low	Low
15. Sales Projection	Low sale low risk	Low sale low risk	Low sale low risk
<i>Decision variables:</i>			
16. Willingness to Take Risk	High	Low	High
17. Product Pricing	High	Low	High
18. Product Decision	No go	Go	No go

Take Risk changes from high to low, Product Pricing from high to low and Product Decision Option one changes from no go to go. This change in the states of the firm variables and product characteristics help us explain why the state of Product Option 1 changes from no go to go. When Scenario 2 is specified, the most likely configuration of the variables changes and the Product Option 1 changes to no go. The change in the configuration is shown in Table 3. The pattern of change in the configuration of the map helps us understand how and why the overall decision changes when firm characteristics change.

Making inferences using sum and max propagation is useful in predicting decision outcomes as well as in evaluating and analyzing decision factors that may have affected specific decision outcomes.

## 6. Summary and conclusions

The main goal of this paper is to propose a formal procedure for making inferences in causal

maps using the Bayesian network approach. Accordingly, we describe a new graphical structure called Bayesian causal maps. Bayesian causal maps combine the strengths of causal maps and Bayesian networks and reduce the limitations of both. Using concepts from the literature on causal modeling and logic, Bayesian causal maps clarify the cause–effect relations depicted in the causal maps. They depict dependence between variables based on causal mapping approach (D-map) as well as a lack of dependence between variables based on the Bayesian network approach (I-map). A Bayesian causal map is therefore a perfect map. Bayesian causal maps consider the reasoning (deductive versus abductive) underlying the cause–effect relations perceived by individuals. This strengthens the validity of the direction of causal relations represented in the map. Bayesian causal maps provide a framework for representing the uncertainty of variables in the map as well as the effect of variables not modeled in the map. Using evidence propagation algorithms, Bayesian causal maps allow us to make inferences about the vari-

ables in the map. Finally, using a case study, we have illustrated how Bayesian causal maps can be constructed starting from a causal map, and how it can be used to make inferences.

There are some interesting implications of our study. First, constructing Bayesian networks has been a big bottleneck in knowledge engineering (Henrion, 1989; Pradhan et al., 1994; Laskey and Mahoney, 1997). If our procedure is tested and refined, it could be used as a basis for constructing Bayesian networks based on an expert's knowledge. Second, most of the emphasis in causal maps has been in its use for knowledge representation. This study enables decision-makers to use causal maps for decision making. Influence diagrams proposed by Howard and Matheson (1981) use Bayesian network models of uncertainty in addition to decision nodes, utility functions, and information constraints. Thus, Bayesian causal maps can be used for normative decision making using the framework of influence diagrams.

### Acknowledgements

We are grateful to Rohini Ahluwalia, V.K. Narayanan, Anurag Pant, and Surendra Singh for their comments and discussions. The research reported in this paper has been supported by two grants from the Kansas University Business School PhD Summer Research Fund, and by a grant from the Kansas University General Research Fund to the second author. The manuscript has benefited from the comments of an anonymous reviewer.

