

A CAUSAL MAPPING APPROACH TO CONSTRUCTING BAYESIAN NETWORKS[†]

Sucheta Nadkarni

College of Business Administration, University of Nebraska-Lincoln

and

Prakash P. Shenoy

School of Business, University of Kansas

Abstract

This paper describes a systematic procedure for constructing Bayesian networks from domain knowledge of experts using the causal mapping approach. We outline how causal knowledge of experts can be represented as causal maps, and how the graphical structure of causal maps can be modified to construct Bayes nets. Probability encoding techniques can be used to assess the numerical parameters of the resulting Bayes nets. We illustrate the construction of a Bayes net starting from a causal map of a systems analyst in the context of an information technology application outsourcing decision.

Key Words: Causal maps, cognitive maps, Bayesian networks, Bayesian causal maps

[†]Please send all correspondence to Sucheta Nadkarni, Department of Management, College of Business Administration, University of Nebraska-Lincoln, Lincoln, NE 68588-0491, <snadkarn@unlnotes.unl.edu>.

Biographies

Sucheta Nadkarni

Sucheta Nadkarni is an Assistant Professor in the Department of Management at the University of Nebraska-Lincoln. She holds a Ph.D. in Business from the University of Kansas. Her articles have been published in journals such as *MIS Quarterly* and *European Journal of Operations Research*. Her research interests include the role of managerial cognition in strategic management, causal mapping and management of technology.

Prakash P. Shenoy

Prakash P. Shenoy is the Ronald G. Harper Distinguished Professor of Artificial Intelligence in the School of Business at the University of Kansas, Lawrence, Kansas. He received his PhD in Operations Research from Cornell University in 1977. He has published extensively in journals such as *Operations Research*, *Management Science*, *Artificial Intelligence*, and the *International Journal of Approximate Reasoning*. His research interests are in the areas of uncertain reasoning and decision analysis. He is the inventor of valuation-based systems, an abstract framework for knowledge representation and inference that unifies disparate subjects such as Bayesian probabilities, Dempster-Shafer theory of belief functions, Spohn's theory of epistemic beliefs, Zadeh's possibility theory, propositional logic, optimization using dynamic programming, and solving systems of equations.

1 Introduction

A Bayesian network [BN] is a graphical model that encodes relationships among variables of interest. When used in conjunction with statistical techniques, a BN has several advantages for data analysis [33]. One, it readily handles situations where some data entries are missing. Two, it can be used to model causal relationships, and hence can be used to gain understanding about a problem domain and to predict the consequences of intervention. Three, because the model has both causal and probabilistic semantics, it is an ideal representation for combining prior knowledge (which often comes in causal form) and data. BNs are especially useful in modeling uncertainty in a domain. BNs have been applied particularly to problems, which require diagnosis of problems from a variety of input data. A few examples of BN applications include medical diagnostic systems, real-time weapons scheduling, and generator monitoring expert system and troubleshooting.

Two different approaches have been used to construct Bayesian networks—data-based approach and knowledge-based approach. The data-based approaches use conditional independence semantics of Bayes nets to induce models from data [17]. The knowledge-based approach use causal knowledge of domain experts in constructing Bayesian networks [25]. The knowledge-based approach is especially useful in situations where domain knowledge is crucial and availability of data is scarce. Elicitation of qualitative knowledge from humans is critical in constructing BNs because humans find it easier to handle qualitative than quantitative data [29]. Moreover, the inference procedures in a BN are more sensitive to the qualitative structure than the quantitative probabilities associated with the structure [33]. Consequently, the most effective BNs are those that combine the qualitative structure based on expert knowledge with the quantitative probabilities identified and revised using hard data. Despite the importance of the

knowledge-based approach and the qualitative structure of the BN in making inference, few systematic techniques exist to construct the qualitative structure of the BNs. In this paper, we propose a causal mapping approach to the construction of BNs based on expert knowledge.

Recently, there has been a growing interest in the use of causal maps to represent domain knowledge of decision-makers [1, 20, 26]. 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 [2] and management sciences [23, 31] to represent salient factors, knowledge, and conditions that influence decision-making. Causal maps have been useful in practice. For example, Axelrod [2] 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 [35] 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.

Causal maps are useful tools to construct Bayesian networks for several reasons. First, causal maps capture *causal* knowledge of experts about a domain that other methods such as protocol analysis and repertory grids cannot capture. Causal knowledge of experts is especially important in the context of decision making because decision problems are described and understood through causal connections. Second, causal maps represent domain knowledge more descriptively than other models such as regression or structural equations. Third, causal mapping is more comprehensive, less time-consuming and causes lesser inconvenience to experts during knowledge elicitation than other techniques such as protocol analysis and repertory grids [5]. Finally, causal maps lend themselves to different types of statistical analysis including matrix

algebra and network analytic methods [2, 3, 6, 13], relation algebra [7], system dynamics [16, 37], decision trees [9] and neural networks [36].

However, despite these advantages, there are some important differences between the network representations of causal maps and Bayesian networks that must be addressed in using causal maps to construct Bayesian Networks [28]. The primary purpose of this paper is to propose a systematic procedure for constructing a causal Bayesian network by combining exploratory and confirmatory methods of constructing causal maps with the technique for converting causal maps to Bayes nets described in [28]. The resulting technique can be used as a semi-formal method for construction of Bayesian networks starting from a domain expert. Following the terminology in [28], we call these graphical structures “Bayesian causal maps.” In this paper, we illustrate this method with a new case study on online ticketing application outsourcing decision in a technology organization.

An outline of the remainder of the paper is as follows. In Section 2, we discuss Bayesian networks, their semantics, and the process of making inferences. In Section 3, we discuss the definition and components of a causal map. In Section 4, we discuss the similarities in and differences between causal maps and Bayesian networks and describe Bayesian Causal Maps and how they are different from causal maps and Bayesian Networks. In Section 5, we describe a procedure for constructing a causal map and its conversion to a Bayesian network. In Section 6, we discuss a case study of an online ticketing application outsourcing decision in a technology organization. In Section 7, we list the advantages and applications of Bayesian causal maps. Finally, in Section 8, we conclude with a summary and a statement of future research.

2 Bayesian Networks

In this section, we briefly describe the definition and semantics of Bayesian networks.

2.1 Definition

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 [33].

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. Figure 1 shows a Bayesian network consisting of four discrete variables: Mileage (M), Brand (B), Car Performance (C), and Purchase the Car (P). 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 Figure 1, e.g., Mileage has two states: ‘high’ and ‘low;’ Brand has two states: ‘Good’ and ‘Bad;’ Car Performance has two states: ‘high’ and ‘low;’ and Purchase the Car has two states: ‘Yes’ and ‘No.’ If there is an arc pointing from X to Y , we say X is a *parent* of Y . In Figure 1, Mileage and Brand have no parents. However, Car Performance has two parents (Mileage and Brand) and Buy the Car has one parent (Car Performance). For each variable, we need to specify a table of conditional probability distributions, one for each configuration of states of its parents. Figure 1 shows these tables of conditional distributions— $P(M)$, $P(B)$, $P(C \mid M, B)$, and $P(P \mid C)$.

Figure 1 goes about here

2.2 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 Figure 1, e.g., we are assuming that

$$P(M, B, C, P) = P(M) \cdot P(B) \cdot P(C \mid M, B) \cdot P(P \mid C),$$

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

$$P(M, P, C, B) = P(M) \cdot P(B \mid M) \cdot P(C \mid M, B) \cdot P(P \mid M, B, C).$$

Comparing the two, we notice that we are making the following assumptions: $P(B \mid M) = P(B)$, i.e., B is independent of M ; and $P(P \mid M, B, C) = P(P \mid C)$, i.e., P is conditionally independent of M and B given C .

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 Figure 1, e.g., one such sequence is $M B C P$. Then, the conditional independence assumptions can be stated as follows. For each variable in the sequence, we are assuming 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 M to B signifies that M is independent of B ; the lack of an arc from M to P and from B to P signifies that P is conditionally independent of M and B given C .

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 [27]. [27] and [19] 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. In our car example, P is conditionally independent of M and B given C . This statement can be interpreted as follows. If the true state of C is known, then in assigning probabilities to states of P , the states of M and B are irrelevant. In other words, if we know that the performance of the car is good, then any knowledge about brand and mileage is irrelevant to the probabilities of purchasing of the car.

In practice, the notion of direct causality is often used to make judgments of conditional independence. Consider in our car example, a situation where M directly causes C and C in turn directly causes P , i.e., the causal effect of M on P is completely mediated by C . Then it is clear that although M is relevant to P , if we know the true state of C , further knowledge of M is irrelevant (for assigning probabilities) to P , i.e., P is conditionally independent of M given C . This situation is represented by the Bayesian network $M \rightarrow C \rightarrow P$ in which there is no arc from M to P . As another example, consider a situation where variable X directly causes variable Y and variable X also directly causes variable 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 $X \perp\!\!\!\perp Z \perp\!\!\!\perp Y$ in which there is no arc from X to Y or vice-versa.

2.3 Making Probabilistic 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 [29]. 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 (un-normalized) *posterior* joint distribution of the variables. Thus, the joint distribution of variables changes each time we learn new information about the variables.

3 Causal Maps

Causal maps, also called cause maps or cognitive maps, are directed graphs that represent the cause-effect relations embedded in experts' thinking. Eden [12] defines a cognitive 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. There are three major components of a causal map: *causal concept*, *causal connection* and *causal value*. Figure 2 shows a part of a causal map of a prospective buyer relating to the decision of whether to buy or not to buy a particular used car.

3.1 Causal concept

A causal concept is a single ideational category [6]. It can be an attribute, issue, factor or variable of a domain, and is represented by a node in the causal map. A concept can be a single word such as ‘Mileage,’ ‘Age,’ and ‘Price;’ a composite word such as ‘Good Brand,’ ‘Fuel Efficiency,’ ‘Car Performance’ and ‘Time Belt Condition;’ or a more complex phrase such as ‘Condition of Car Parts,’ ‘Accident Record of the Car’ and ‘Buy the Car.’

