

VALUATION-BASED SYSTEMS FOR BAYESIAN DECISION ANALYSIS

PRAKASH P. SHENOY

The University of Kansas, Lawrence, Kansas

(Received June 1990; revision received May 1991; accepted May 1991)

This paper proposes a new method for representing and solving Bayesian decision problems. The representation is called a valuation-based system and has some similarities to influence diagrams. However, unlike influence diagrams which emphasize conditional independence among random variables, valuation-based systems emphasize factorizations of joint probability distributions. Also, whereas influence diagram representation allows only conditional probabilities, valuation-based system representation allows all probabilities. The solution method is a hybrid of local computational methods for the computation of marginals of joint probability distributions and the local computational methods for discrete optimization problems. We briefly compare our representation and solution methods to those of influence diagrams.

The main goal of this paper is to propose a new method for representing and solving Bayesian decision problems. We propose a new representation of a decision problem called a *valuation-based system*. A graphical depiction of a valuation-based system is called a *valuation network*. Valuation networks are similar in some respects to influence diagrams. Like influence diagrams, valuation networks are a compact representation emphasizing qualitative features of decision problems. Also, like influence diagrams, valuation networks allow representation of decision problems without any preprocessing. However, there are some differences. Whereas influence diagrams emphasize conditional independence among random variables, valuation networks emphasize factorizations of joint probability distributions. Also, the representation method of influence diagrams allows only conditional probabilities. While conditional probabilities are readily available in pure causal models, they are not always readily available in other graphical models (see, for example, Darroch, Lauritzen and Speed 1980). The representation method of valuation-based systems is more general and allows direct representation of all probability models.

The solution method proposed is a hybrid of local computational methods for computations of the marginals of joint probability distributions and local

computational methods for discrete optimization. Local computational methods for computation of the marginals of joint probability distributions have been proposed by, e.g., Pearl (1988), Lauritzen and Spiegelhalter (1988), Shafer and Shenoy (1988, 1990), Jensen, Lauritzen and Oleson (1990), and Jensen, Olesen and Anderson (1990). Local computational methods for discrete optimization are also called nonserial dynamic programming (Bellman 1957, Bertele and Brioschi 1972). Viewed abstractly using the framework of valuation-based systems, these two local computational methods are actually quite similar. Shenoy and Shafer (1990) and Shenoy (1991b) show that the same three axioms justify the use of local computation in both cases.

Valuation-based systems are described in Shenoy (1989, 1991c). In valuation-based systems, we represent knowledge by functions called *valuations*. We draw inferences from such systems using two operations called *combination* and *marginalization*. Drawing inferences can be described simply as marginalizing all variables out of the joint valuation. The joint valuation is the result of combining all valuations. The framework of valuation-based systems is powerful enough to also include Dempster-Shafer theory of belief functions (Shenoy 1991c), Spohn's theory of epistemic beliefs (Shenoy 1991a, c),

Subject classifications: Decision analysis, theory: representation, and solution using local computation. Dynamic programming, Markov, finite state: solution using local computation. Networks/graphs: representation for decision, optimization and probabilistic inference problems.

Area of review: DECISION ANALYSIS, BARGAINING AND NEGOTIATION.

possibility theory (Dubois and Prade 1990), propositional logic (Shenoy 1990a), and constraint satisfaction problems (Shenoy and Shafer 1988).

Traditional methods of representing problems in decision theory are payoff matrices and decision trees. The payoff matrix representation has its origins in the work of von Neumann and Morgenstern (1953) on normal form games and was made popular by Savage (1950). It is a convenient representation in problems where there is one decision to be made and one common uncertainty for all acts. In multistage decision problems, a payoff matrix representation requires the enumeration of possible strategies and the computation of the joint distribution of all random variables. Each of these two tasks can be computationally intractable due to combinatorial explosion.

A decision tree is a more flexible and graphic representation tool especially for multistage decision problems. Decision trees have their genesis in the work of von Neumann and Morgenstern on extensive form games. One distinct advantage of decision trees is that an optimal strategy can be identified using dynamic programming methods without enumerating all possible strategies (Zermelo 1913, Bellman 1957, Raiffa and Schlaifer 1961). Raiffa and Schlaifer call the dynamic programming method for solving decision trees "averaging-out-and-folding-back."

Although decision trees are more expressive and computationally more efficient than payoff matrices, decision trees have several drawbacks. First, since decision trees explicitly represent acts and events, the trees grow too fast in many problems. Thus, an n -stage decision problem with m choices or events at each stage has at least m^n endpoints.

Second, probabilities of events may not be available in the form that decision tree methodology requires. In such cases, it is necessary to compute these probabilities from the available probabilities using the laws of probability theory. This is a major drawback of decision trees. There should be a cleaner way of separating a representation of a problem from its solution. The former is hard to automate while the latter is easy. Decision trees mix these two tasks making automation difficult. The problem of finding marginals from a factored representation of the joint distribution has received much attention in the literature on uncertainty in artificial intelligence (e.g., Pearl 1988, Shafer and Shenoy 1988, 1990, Lauritzen and Spiegelhalter 1988, Jensen, Olesen and Andersen 1990, and Jensen, Lauritzen and Olesen 1990).