### References

- Axelrod, R., 1976. *Structure of Decision: The Cognitive Maps of Political Elites*. Princeton University press, Princeton, NJ.
- Bougon, M., 1983. Uncovering cognitive maps: The 'self-Q' technique. In: Morgan, G. (Ed.), *Beyond Method: A Study of Organizational Research Strategies*. Sage, Beverley Hills, CA.
- Bougon, M., 1992. Congregate cognitive maps: A unified dynamic theory of organization and strategy. *Journal of Management Studies* 29, 369–389.
- Bougon, M.K., Weick, K.E., Binkhorst, D., 1977. Cognition in organizations: An analysis of the Utrecht Jazz Orchestra. *Administrative Science Quarterly* 22, 606–639.
- Carley, K., 1997. Extracting team mental models through textual analysis. *Journal of Organizational Behaviour* 18, 533–558.
- Carley, K., Palmquist, M., 1992. Extracting, representing and analyzing mental models. *Social Forces* 70 (3), 601–636.
- Charniak, E., McDermott, D., 1985. *Introduction to Artificial Intelligence*. Addison-Wesley, Reading, MA.
- Chen, K., Mathes, J.C., Jarboe, K., Wolfe, J., 1979. Value oriented social decision analysis: Enhancing mutual understanding to resolve public policy issues. *IEEE Transactions on Systems, Man and Cybernetics* 9, 567–580.
- Churchman, C.W., Ackoff, R.L., Arnoff, E.L., 1957. *Introduction to Operations Research*. Wiley, New York.
- Clemen, R.T., 1996. *Making Hard Decisions – An Introduction to Decision Analysis*. Duxbury Press, Belmont, CA.
- Darwiche, A., Goldszmidt, M., 1994. On the relation between kappa calculus and probabilistic reasoning. In: Mantaras, R.L., Poole, D. (Eds.), *Uncertainty in Artificial Intelligence: Proceedings of the 10th Conference*. Morgan Kaufmann, San Francisco, CA, pp. 145–153.
- Dawid, A.P., 1979. Conditional independence in statistical theory (with discussion). *Journal of the Royal Statistical Society Series B* 41 (1), 1–31.
- Eden, C., 1991. Working on problems using cognitive mapping. In: Littlechild, S.C., Shutler, M. (Eds.), *Operations Research Management*. Prentice-Hall, London, pp. 236–259.
- Eden, C., 1992. On the nature of cognitive maps. *Journal of Management Studies* 29, 261–266.
- Eden, C., Jones, S., Sims, D., 1979. *Thinking in Organizations*. Macmillan, London.
- Eden, C., Jones, S., 1980. Publish or perish?: A case study. *Journal of Operational Research Society* 31, 131–139.
- Eden, C.F., Ackermann, F., Cropper, S., 1992. The analysis of cause maps. *Journal of Management Studies* 29 (3), 309–323.
- Eden, C., Ackermann, F., 1993. Evaluating strategy: Its role within the context of strategic control. *Journal of Operational Research Society* 44, 853–865.
- Fiol, M., 1990. Explaining strategic alliance in the chemical industry. In: Huff, A.S. (Ed.), *Mapping Strategic Thought*. Wiley, Chichester, UK, pp. 227–249.
- Fiol, M., Huff, A.S., 1992. Maps for managers: Where are we? Where do we go from here? *Journal of Management Studies* 29, 269–285.
- Forrester, J.W., 1961. *Industrial Dynamics*. MIT Press, Cambridge, MA.
- Heckerman, D., 1996. Bayesian networks for data mining. *Data Mining and Knowledge Discovery* 1, 79–119.
- Henrion, M., 1989. Some practical issues in constructing belief networks. In: Kanal, L.N., Levitt, T.S., Lemmer, J.F. (Eds.), *Uncertainty in Artificial Intelligence*, vol. 3. North-Holland, Amsterdam, pp. 161–173.
- Henrion, M., Provan, G., Favero, B.D., Sanders, G., 1994. An experimental comparison of numerical and qualitative