Figure 2 goes about here

3.2 Causal connection

A causal connection is a tie that links two concepts in the map and is represented with a unidirectional arrow. It depicts an *antecedent-consequence* relation between two concepts. The concept at the tail of an arrow is taken to *cause* the concept at the head of the arrow. In Figure 2, ‘Mileage,’ ‘Age,’ ‘Brand’ and ‘Fuel Efficiency’ determine the performance as well as the price of the car. ‘Accident Record of the Car’ determines the ‘Engine Condition,’ and the ‘Transmission Condition.’ Similarly, the ‘Engine Condition,’ ‘Transmission Condition,’ and ‘Timing Belt Condition’ determine the condition of the parts of the car. Finally, the decision to buy the car is a consequence of ‘Price,’ ‘Performance,’ and ‘Car Condition.’

A causal connection can be positive or negative. A positive connection indicates that an increase in the causal concept leads to an increase in the effect concept, whereas a negative connection indicates that an increase in the causal concepts leads to a decrease in the effect

concept. In Figure 2, for example, ‘Fuel Efficiency’ and ‘Good Brand’ exert a positive influence on the ‘Car Performance.’ Thus, the higher the fuel efficiency and better the brand, the higher will be the performance of the car. On the other hand, ‘Age’ and ‘Mileage’ have a negative influence on ‘Car Performance.’ Thus, higher the age and mileage of the car, lower the performance of the car.

3.3 Causal value

A causal value represents the strength of the causal connection. Different techniques have been used to determine the causal value including social networks and matrix algebra [2, 6], system dynamics [16], relation algebra [7], and neural networks [36]. The choice of techniques used to determine the causal value is determined by the purpose of analysis. In this study, we use causal maps to construct Bayesian networks and represent the causal values as Bayesian probabilities.

4 Transforming Causal Maps to Bayesian Networks

Although Bayesian networks and causal maps are causal models that represent cause-effect beliefs of experts, there are some differences in the two approaches to modeling that need to be addressed if we are to transform causal maps to Bayesian Networks. These differences are discussed in the following paragraphs. Most of the discussion in this section is taken from [28].

4.1 Conditional independencies

A network model can be either a dependence map (D-map) or an independence map (I-map) [29]. A D-map guarantees that concepts found to be connected are indeed dependent; however, it may display a pair of dependent concepts as a pair of separated concepts. In other words, in a D-map, a link or arrow between two nodes in the model implies that the two nodes are related. However, a lack of an arrow between nodes does not necessarily imply independence between

the two nodes. An I-map, on the other hand, guarantees that concepts found to be separated are indeed conditionally independent, given other variables. However, it may display a pair of independent concepts as connected concepts. Thus, in an I-Map, lack of an arrow implies independence between two nodes, whereas the presence of an arrow between two nodes does not necessarily imply that the two nodes are related.

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. However, 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. This is because the process for deriving causal maps is exploratory. A lack of an arrow may result from the lack of articulation of an arc on the part of the expert. It does not imply that the expert believes the nodes to be independent.

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 ensure that the lack of links between the concepts in the causal maps implies independence and the presence of links between concepts implies dependence. In other words, we need to make causal maps both D-maps and I-maps.

An example of this is shown in Figure 3. The solid arrows in the figure represent links identified in the original causal map based on the narrative yielded by an open ended exploratory interview conducted with the buyer. However, when the expert was shown the original causal

map based on the first interview and asked if the map was accurate or if the buyer wanted to make any further changes to the causal map in terms missing links, redundant links and wrong direction of the links, the buyer added three more links shown by the dotted lines in Figure 3: From Mileage to Car Performance, from Car History to Car Performance and from Brand to Price of the Car. The addition of the three new links changes the inference about the variables in the map. For example, car performance was independent of Mileage and Car History in the original map. However, in the new map Mileage and Car History are relevant in making inferences about Car Performance. Similarly, in the original map, Brand was not relevant to making inferences about the Price of the Car. However, in the new map, Price of Car is dependent not only on Mileage but also on Brand. In short, the new links change the conditional independence assumptions about variables in the map.

Figure 3 goes about here

4.2 Reasoning underlying cause-effect relations

Causal maps identify individuals' perceptions of cause-effect relationships between variables based on language rather than the reasoning processes [6]. Studies in managerial cognition indicate that individuals reason by accumulating possibly significant pieces of information and organizing them in relation to each other so as to be able to combine them into a conclusion and decision [8]. 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* [8]. 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 Figure 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 'Car been in an accident' to 'Dent in the body of the car' of the car. Causal statement 2 involves abductive reasoning. Since information about whether the car has been in an accident is not known to the buyer (in this case), the buyer is making inference about this unknown variable based on his/her observation of the Dent in the body of the car. This does not imply that dent in the car causes the car to be have been in an accident. 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 Figure 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 reasoning underlying the causal statements rather than the language used.

Figure 4 goes about here

4.3 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 [12, 13]. For example, a direct link between two concepts in the causal map does not guarantee a direct relationship between the two concepts. It just implies a relation between the two concepts that can be either direct or indirect. This distinction is important to identify conditional independencies in the causal maps. Figure 5 depicts how a lack of distinction between direct and indirect relationship affects conditional independence assumptions in a causal map.

Figure 5 goes about here

In Figure 5, both ‘accident record of the car’ and ‘performance of the car’ affect the decision of whether to ‘buy/not buy the car.’ In the modified Bayesian causal map, there is no linkage between accident record of the car and Buy/not buy implying that accident record of the car impacts the decision to buy the car strictly through performance of the car. If we have complete information on performance of the car, any additional information on accident Record of the car would be irrelevant in making inferences about the decision to buy/not buy the car.

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 Accident History of the Car (in Figure 5) affects the decision to Buy/Not Buy the Car only through Performance of the Car, then an arrow from Accident History of the Car to the decision to Buy/Not Buy the Car 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.

4.4 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 [4, 12, 20]. 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. In addition to coding mistakes, feedback loops may indicate dynamic relations between variables over time. 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, disaggregating the variables into two time frames can often solve the problem of circularity.

Figure 6 goes about here

For example, Figure 6 shows a reciprocal causal relation between Accident Record of the Car and Car Performance and reasoning underlying this circular relation. Arrow t_1 implies that the prior or past accident record of the car affects the future car performance. In other words, if the car has been in an accident, then there may be problems with the car that may affect its current performance. Arrow t_2 implies that current car performance can affect the future accident record of the car. The circular relation has resulted from aggregation of the variable Accident Record of the Car across two time frames: t_1 and t_2 . After de-aggregating Accident Record of the Car into two time frames, we get an acyclic relation between the three variables. To make the causal map acyclic, we can either include both the arrows in the map or we can arbitrarily retain one of the two relations and exclude the other from the causal map. This will depend on the time frame of the decision being modeled. 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.

5 Constructing Bayesian Causal Maps

In this section, we propose a systematic procedure to construct Bayesian Causal Maps. The procedure comprises four main steps:

1. Data elicitation
2. Derivation of causal maps
3. Modification of causal maps to construct Bayesian Causal Maps
4. Derivation of the Parameters of Bayesian Causal Maps

In the first step—data elicitation—an individual domain expert is interviewed using qualitative interview to elicit his/her domain knowledge and the experts' response to the

interview is transcribed to get a text that we call a ‘narrative.’ In the second step, the narrative obtained in the first step is analyzed using a systematic content analysis technique to represent the narrative in the form of a causal map of the expert. In the third step, the causal map of the expert is modified to eliminate biases that result from the use of textual analysis and to make the structure of the causal maps compatible with Bayesian Networks. In the final step, the parameters of the Bayesian Causal Maps are derived using probability-encoding techniques.

5.1 Data Elicitation

In this step, domain information is elicited from the expert. Two different types of elicitation techniques are typically employed to capture domain information: *structured* and *unstructured*. The structured techniques are based on a confirmatory approach to data elicitation whereas unstructured methods are based on the exploratory approach. These two types of techniques differ in terms of *purpose* and *type of knowledge elicited* [6]. In the structured techniques, experts are provided with a list of pre-defined concepts and are asked to specify the direction and sign (positive and negative) between the concepts. Structured methods are more suitable for confirming and validating expert knowledge rather than for eliciting expert knowledge for domains that are not clearly defined. On the other hand, the purpose of an unstructured approach is to inductively explore a new or unfamiliar domain by posing questions such as: “*What are the factors relevant to the decision?*” The unstructured approach yields a richer understanding of the processes that individuals engage in decision-making as well as helps gather important insights into the general knowledge that individuals have on the domain being evaluated.

The choice of elicitation methods affects the data elicitation process as well as the coding process in constructing cognitive maps. In the structured methods, the concepts in the cognitive maps are defined *a priori* by modelers and these concepts are imposed on the experts from whom

the knowledge is elicited. Hence, in the structured methods, researchers know the number of concepts in the cognitive map. On the other hand, in the unstructured methods, the concepts emerge from the data or the narrative of the expert. In this paper, we propose a combination of the two techniques to facilitate the elicitation of data using an inductive approach and the validation of data using a confirmatory approach. The structured methods are discussed in greater detail in the validation section.

In unstructured methods, in-depth qualitative and open-ended questions are posed to the expert to obtain raw data in the form of a narrative. This narrative is then used to construct cognitive maps using textual analysis. Unstructured methods are most appropriate for eliciting expert knowledge because they are exploratory and less intrusive. This is because the concepts and the links between concepts are allowed to emerge in the process of interviews by sequencing the interview questions based on the responses of the expert. These methods are particularly suitable for eliciting expert knowledge for unknown and ill-structured domains. The knowledge elicited through unstructured methods can be validated using structured methods. A widely used qualitative interview technique that can be used to elicit a narrative is open interview with probes [32]. This interview consists of three different types of questions: *broad-open ended questions*, *probing questions* and *closed questions*.