Third, the decision tree representation demands conditional probability distributions at each random variable node. This demand often necessitates division

operations (in the preprocessing of probabilities) that may be unnecessary. The unnecessary divisions are compensated for by multiplications that neutralize the divisions. These unnecessary divisions and multiplications make the decision tree solution process inefficient.

Influence diagrams were initially proposed as an alternative to decision trees for represent decision problems (Miller et al. 1976, Howard and Matheson 1984). Subsequently, Olmsted (1983) and Shachter (1986) devised methods for solving influence diagrams directly (without having to convert them to decision trees). In the last decade, influence diagrams have become popular for representing and solving decision problems (Oliver and Smith 1990).

Influence diagrams do not share some of the drawbacks of decision trees mentioned above. First, since acts and events are not graphically depicted, influence diagrams do not grow as fast as decision trees. Second, users can input conditional probabilities directly in the form they are available without having to compute the posteriors. The computation of the posteriors is part of the process of solving influence diagrams (Olmsted 1983, Shachter 1986, Tatman 1986, Ezawa 1986, Tatman and Shachter 1990). The process of solving influence diagrams involves arc reversals and node removals. Although the process of solving influence diagrams is more complex than the process of solving decision trees, the process can easily be automated (Shachter 1988).

The representation method of influence diagrams allows only conditional probabilities. While conditional probabilities are readily available in pure causal models, they are not always available in other graphical models (see, e.g., Darroch, Lauritzen and Speed 1980, Wermuth and Lauritzen 1983, Edwards and Kreiner 1983, Kiiveri, Speed and Carlin 1984). In such cases, the probabilities have to be preprocessed before they can be represented in influence diagrams. The representation method of valuation-based systems is more general—all probability models can be represented directly without any preprocessing.

The arc reversal operation in influence diagrams involves unnecessary divisions. As in decision trees, these unnecessary operations are the result of influence diagram representation requirements that allow only conditional probabilities. The solution method for valuation-based systems described in this paper involves minimal division. If there is only one joint utility function (in unfactored form), then our method involves no divisions. This assumption of one joint utility function is similar to the assumption in influence diagrams that there is only one value node (see,

e.g., Shachter 1986). Also, if we have a factorization of the joint probability distributions into conditionals and the ordering of the variables is consistent in some sense with the information constraints, then again, our method involves no divisions. Finally, if the combination of valuations is associative, as, for example, it is if the joint utility function factors multiplicatively, then again no divisions are necessary. In general, if we want to take computational advantage of an additive factorization of the joint utility function, then divisions may be necessary.

An outline of this paper is as follows. In Section 1, we give a complete statement, decision tree representation, and decision tree solution of the oil wildcatter's problem (Raiffa 1968). We use the oil wildcatter's problem to illustrate all definitions in the paper. In Section 2, we describe valuation-based systems specifically designed for the representation and solution of Bayesian decision problems. In Section 3, we define combination and marginalization operations used for solving valuation-based systems. In Section 4, we describe the semantics of a valuation-based system representation, and the conditions under which such a representation is well defined. In Section 5, we describe a fusion algorithm for solving valuation-based systems using local computation. In Section 6, we compare valuation-based systems to influence diagrams. In Section 7, we make some concluding remarks. Finally, Section 8 contains a proof of the main theorem in the paper.

1. THE OIL WILDCATTER'S PROBLEM

The oil wildcatter's problem is reproduced with minor modifications from Raiffa. An oil wildcatter must decide either to drill (d) or not drill ($\sim d$). He is uncertain whether the hole is dry (dr), wet (we) or soaking (so). Table I gives his monetary payoffs and his subjective probabilities of the various states. The cost of drilling is \$70,000. The net return associated with the d - we pair is \$50,000, which is interpreted as a return of \$120,000 less the \$70,000 cost of drilling. Similarly, the \$200,000 associated with the d - so pair

Table I
The Payoff Matrix for the Oil Wildcatter's Problem

State and Wildcatter's Profit, \$ (π)	Act		Probability of State
	Drill (d)	Not Drill ($\sim d$)	
Dry (dr)	-70,000	0	0.500
Wet (we)	50,000	0	0.300
Soaking (so)	200,000	0	0.200

Table II

Probabilities of Seismic Test Results Conditional on the Amount of Oil

Amount of Oil (O) $P(R O)$	Seismic Test Results (R)		
	No Structure (ns)	Open Structure (os)	Closed Structure (cs)
Dry (dr)	0.600	0.300	0.100
Wet (we)	0.300	0.400	0.300
Soaking (so)	0.100	0.400	0.500

is a net return (a return of \$270,000 less the \$70,000 cost of drilling).

At a cost of \$10,000, the wildcatter could take seismic soundings which will help determine the geological structure at the site. The soundings will disclose whether the terrain below has no structure (ns)—(that's bad), or an open structure (os) (that's so-so), or a closed structure (cs) (that's really hopeful). The experts have provided us with Table II which shows the probabilities of seismic test results conditional on the amount of oil.