- probabilistic reasoning. In: Mantaras, R.L., Poole, D. (Eds.), *Uncertainty in Artificial Intelligence: Proceedings of the 10th Conference*. Morgan Kaufmann, San Francisco, CA, pp. 319–326.
- Howard, R., Matheson, J., 1981. Influence diagrams. In: Howard, R., Matheson, J. (Eds.), *Readings on the Principles and Applications of Decision Analysis*, vol. 2. Strategic Decisions Group, Menlo Park, CA, pp. 721–762.
- Huff, A.S., 1990. *Mapping Strategic Thought*, Wiley, Chichester, UK.
- Jaques, E., Clement, S.C., 1996. *Executive leadership – A practical guide to managing complexity*. Cason, Hall and Company.
- Jensen, F.V., Lauritzen, S.L., Olesen, K.G., 1990. Bayesian updating in causal probabilistic networks by local computation. *Computational Statistics Quarterly* 4, 269–282.
- Klein, J.H., Cooper, D.F., 1982. Cognitive maps of decision workers in complex game. *Journal of Operational Research Society* 33, 63–71.
- Knoke, K., Kuklinski, J.H., 1982. *Network Analysis*. Sage, Beverley Hills, CA.
- Laskey, K.B., Mahoney, S.M., 1997. Network fragments: Representing knowledge or constructing probabilistic models. In: Geiger, D., Shenoy, P.P. (Eds.), *Uncertainty in Artificial Intelligence: Proceedings of the 13th Conference*. Morgan Kaufmann, San Francisco, CA, pp. 334–341.
- Laukkanen, M., 1996. Comparative cause mapping of organizational cognition. In: Meindl, J.R., Stubbart C., Porac J.F. (Eds.), *Cognition Within and Between Organizations*. Sage, Beverley Hills, CA, pp. 1–45.
- Lauritzen, S.L., Spiegelhalter, D.J., 1988. Local computations with probabilities on graphical structures and their application to expert systems (with discussion). *Journal of Royal Statistical Society Series B* 50 (2), 157–224.
- Lauritzen, S.L., Dawid, A.P., Larsen, B.N., Leimer, H.-G., 1990. Independence properties of directed Markov fields. *Networks* 20 (5), 491–505.
- Pearl, J., 1986. Fusion, propagation, and structuring in belief networks. *Artificial Intelligence* 29, 241–288.
- Pearl, J., 1988. *Probabilistic Reasoning in Intelligent Systems: Networks of Plausible Inference*. Morgan Kaufmann, San Mateo, CA.
- Pradhan, M., Provan, G., Middleton B., Henrion, M., 1994. Knowledge engineering for large belief networks. In: Lopez de Mantaras, R., Poole, D. (Eds.), *Uncertainty in Artificial Intelligence: Proceedings of the 10th Conference*. Morgan Kaufmann, San Francisco, CA, pp. 484–490.
- Ross, L.L., Hall, R.I., 1980. Influence diagrams and organizational power. *Administrative Science Quarterly* 25, 57–71.
- Shachter, R.D., Heckerman, D.E., 1987. A backwards view for assessment. *AI Magazine* 8 (3), 55–61.
- Shenoy, P.P., Shafer, G., 1990. Axioms for probability and belief-function propagation. In: Shachter, R.D., Levitt, T.S., Lemmer, J.F., Kanal, L.N. (Eds.), *Uncertainty in Artificial Intelligence*, vol. 4. North-Holland, Amsterdam, pp. 169–198. Reprinted in: Shafer, G., Pearl, J. (Eds.), *Readings in Uncertain Reasoning*. Morgan Kaufmann, San Mateo, CA, 1990, pp. 575–610.
- Speigelhalter, D.J., Dawid, A.P., Lauritzen, S.L., Cowell, R.G., 1993. Bayesian analysis in expert systems. *Statistical Science* 8 (3), 219–283.
- Spetzler, C.S., Staël von Holstein, C.S., 1975. Probability encoding in decision analysis. *Management Science* 22 (3), 340–358.
- Swan, J.A., 1995. Exploring knowledge and cognitions in decisions about technological innovation: Mapping managerial cognitions. *Human Relations* 48 (11), 1241–1270.
- Wang, S., 1996. A dynamic perspective of differences between cognitive maps. *Journal of the Operational Research Society* 47, 538–549.
- Wolstenholme, E.F., Coyle, R.G., 1983. The development of system dynamics as a methodology for systems description and qualitative analysis. *Journal of Operational Research Society* 34, 569–681.
- Winston, P.H., 1984. *Artificial Intelligence*. Addison-Wesley, Reading, MA.