An example of an open interview with probes conducted with the prospective buyer (university graduate student) relating to the decision of whether or not to buy a 1995 Honda Accord car is presented in Figure 7. As shown in the Figure, the interview starts by posing a broad question to extract the general decision variables such as: “What factors would you consider in deciding whether or not to buy a 1995 Honda Accord car LX?” The subject’s answer to this question can then be used to identify ‘probes’ or key phrases identified by the subject.

Subsequent questions presented to the subject relate to each of these probes in terms of direct questions as well as indirect relationships with other probes offered by the subject. Closed questions are very specific and require the subject to answer as either ‘yes’ or ‘no.’ Closed questions are primarily used for clarification purposes. The bold phrases in the prospective buyer’s response represent the probes identified by the interviewer. For example in the interview shown in Figure 7, “performance of the car” and “condition of car parts” are probes used by the interviewer to get more detailed factors that determine the performance of car. The probing question on car performance (question 2) yielded 3 additional probes: “age,” “mileage” and “fuel”; the probing question on condition of car parts also yielded 3 probes: “engine,” “transmission,” and “time belt.” These probes were used to get more detailed information about each of these concepts. This probing continues till the prospective buyer has exhausted the list of factors that make-up a domain and he/she cannot think of any additional factors. The responses of the expert to the open interview can be transcribed to yield a ‘narrative’ or a ‘text.’ This narrative or text is then analyzed using a systematic procedure of textual analysis to derive causal maps. This procedure is described in the next section.

Figure 7 goes about here

5.2 Derivation of causal maps.

There are four different steps in deriving causal maps using narrative or text yielded by the interview [2, 20]. These steps are shown in Figure 8 and discussed in detail in the following paragraphs.

Figure 8 goes about here

5.2.1 Identify Causal Statements in the Narrative

The first step in constructing causal maps is to identify causal statements in the narrative. Causal statements are statements in the narrative that explicitly contain a cause-effect relationship. A causal statement links two different concepts through a causal connector. An important consideration in identifying causal statements in a narrative is to define rules for recognizing causal connectors. This involves developing a comprehensive dictionary of words or phrases that can be considered as causal connectors. Examples of words used to represent causal connectors include 'if-then', 'because', 'so,' 'as,' 'therefore' etc. Each statement containing a causal connector can be identified as a 'causal statement.' This can be done either manually or can be automated. In the manual procedure, multiple raters can develop a comprehensive dictionary of causal connectors before going through the narrative or the text yielded by the open interview. They can then recognize the causal connectors in the narrative to identify causal statements. The advantage of the manual procedure is that raters can add new causal connectors to the pre-defined list of causal connectors while going through the narrative and hence the chance of missing a uniquely worded causal statement is low. But at the same time, the manual procedure is labor intensive and time consuming.

Alternatively, causal statements can be identified using an automated process. In an automated process, two types of files are created: narrative files and causal connector file. A separate text file is created for each expert narrative (response to the open interview). The causal connector file contains the list of causal connectors. The advantage of an automated process is that it is time saving and is not labor intensive. However, the disadvantage of the automated process is that the list of causal connectors needs to be defined before hand. Causal connectors cannot be added to the predefined list as is done in the case of manual process. This may result in

the loss of some peculiarly worded causal statements in the narrative that do not contain pre-defined causal connectors. The choice of methods may depend on pragmatic factors such as the length of the text, the complexity of the domain etc. Figure 8 shows two causal statements identified from the narrative of the subject interviewed. These two statements were identified as causal statements because they contain words identified as causal connectors: ‘leads to’ and ‘if-then.’

5.2.2 Construct Raw Causal Maps

Once the causal statements are identified, they are broken into causal phrases, causal connectors and effect phrases to derive the raw cognitive maps. Again, this can be done either manually or can be automated. Figure 8 shows how the two causal statements identified in step 1 are broken into raw cognitive maps. This process can also be automated by defining the rule for classifying phrases in the causal statements into cause and effect phrases. A separate rule needs to be defined for each causal connector. For example, a phrase immediately following ‘if’ can be classified as a causal phrase, whereas a phrase following ‘then’ can be classified as an effect phrase. Similarly, the phrase immediately before ‘leads to’ can be classified as a causal phrase, whereas the phrase following ‘leads to’ can be classified as an effect phrase. The automated process is less labor intensive, less time consuming and more reliable than is the manual process. However, it may result in misclassification of some peculiarly worded statements.

5.2.3 Design Coding scheme

The raw causal maps derived in step 2 are cast in the language of the expert. In spite of their usefulness, the raw maps obscure analysis because of their complexity. Hence, there is a need to design a coding scheme to recast the raw causal maps into the final cognitive maps. This process of coding is called *filtering* or *aggregation*. Aggregation is the process of determining which part

of the text to code, and what words to use in the coding scheme. Aggregating phrases in the raw causal maps into generalized concepts can be used to move the coded text beyond explicitly articulated idea to implied or tacit ideas. Aggregation can also be used to avoid misclassification of concepts due to peculiar wording on the part of individuals.

In the process of coding, the raters have to decide which words in the raw causal phrases to retain and which words to delete. The raters also have to decide which part of the phrase needs to be reworded. The raw causal phrase can be changed into a coded concept that may be a single word, a composite word or a complex phrase. This process requires human interpretation and it is recommended that it be done manually. Multiple raters can code raw phrases into coded concepts using the ‘majority’ rule or the ‘consensus’ rule. In other words, all or a majority of the raters must agree on the coded concept used to represent the raw phrase.

Additionally, it is important to confirm that the coded concepts capture the meaning implied by the raw causal phrases to avoid inconsistencies between raw cognitive maps and coded cognitive maps. This can be done through a close collaboration of the expert whose causal map is being constructed. The expert can be shown the coded concepts used to recast the phrases used by him/her in the interview. The input of the original experts is crucial in this stage to avoid inconsistencies between raw cognitive maps and coded cognitive maps. Figure 8 shows how phrases used by the prospective buyer in the raw cognitive maps are coded into generalized concepts. The coding scheme was developed by two raters using the consensus rule and was confirmed with the buyer who was interviewed in step 1.

5.2.4 Convert Raw Causal Maps into Coded Causal Maps

Finally, the coding scheme developed in step 3 is used to recast the raw cognitive maps into coded maps. A coded cognitive map is a network of concepts formed from causal statements in a

narrative depicting directionality (cause-effect) and sign (positive and negative) of the relations between the concepts. Two statements are linked if they share one concept. For example, causal statement 1 and causal statement 2 in Figure 8 share the concept “Car Performance” thus resulting in the network of “Mileage → Car Performance → Buy the Car.”

5.3 Modification of Causal Maps

As discussed earlier, the structure of the causal maps requires modification to make it compatible with the Bayesian network by paying attention to four major modeling issues: conditional independencies, reasoning underlying the link between concepts, distinction between direct and indirect relations and eliminating circular relations. Structured methods are appropriate tools to eliminate the four biases discussed above. Two most widely used structured methods are *structured interviews* and *adjacency matrices*. In structured interviews, the experts are provided a list of paired concepts as well as different alternative specifications of the relation between the concepts in the original map. The experts are then instructed to choose an alternative to specify the direct relation between the pair of concepts. Figure 9 is an illustration of a part of a structured interview filled by a prospective car buyer.

Figure 9 goes about here

Alternatively, experts can be provided the concepts in the form of an adjacency matrix (shown in Figure 10), where the rows represent causes and columns represent effects. The experts are asked to enter ‘0’ (no relation), ‘+’ (positive relation) or ‘-’ (negative relation) in each cell to specify the relation between two concepts in the matrix. These two structured methods help in removing the four modeling biases relating to the construction of Bayesian Causal Maps.

Figure 10 goes about here

5.4 Deriving Parameters

Once the structure of the Bayesian Causal Maps is constructed, numerical parameters of this modified structure need to be assessed so that the propagation algorithms in the Bayesian network can be used to make inferences.

The causal map has been used primarily to qualitatively describe the variables used by experts to describe a particular decision domain. The focus of causal maps is to analyze the structure of the map using network analysis techniques [24]. Consequently, the uncertainty associated with the different variables in causal map is not captured by a causal map. All variables are assumed to have the same level of uncertainty. A Bayesian network allows a decision-maker to make inferences on the different variables in the network based on the information about other variables in the network. In order to be able to make inferences, we need to assess uncertainty associated with every variable in the map and the interactive effects of multiple causal variables on effect variables.

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*.

The parameters of the Bayesian causal maps can be derived in two steps: identification of state space of each variable in the Bayesian causal map and derivation of the conditional probabilities associated with the variables in the map. To identify the state space of each

variable, it is very important to develop precise definitions of each concept [28]. This is especially important because the meaning associated with the variables in the causal maps is not universal and depends on the perceptions of the experts. Precise definitions also specify the scope of each variable that is especially useful in making inferences based on the Bayesian causal maps. Experts can be asked to define each variable in the map. These definitions can be further modified and validated through subsequent interviews aimed at clarifications. For example, our expert defined car performance in terms of “reliability, safety and driving pleasure.” Based on the definitions provided by the experts, the state space of the variables can be established. For example, the five variables in Figure 10 have the following states: Mileage (High, Low), Age (Old, New), Fuel Efficiency (High, Low), Car Performance (Good, Bad), Brand Quality (High, Low).

Once the states of the variables in the causal maps are specified, the conditional probabilities associated with the variables can be derived using probability-encoding techniques. Many different probability-encoding techniques are available (for a detailed review see [34]) 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. [34] describes 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. [18, 30]. Once the parameters of the causal map are identified, propagation algorithms can be used to make inferences about the variables in the causal maps.

6 A Case Study: IT Application Outsourcing Decision

This section describes a case study of a construction of a Bayesian causal map. First, we illustrate how starting from a causal map, we constructed 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. Second, we show how Bayesian network software can be used to draw probabilistic inferences in a Bayesian causal map.

6.1 Decision context

We used a real-time IT application outsourcing decision analyzed by a systems analyst of a big-five consulting firm. We chose this decision context for two major reasons. First, IT application outsourcing is an emerging domain and the boundaries of this domain are not clearly defined. Therefore, this domain is appropriate for the exploratory approach of constructing causal maps. Second, the decision alternatives, outcomes and application environmental factors involved uncertainty and required the systems analyst to use his/her intuition. It allowed the analyst to

develop his/her own framework in diagnosis, analyses, and recommendations of decision options.

The IT application outsourcing decision was as follows. A major airline company had recently decided to develop an online ticketing system. Although the company is a well-established company with a well-developed regular ticketing processes, online ticketing is a totally new concept to the company. The decision faced by the airline company is whether to develop the online ticketing system in-house or to use an application service provider's (ASP) application. The role of the system analyst was to analyze the decision and suggest a recommendation.

6.2 Subject

The subject was a systems analyst in a big-five consulting firm. She had an MBA in information systems and 2 years of experience in systems analysis and design. The subject was a part of the team that analyzed the online ticketing outsourcing decision.

6.3 Procedure for Constructing a Bayesian Causal Map.

6.3.1 Step 1: Data Elicitation

The subject was interviewed using an open-ended interview with probes. The interview lasted about two hours. The interview began with a very broad question: "What do you think are the key factors affecting the Online Ticketing Application Outsourcing Decision at ABC Airlines?"

The subsequent 'probes' were based on the factors suggested by the expert. The probing continued till a comprehensive list of factors relating to the outsourcing decision was elicited and the subject could not think of any additional factors.

6.3.2 Step 2: Derivation of Causal Map

The causal map of the expert was constructed from the narratives yielded by the interviews using the four-step procedure described in Section 5.

Identifying causal statements. Two raters identified the causal statements based on a comprehensive list of causal connectors developed by the two raters. A statement was identified as causal if it contained one of the words listed as causal connectors. Examples of words contained in the list of causal connectors include ‘if-then,’ ‘because,’ ‘so,’ etc.

Raw cognitive maps. The causal statements identified in step 1 were broken into causal phrases, causal connectors and effect phrases to derive the raw causal maps.

Coding scheme. The phrases used by the expert were coded into generalized concepts by two raters. We closely consulted the expert to ensure that the coded concepts did not deviate from the original cause and effect phrases used by the expert in his/her interview response.

Coded causal map. The coding scheme developed in the previous step was used to recast the raw causal map into coded map. The Net-analysis program was used to construct the causal map of the expert. The input file contained all the *causal pairs* identified by the expert in the form of causal concept, effect concept, direction the link and sign of the link. The program then identifies the common concepts between different causal pairs and links these causal pairs. It provides the output in the form of an adjacency matrix that includes all the links between pairs of concepts in the map.

The original causal map shown in Figure 11 describes the subject’s causal perceptions of the decision problem in the online ticketing decision. There are 23 variables in the map that can be broadly classified as in-house application cost variables, ASP outsourcing cost variables, risk determinant variables, application environment variables, and decision variables. A brief

definition and the possible states of each variable are shown in Tables 1 and 2. The variable *In-house Feasibility* was not captured in the original causal map. It was identified in the follow-up interviews conducted with the subject.

Figure 11 goes about here

6.3.3 Step 3: Modification of Original Causal map

The original causal map was modified using the structured questionnaire. The expert was provided with a list of paired concepts and alternative specifications of the direction ($\rightarrow/\leftarrow/0$) and sign (+: positive relation, -: negative relation, 0: no relation) of the relation between the concepts (as shown in Figure 9). The expert was then instructed to choose the alternative that best specified the relation between the pair of concepts. This procedure eliminated the four modeling biases discussed below. The modified causal map is shown in Figure 12.

Figure 12 goes about here

1. Direct causality between variables. The input expert's response to the structured interview resulted in two major changes relating to direct versus indirect relation between concepts in the original map. First, in the original causal map, Business Domain was directly related to In-house or ASP variable. However, the subject clarified this relation in the structured questionnaire interview and suggested an indirect relationship between Business Domain and In-house/ASP variable. In the modified map, Business Domain indirectly affects In-house/ASP through two different variables: Risk Preferences and Application Cost. Second, there was a direct relation between Application Maturity and In-house/ASP variable in the original causal

maps. However, in the follow-up interview, it was found that Application Maturity also affects In-house/ASP variable through Risk Preferences and Application Cost.

Table 1 goes about here

Table 2 goes about here

2. Conditional Independence. The expert's response to the structured interview yielded five additional links in the original map. Overall Feasibility was a new variable added to the map. This variable did not exist in the original map. First is the link between Labor Market and Overall Feasibility that did not exist in the original causal map. Second is the link between Knowledge, Skills and Abilities and Feasibility. Third is the link between Process Maturity and Product Customization. These variables were shown as conditionally independent in the original map. Similarly, the modified map shows links between Product Customization and Maintaining Application Currency and Feasibility and In-house/ASP that did not exist in the original causal map.

3. Deductive reasoning. The original map shows a link from Labor Cost to Knowledge, Skills & Abilities. However, the direction of the arrow is based on abductive reasoning from the observable measure (Labor) to the latent cause (Knowledge, Skills & Abilities). The direction of this arrow changed in the modified causal map from Knowledge, Skills and Abilities to Labor Cost.

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 online ticketing application decision was reached (to eliminate circular relations due to relations pertaining to

different periods). This resulted in the elimination of two reciprocal relationships in the original causal map. First is the two-way relation between In-house Application Cost and Online Ticketing and second is the circular relation between In-house/ASP variable and Net Application Value-added. The resultant structure of the Bayesian Causal Map is shown in Figure 12.

6.3.4 Step 4: Assessing Parameters

In this step, the parameters of the Bayesian Causal Map 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.

1. To rate the marginal and conditional probabilities on a discrete scale (0 to 10); and
2. 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 [18, 30].

6.4 Validating the Bayes Net Model

We used Netica [www.norsys.com] to make probabilistic inferences using 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 Online Ticketing Application option under different scenarios. The scenarios were defined in consultation with two IS experts (including the subject), and they represent situations in which there are unambiguous prescriptions for Application Outsourcing decisions in the IS literature. We illustrate how predictions can be made about our subject's perceptions of In-

house/ASP decision under different information conditions. The decision prescriptions yielded by the Bayesian Causal Maps were checked for validity with the expert

We specify two different scenarios and show how our inferences about Application decision option change depending on the external application cost, application feasibility and environment factors. In the first scenario, we consider a favorable labor market and an unstable and complex application environment. In terms of application environment, we consider an unstable market environment as well as a high sensitivity of passenger data. Accordingly we specify the states of the following four firm variables in the map: Knowledge, Skills & Abilities = low, Labor Market = high, Fluctuations in Sales = low, and Sensitivity of Passenger Data = low. Since the case specifically mentioned that online ticketing was a new process for the airline company, we specified Application Maturity = New. 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 Application Cost, In-house Feasibility, Net Application Value-added, and In-house/ASP Decision is shown in Table 3. When additional information is received about application cost and environment factors, the posterior probability of Application Cost = In-house < ASP increases from 0.46 to 0.67, that of (in-house) Feasibility = high increases from 0.55 to 0.85 and that of Net Application Value-added = In-house > ASP increases from 0.61 to 0.66. These posterior marginals change our inference about the state of In-house/ASP Decision. The posterior probability of In-house Application is 0.54 in comparison to a prior of 0.49. Under the conditions described in scenario 1, our Systems Expert is likely to select In-house Application Development because the cost of in-house development is lower than outsourcing to ASP, the Net Value-added is much higher for In-house development and the feasibility of In-house development is high.

In the second scenario, we considered an adverse labor market and stable and less complex application environment. Accordingly, the states of the two application cost variables and the two application environment variables in the map are specified as follows: Knowledge, Skills & Abilities = high, Labor Market = low, Fluctuations in Sales = low and Sensitivity of Passenger Data = low. As shown in Table 3, the posterior probability of (in-house) Feasibility = low increases from 0.45 to 0.80, that of Application Cost = In-house > ASP increases from 0.54 to 0.56, and that of Net Application Value-added = In-house < ASP increases 0.39 from to 0.58. This implies that in scenario two, our systems expert is more likely to reject In-house Application and select the option of using ASP's application.

Table 3 goes about here

7 Advantages and Applications of Causal Bayesian Networks

By integrating Bayesian networks and causal maps, causal Bayesian networks combine the advantages of the two methods and reduce the limitations of either. First, since causal Bayesian networks are both a D-map and an I-map, they represent relationships between variables more comprehensively than either the Bayesian network or causal maps. Second, causal Bayesian networks allow a robust probabilistic inference based on both causality and conditional independence. This is especially important in decision context, where cause-effect relations are critical in making inferences about decision outcomes. Third, the proposed method imposes no a priori assumptions about the orthogonality of variables or the (non-) existence of interaction terms, symmetry, or linearity. Thus, complex interdependencies can be modeled.

We propose causal Bayesian networks as a useful tool to complement other decision modeling methods and decision aids. First, they can be used to support initial decision-making

[22]. Evidence from prior studies suggests that in the vast majority of cases, decision makers are outperformed by their own bootstrap models due to the elimination of unsystematic errors. However, decision aids are just that, aids, and immense skills are still required in assessing the states of the variables, many of which are intangible and/or difficult to observe. However, a causal Bayesian network makes transparent the drivers behind the overall assessment and the software implementation allows easy what-if and sensitivity analysis by changing variables states and observing the automatically updated decision outcomes. Second, the mere act of explicating and formalizing hitherto tacit decision models surfaces hidden assumptions that can now be scrutinized. Research on causal mapping suggests, that the act of drawing a causal map in itself can reduce decision bias (e.g. Hodgkinson, *et al.*, 1999). Third, novices become experts over time by learning from experience. Bayesian causal maps support this process by making variable assessments and decision drivers explicit and storing them so that they can be compared to reality later. Entering actual rather than predicted values for variables and comparing them, as well as the revised vs. the initial investment probability, can uncover both errors in variable state assessments and errors in the decision model, which can then be refined. Finally, causal Bayesian networks can be used as reference and departure in coaching, teaching, training, and in collaborative learning contexts.

8 Summary and Conclusions

The main goal of this paper is to propose a semi-formal method for constructing the graphical structure of a Bayesian network based on domain knowledge. Our method consists of first constructing a causal map and then converting it to a Bayesian network. We call such Bayesian networks, 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. Finally, using evidence propagation algorithms, Bayesian causal maps allow us to make inferences about the variables in the map. We have illustrated how Bayesian causal maps can be constructed starting from a causal map, and how it can be used to make inferences about a new product decision in different scenarios.

There are some interesting implications of our study. This study enables decision-makers to use causal maps for decision-making. Influence diagrams proposed in [19] use Bayesian network models of uncertainty in addition to decision nodes, utility functions, and information constraints. Thus, Bayesian causal maps can be use for normative decision-making using the framework of influence diagrams.

Acknowledgements

The research reported in this paper has been supported by two grants from the Kansas University Business School PhD Summer Research Fund to both authors, and by a grant from the Kansas University General Research Fund to the second author. An earlier version of this paper was presented at INFORMS-Philadelphia in November 1999. We are grateful to Sonali Murthy, Aditi Chakravorty and Greg Freix, for their comments and discussions.

References

- [1]. J. C. Anderson and D. W. Gerbing, Some methods for respecifying measurement models to obtain unidimensional construct measurement, *Journal of Marketing Research*, 19, (1982).
- [2]. R. Axelrod, *Structure of Decision: The Cognitive Maps of Political Elites* (Princeton University Press, Princeton, NJ, 1976).
- [3]. M. K. Bougon, K. E. Weick, and D. Binkhorst, Cognition in organizations: An analysis of the Utrecht Jazz Orchestra, *Administrative Science Quarterly*, 22 (1977).
- [4]. M. Bougon, Uncovering cognitive maps: The 'self-Q' technique, in Morgan, G. (ed.), *Beyond Method: A Study of Organizational Research Strategies* (Sage, CA, 1983).
- [5]. S. Brown, Cognitive Mapping and Repertory Grids for Qualitative Survey Research: Some Comparative Observations, *Journal of Management Studies*, 29 (1992).
- [6]. K. Carley and M. Palmquist, Extracting, representing and analyzing mental models, *Social Forces*, 70, No., (1992).
- [7]. B. Chaib-Draa, Causal maps: Theory, implementation, and practical application in Multiagent environments, *IEEE Transactions on Knowledge and Data Engineering*, 14, No. 6 (2002).
- [8]. E. Charniak and D. McDermott, *Introduction to Artificial Intelligence* (Addison-Wesley, Reading, MA, 1985).
- [9]. K. Chen, J. C. Mathes, K. Jarboe, and J. Wolfe, Value oriented social decision analysis: Enhancing mutual understanding to resolve public policy issues, *IEEE Transactions on Systems, Man and Cybernetics*, 9, No. 9 (1979).
- [10]. C. W. Churchman, R. L. Ackoff, and E. L. Arnoff, *Introduction to Operations Research* (Wiley, NY, 1957).

- [11]. A. P. Dawid, Conditional independence in statistical theory (with discussion), *Journal of the Royal Statistical Society, Series B*, 41, No. 1, (1979).
- [12]. C. Eden, F. Ackermann and S. Cropper, The analysis of cause maps, *Journal of Management Studies*, 29, No. 3 (1992).
- [13]. C. Eden, S. Jones, and D. Sims, *Thinking in Organizations* (Macmillan, London, UK, 1979).
- [14]. L. Fahey. and V. K. Narayanan, Linking changes in revealed causal maps and environmental change: An empirical study, *Journal of Management Studies*, 26 (1989).
- [15]. M. Fiol and A. S. Huff, Maps for managers: Where are we? Where do we go from here? *Journal of Management Studies*, 29 (1992).
- [16]. J. W. Forrester, *Industrial Dynamics* (MIT Press, Cambridge, MA, 1961).
- [17]. D. Heckerman, Bayesian networks for data mining, *Data Mining and Knowledge Discovery*, 1 (1996).
- [18]. M. Henrion, Some practical issues in constructing belief networks, in L. N. Kanal, T. S. Levitt and J. F. Lemmer (eds.), *Uncertainty in Artificial Intelligence 3*, North-Holland, Amsterdam, 1989).
- [19]. R. Howard and J. Matheson, Influence diagrams, in R. Howard and J. Matheson (eds.) (1984), *Readings on the Principles and Applications of Decision Analysis*, 2, (Strategic Decisions Group, Menlo Park, CA, 1981).
- [20]. A. S. Huff, *Mapping Strategic Thought* (Wiley, Chichester, UK, 1990).
- [21]. E. Jaques and S. C. Clement, *Executive Leadership—A Practical Guide to Managing Complexity* (Cason, Hall and Company, 1996).

- [22]. B. Kemmerer, S. Mishra and, P. P. Shenoy, Bayesian causal maps as decision aids in venture capital decision making: methods and applications, *Academy of Management Proceedings* (2002).
- [23]. J. H. Klein and D. F. Cooper, Cognitive maps of decision workers in complex game, *Journal of Operational Research Society*, 33 (1982).
- [24]. K. Knoke and J. H. Kuklinski, *Network Analysis* (Sage, CA, 1982).
- [25]. K. B. Laskey and S. M. Mahoney, Network fragments: Representing knowledge for constructing probabilistic models, in D. Geiger and P. P. Shenoy, (eds.), *Uncertainty in Artificial Intelligence: Proceedings of the Thirteenth Conference*, (Morgan Kaufmann, CA, 1997).
- [26]. M. Laukkanen, Comparative cause mapping of organizational cognition, in J. R. Meindl, C. Stubbart and J. F. Porac (eds.), *Cognition Within and Between Organizations* (Sage, CA, 1996).
- [27]. S. L. Lauritzen, A. P. Dawid, B. N. Larsen and H.-G. Leimer, Independence properties of directed Markov fields, *Networks*, 20, No. 5, (1990).
- [28]. S. Nadkarni and P. P. Shenoy, A Bayesian network approach to making inferences in causal maps, *European Journal of Operational Research*, 128, No. 3 (2001).
- [29]. J. Pearl, *Probabilistic Reasoning in Intelligent Systems: Networks of Plausible Inference* (Morgan Kaufmann, San Mateo, CA., 1988).
- [30]. M. Pradhan, G. Provan, B. Middleton and M. Henrion, Knowledge engineering for large belief networks, in R. Lopez de Mantaras and D. Poole (eds.), *Uncertainty in Artificial Intelligence: Proceedings of the Tenth Conference*, (Morgan Kaufmann, San Francisco, CA, 1994).

- [31]. L. L. Ross and R. I. Hall, Influence diagrams and organizational power, *Administrative Science Quarterly*, 25 (1980).
- [32]. P. H. Rossi, J. D. Wright and A. B. Anderson, *Handbook of Research Methods* (Academic Press Inc., Orlando FL, 1983).
- [33]. D. J. Spiegelhalter, A. P. Dawid, S. L. Lauritzen and R. G. Cowell, Bayesian analysis in expert systems, *Statistical Science*, 8, No. 3 (1993).
- [34]. C. S. Spetzler and C. S. Staël von Holstein, Probability encoding in decision analysis, *Management Science*, 22, No. 3 (1975).
- [35]. J. A. Swan, Exploring knowledge and cognitions in decisions about technological innovation: Mapping managerial cognitions, *Human Relations*, 48, No. 11 (1995).
- [36]. S. Wang, A dynamic perspective of differences between cognitive maps, *Journal of the Operational Research Society*, 47 (1996).
- [37]. E. F. Wolstenholme and R. G. Coyle, The development of system dynamics as a methodology for systems description and qualitative analysis, *Journal of Operational Research Society*, 34 (1983).

Figure 1. A Bayesian Network with Conditional Probability Tables

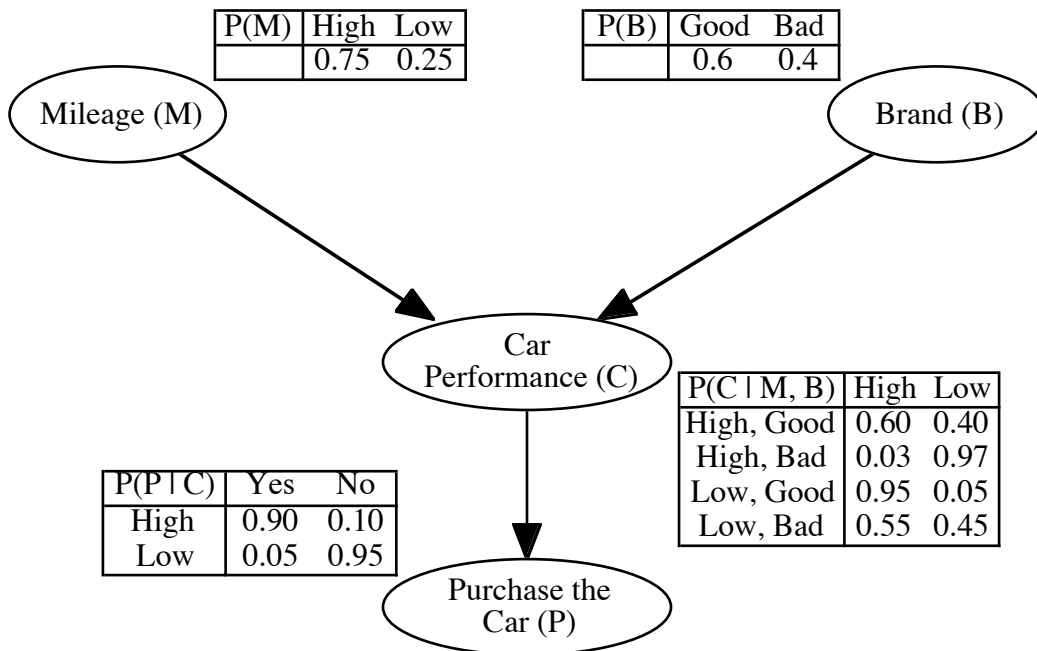


Figure 2. Causal map of a Prospective Buyer Relating to a Used Car Buying Decision

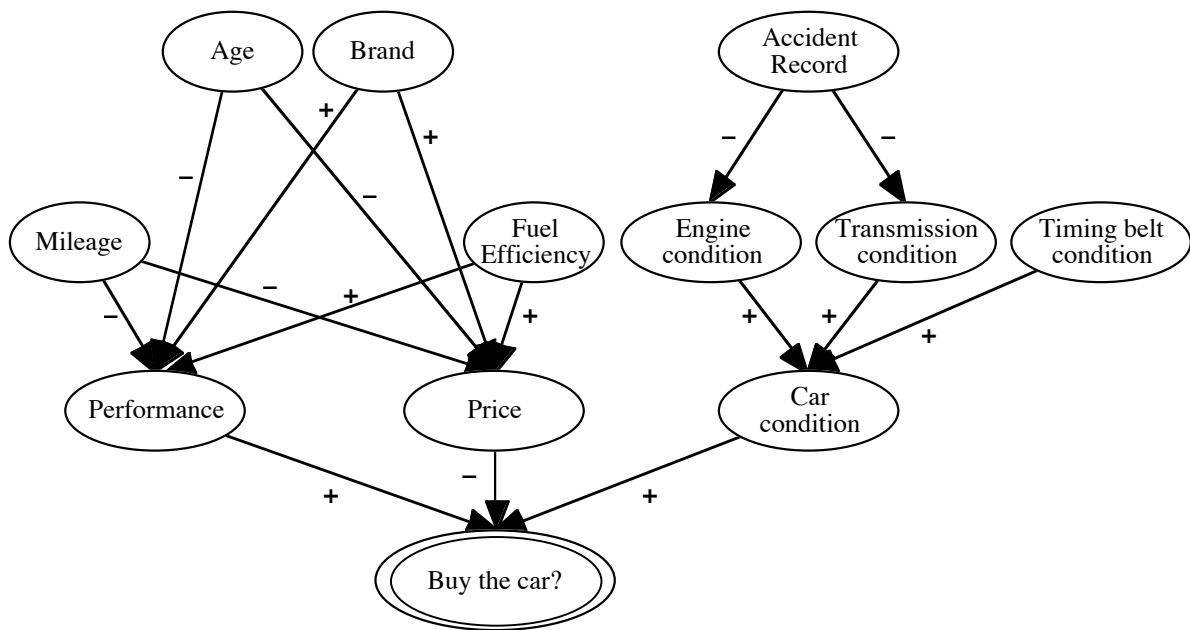


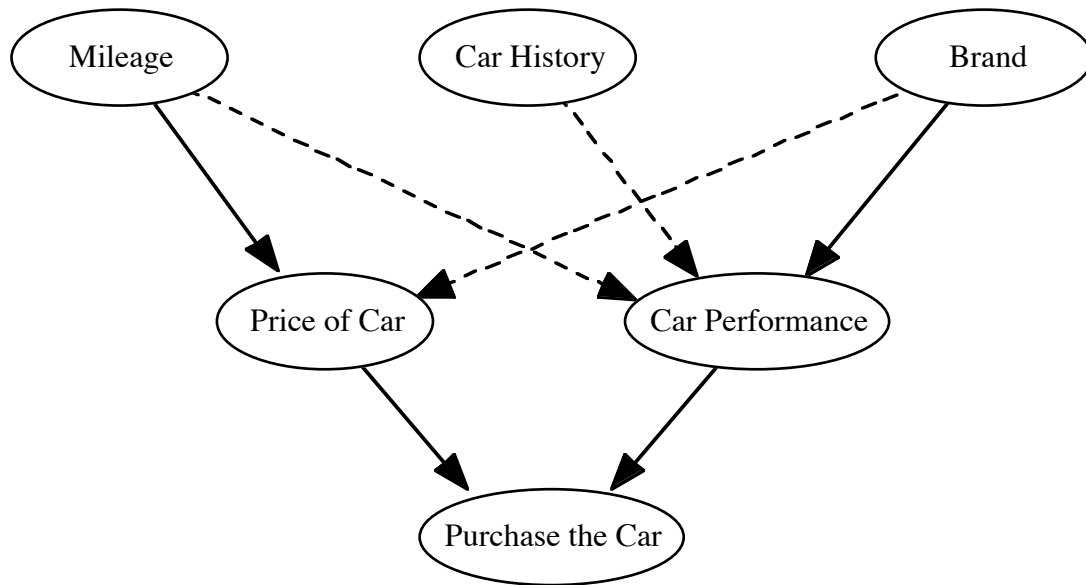

Figure 3. Making a causal map a D-map and an I-map

Figure 4. Distinguishing Between Deductive and Abductive Reasoning

Statement 1. If the car has been in an accident, then there will be a dent in the body of the car.
Causal map based on statement 1: Correct Causal Map



Statement 2. There is a dent in the body of the car, therefore the car may have been in an accident.
Causal map based on statement 2: Incorrect Causal Map




Figure 5. Distinguishing between Direct and Indirect Relations

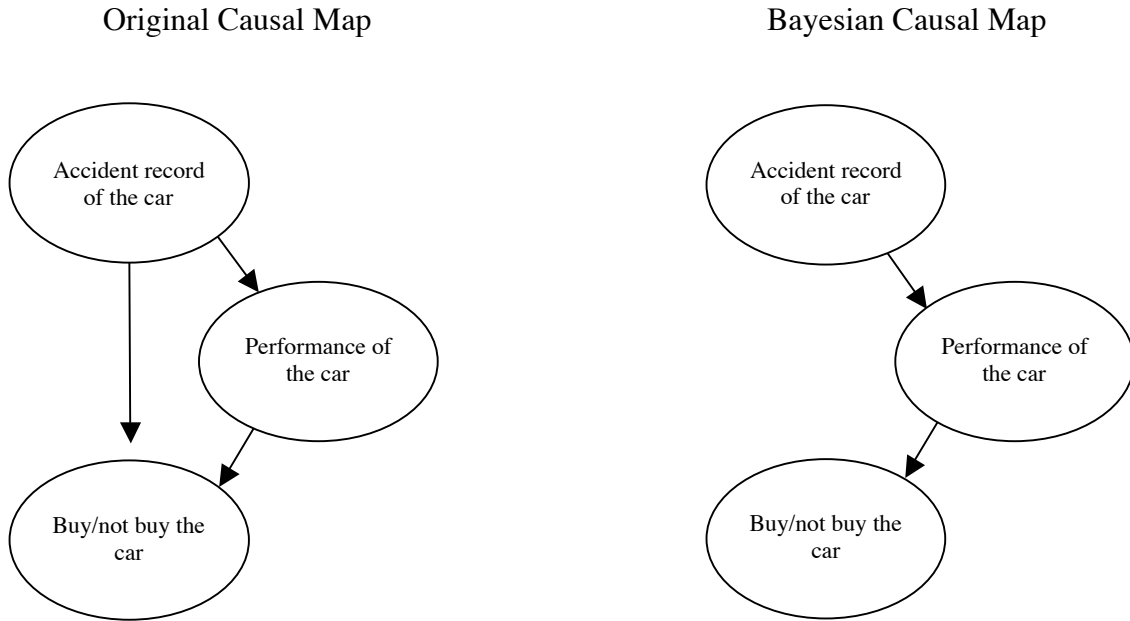


Figure 6. Disaggregating Variables over Time

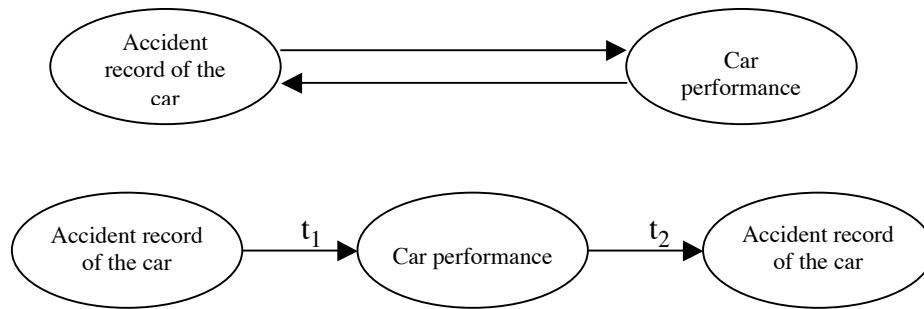


Figure 7. A Part of the Open Ended Interview with Probes Conducted with a Prospective Used Car Buyer

Question 1: What factors would you consider in deciding whether or not to buy the 1995 Honda Accord LX?

Prospective Buyer's response: I would consider **how good is the performance of the car**, ...**how good is the condition of the car parts** and of course the **price of the car**...

Question 2: You mentioned **performance of the car**. What specific factors determine the performance of a car?

Prospective Buyer's response: **Mileage** and **age** of the car definitely affect car performance...I think **fuel efficiency** is also important and affects the performance of a car

Question 3: You mentioned **the condition of car parts**. Which car parts are most important to you?

Prospective Buyer's response: ...**engine** and **transmission** are the most important...Well, **time belt** is also quite important...

Figure 8. Illustration of the Procedure for Deriving Causal Maps

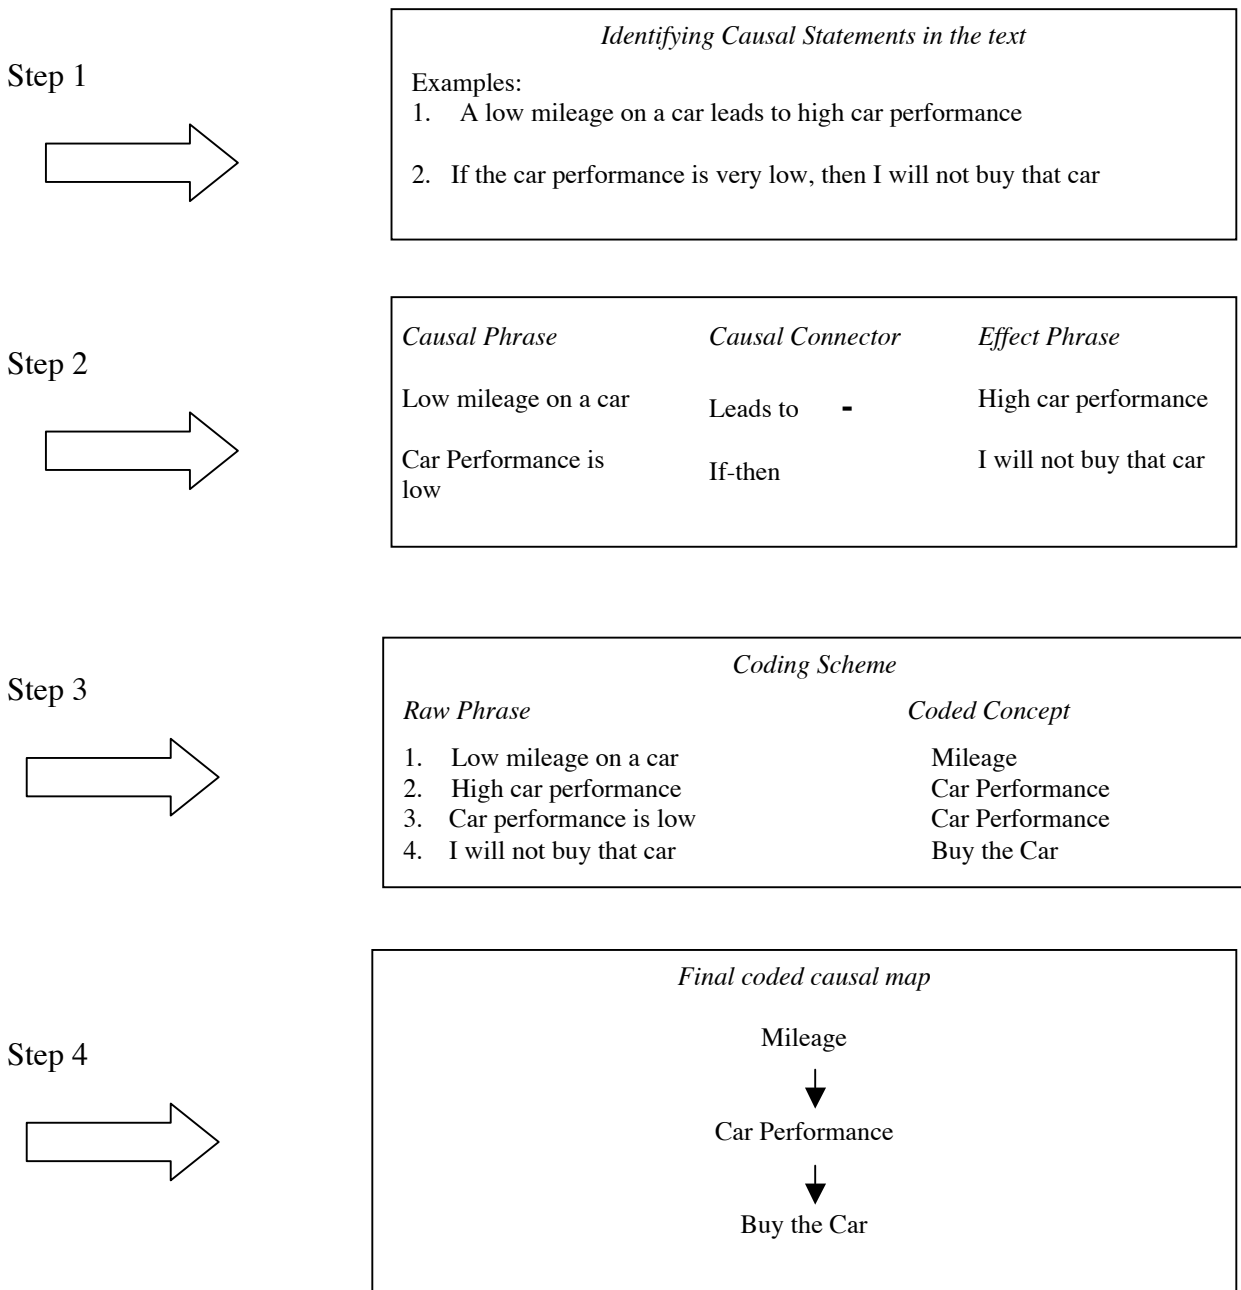


Figure 9. An Illustration of a Part of the Structured Interview

Directions: Please circle one of the four alternatives provided to specify the type of direct relation between the concepts listed below. Also circle the sign associated with the relation.

Mileage	None	<input checked="" type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	Performance
		+	-		
Mileage	None	<input type="checkbox"/>	<input checked="" type="checkbox"/>	<input type="checkbox"/>	Age
		+	-		
Mileage	None	<input checked="" type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	Fuel Efficiency
		+	-		
Mileage	<input checked="" type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	Brand Quality
		+	-		
Age	None	<input checked="" type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	Performance
		+	-		
Age	None	<input checked="" type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	Fuel Efficiency
		+	-		
Age	<input checked="" type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	Brand Quality
		+	-		
Fuel Efficiency	None	<input checked="" type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	Performance
		+	-		
Fuel Efficiency	None	<input type="checkbox"/>	<input checked="" type="checkbox"/>	<input type="checkbox"/>	Brand Quality
		+	-		
Brand Quality	None	<input checked="" type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	Performance
		+	-		

Figure 10. Illustration of An Adjacency Matrix

Causes Effects	Mileage	Age	Fuel Efficiency	Performance	Brand Quality
1. Mileage		+	0	0	0
2. Age	0		0	0	0
3. Fuel Efficiency	□	□		0	+
4. Performance	□	□	+		+
5. Brand Quality	0	0	0	0	

Figure 11. The Original Causal Map of a Systems Analyst for the Online Ticketing Application

Outsourcing Decision

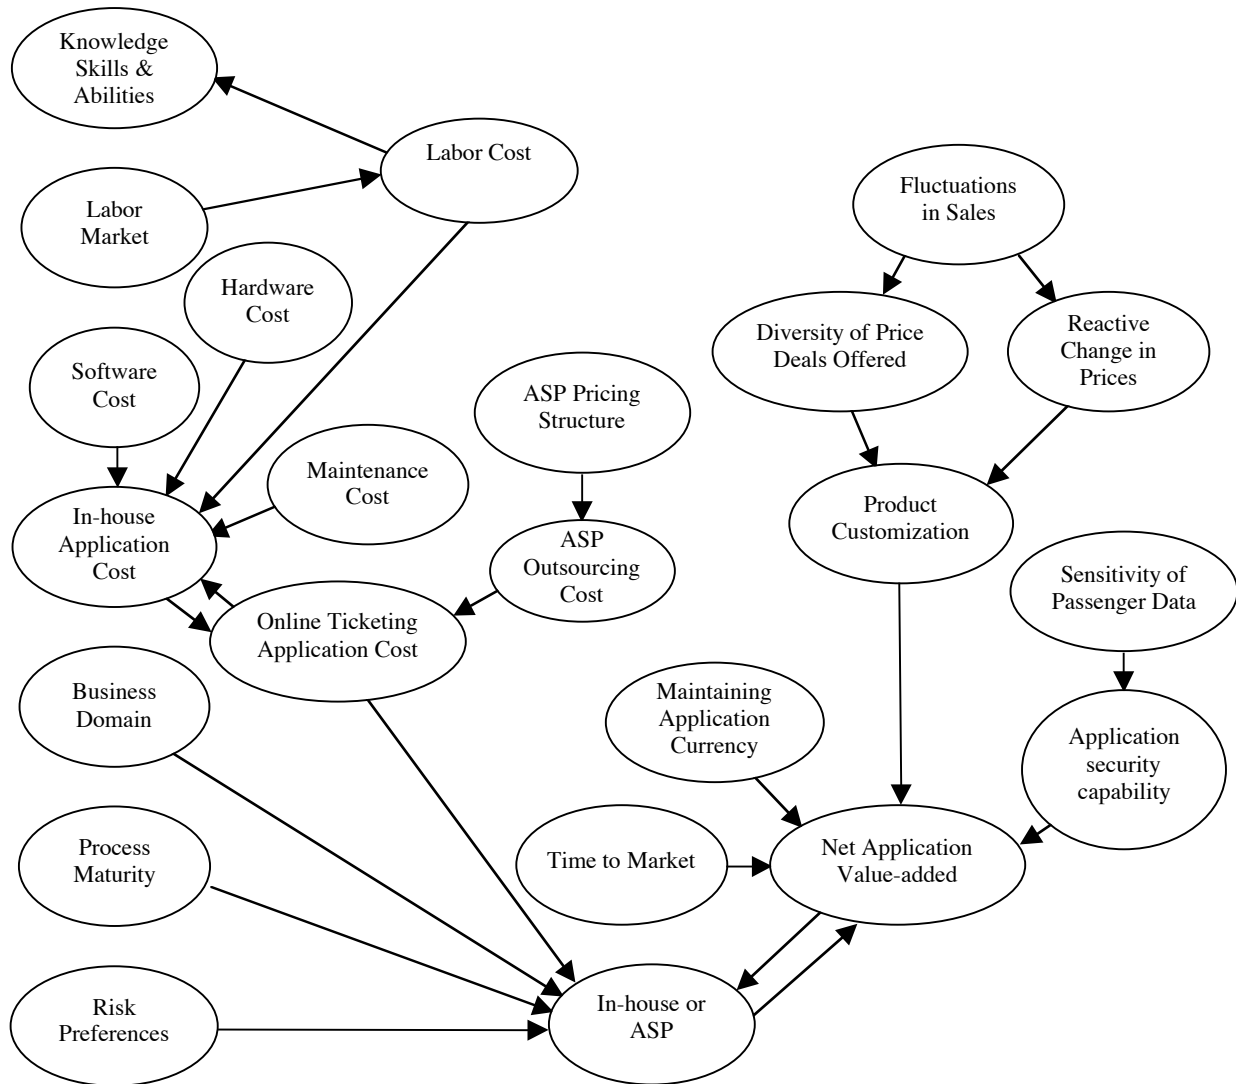


Figure 12. The Modified Causal Map of the Online Ticketing Application Decision

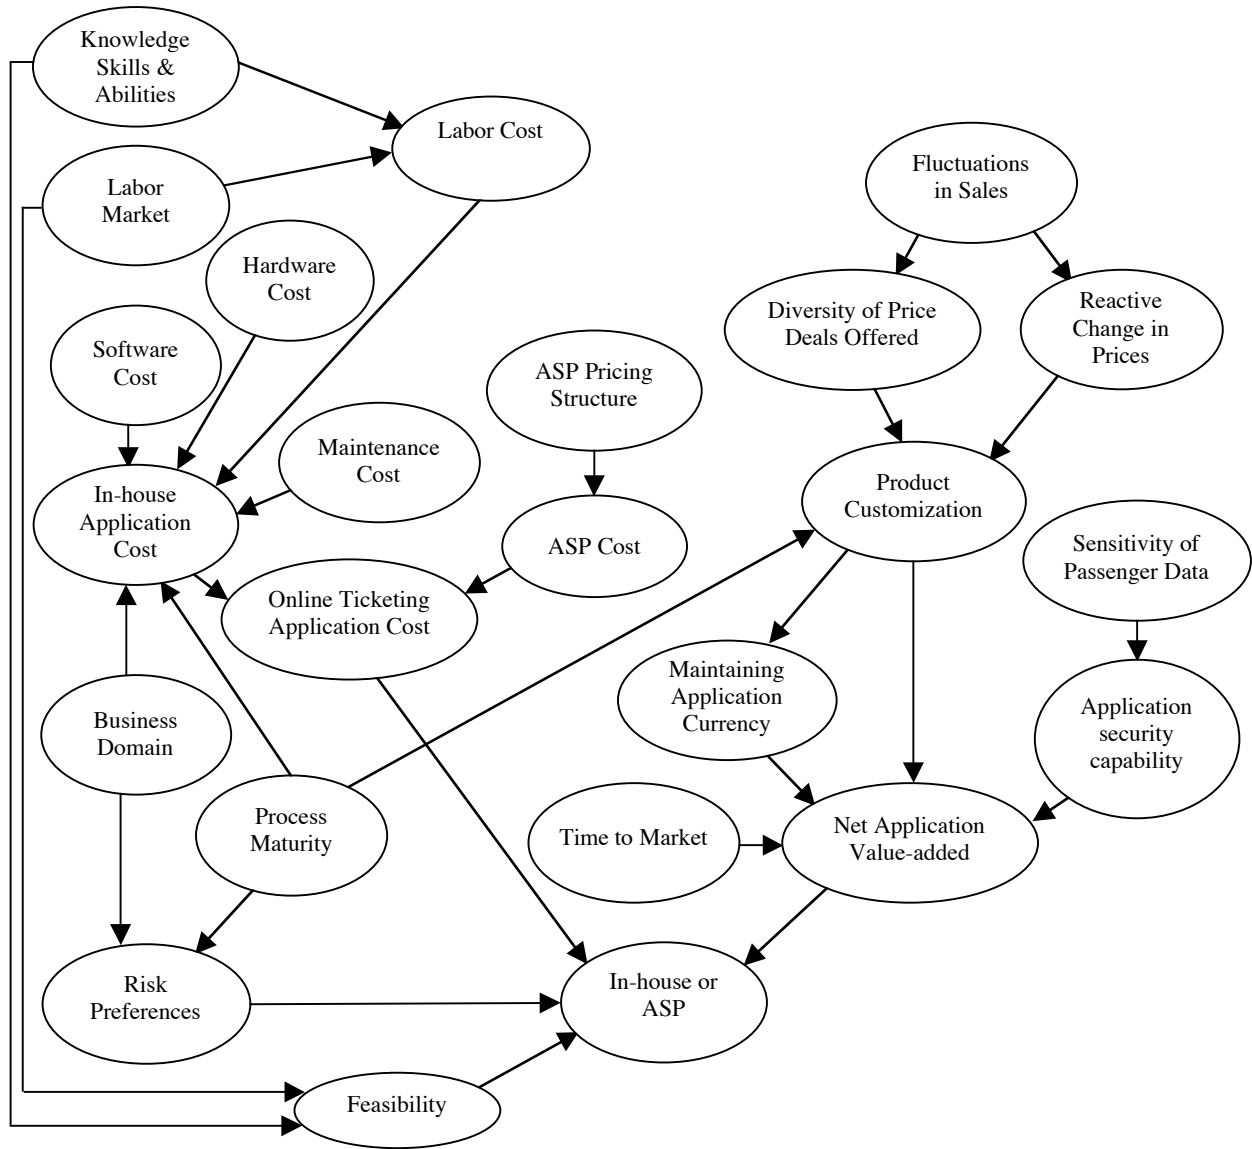


Table 1. Definition and States of Variables in the Causal Map

<i>Variable</i>	<i>Definition</i>	<i>States</i>
In-house Application Cost Variables		
1. Knowledge, Skills & Abilities (KSAs)	The set of skills required to successfully design, develop, and implement the online ticketing application in-house such as language skills, programming skills, database skills, process skills, and project management skills.	Broad, Narrow
2. Labor Market	The total number of IS persons available for hiring who have the knowledge, skills, and abilities relevant to the current application domain.	Large, Small
3. Labor Cost	The total cost of hiring individuals possessing the requisite KSAs for designing, developing, and implementing in-house application for inline ticketing.	High, Low
4. Software Cost	The total cost of purchasing and developing the software necessary to design, develop, and implement in-house application including off-the-shelf software modules, developing CASE Tools, and ancillary software.	High, Low
5. Hardware Cost	The total cost of buying the necessary hardware equipment for developing in-house application such as computers, workstations, and networking equipment.	High, Low
6. Maintenance Cost	The total cost of modifying, updating and system testing the necessary application hardware & software.	High, Low
7. In-house Application Cost	The sum total of software cost, hardware cost, maintenance cost, and labor cost.	Relatively High, Relatively Low
ASP Outsourcing Cost Variables		
8. ASP Price Structure	The form of fee such as a licensing fee charged by the application service providers (ASPs) for the use of their application.	High, Low
9. ASP Outsourcing Cost	The total cost incurred to outsource the application to an ASP and includes the fee charged by the ASP and any hardware or networking cost incurred.	Relatively High, Relatively Low
10. Online Ticketing Application Cost	A comparison of the cost of developing in-house application to the cost of outsourcing the application to an ASP to determine which alternative is better cost-wise.	In-house>ASP, ASP>In-house
Risk Determinant Variables		
11. Business Domain	Can be either primary to the value chain of the business or it may support the primary activities. Primary activities such as in-bound logistics, manufacturing and sales are central to the business. Supporting activities such as R&D, human resources indirectly affect value creation.	Primary, Supporting
12. Process Maturity	Whether the online ticketing process already exists in the company or whether it is totally new to the company.	Existing, New
13. Risk Preferences	The risk orientation of the decision-makers and/or corporate culture within the company in terms of willingness to take risk.	High risk, Low risk
14. Feasibility	The degree to which development of in-house online ticketing application is technically, legally, and organizationally feasible.	High, Low

Table 2. Definition and States of Variables in the Causal Map (continued from Table 1).

<i>Variable</i>	<i>Definition</i>	<i>States</i>
Application Environment Variables		
15. Fluctuation in Sales	The periodic variations in the sales, e.g., daily, weekly and monthly.	High, Low
16. Diversity of Price Deals Offered	The total number of and the complexity of the tariff rates offered by the airline company and the extent to which these are different from each other.	High, Low
17. Reactive Changes in Prices	The degree to which airline prices change in response to external environmental (e.g. change in fuel prices) and competitive forces.	Frequent, Non-frequent
18. Product Customization	Refers to the degree to which the online ticketing application needs to be adapted to and supportive of diversity of price deals offered by the airline company and the external environmental factors affecting changes in airline prices.	High, Low
19. Sensitivity of Passenger Data	The degree to which the customer data is sensitive to issues of privacy, confidentiality, transmission, storage, and security.	High, Low
20. Application Security Capability	The extent to which the application is capable of addressing the data security issues, e.g., restriction of the number of parties who have access to the data processed by and stored in the application.	High. Low
21. Maintaining Application Currency	The cost and level of effort required to maintain the currency of the application through revisions and upgrades.	High, Low
22. Time to Market	The quickness and agility with which the service can be marketed to the customer. The shorter the time to market, the more efficient the system.	Short, Long
23. Net Application Value-added (NVA)	Comparison of the net value-added (NVA) through either in-house development or ASP's application.	NVA In-house, NVA ASP
Decision Variables		
24. Online Ticketing Outsourcing Decision	The decision whether to develop in-house application or use the ASP's application.	In-house, ASP

Table 3. Prior and Posterior Marginal Probabilities under Two Different Scenarios

<i>Variable States</i>		<i>Prior Marginals</i>	<i>Posterior Marginals in Scenario 1</i>	<i>Posterior Marginals in Scenario 2</i>
1. Application Cost:	In-house > ASP	0.54	0.33	0.56
	In-house < ASP	0.46	0.67	0.44
2. Feasibility:	High	0.55	0.85	0.20
	Low	0.45	0.15	0.80
3. Net-Application Value-added:	In-house > ASP	0.61	0.66	0.42
	In-house < ASP	0.39	0.34	0.58
4. In-House/ASP Decision:	In-house	0.49	0.54	0.43
	ASP	0.51	0.46	0.57
<p><i>Scenario 1:</i> Knowledge, Skills & Abilities = low, Labor Market = high, ASP Pricing Structure = high, Fluctuations in Sales = high, Sensitivity of Passenger Data = high.</p> <p><i>Scenario 2:</i> Knowledge, Skills & Abilities = high, Labor Market = low, ASP Pricing Structure = low, Fluctuations in Sales = low, Sensitivity of Passenger Data = low.</p>				