Figure 1. The preprocessing of probabilities in the oil wildcatter's problem.

Figures 1 and 2 show a decision tree representation and solution of this problem. Note that even before the decision tree can be specified completely the conditional probabilities in the tree have to be computed from those specified in the problem, as in Figure 1. In Figure 1, the probability tree on the left is used to compute the joint probabilities, and the probability tree on the right is used to compute the marginals for test results and the conditional probabilities of the amount of oil given test results. As we will see later, all computations in the probability tree on the right are unnecessary.

Figure 2 shows the solution of the oil wildcatter's problem. The optimal strategy is to do a seismic test, not drill if the seismic test reveals no structure, and

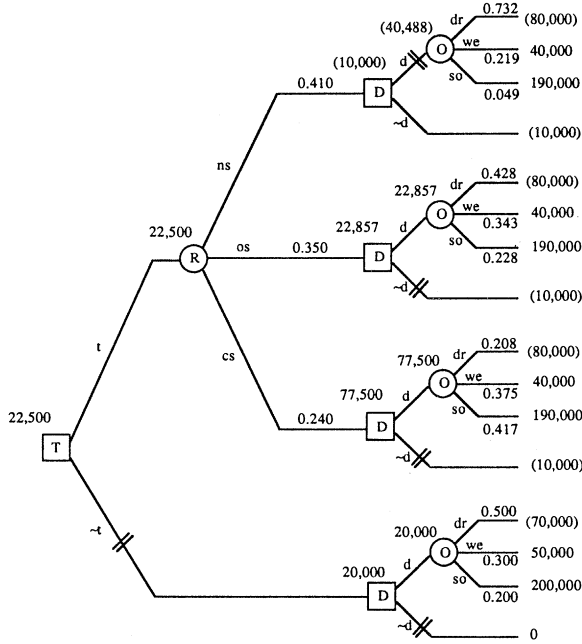


Figure 2. A decision tree representation and solution of the oil wildcatter's problem.

drill if the seismic test reveals either open or closed structure. The expected profit associated with this strategy is \$22,500.

2. VALUATION-BASED SYSTEM REPRESENTATION

A valuation-based system representation of decision problems uses decision variables, random variables, frames, payoff valuations, potentials, and precedence constraints. We will discuss each of these in detail. A graphical representation of these objects is called a *valuation network*. Figure 3 shows a valuation network for the oil wildcatter's problem.

2.1. Variables, Frames and Configurations

A decision node is represented as a variable. The possible values of a decision variable represent the acts available at that point. We use the symbol \mathcal{W}_D for the set of possible values of decision variable D . We assume that the decision-maker has to pick one and only one of the elements of \mathcal{W}_D as their decision. We call \mathcal{W}_D the *frame* for D . Decision variables are represented in a valuation network by rectangular nodes.

In the oil wildcatter's problem, there are two decision nodes D and T . The frame for D has two acts: drill (d), and not drill ($\sim d$). The frame for T also has

two elements: do a seismic test (t), and not do a seismic test ($\sim t$).

If R is a random variable, we use the symbol \mathcal{W}_R to denote its possible values. We assume that one and only one of the elements of \mathcal{W}_R can be the true value of R . We call \mathcal{W}_R the *frame* for R . Random variables are denoted in valuation networks by circular nodes.

In the oil wildcatter's problem, there are two random variables: amount of oil (O) and seismic test results (R). The frame for O has three elements: dry (dr), wet (we), and soaking (so). The frame for R has four elements: no result (nr), no structure (ns), open structure (os), and closed structure (cs).

Let \mathcal{X}_D denote the set of all decision variables, let \mathcal{X}_R denote the set of all random variables, and let $\mathcal{X} = \mathcal{X}_D \cup \mathcal{X}_R$ denote the set of all variables. In this paper, we are concerned only with the case where \mathcal{X} is finite. We also assume that all the variables in \mathcal{X} have finite frames. We use upper-case Roman alphabets to denote variables.

We often deal with nonempty subsets of variables in \mathcal{X} . Given a nonempty subset h of \mathcal{X} , let \mathcal{W}_h denote the Cartesian product of \mathcal{W}_X for X in h , i.e., $\mathcal{W}_h = \times\{\mathcal{W}_X \mid X \in h\}$. We can think of \mathcal{W}_h as the set of possible values of the joint variable h . Accordingly, we call \mathcal{W}_h the *frame* for h . Also, we refer to elements of \mathcal{W}_h as *configurations* of h . We use this terminology even when h consists of a single variable, say X . Thus, we refer to elements of \mathcal{W}_X as *configurations* of X . We use lower case, boldface letters, such as x and y , to denote configurations. Also, if x is a configuration of g , y is a configuration of h , and $g \cap h = \emptyset$, then (x, y) denotes a configuration of $g \cup h$.

It is convenient to extend this terminology to the case where the set of variables h is empty. We adopt the convention that the frame for the empty set \emptyset consists of a single configuration, and we use the symbol \diamond to name that configuration; $\mathcal{W}_\emptyset = \{\diamond\}$. To




Figure 3. A valuation network for the oil wildcatter's problem.

be consistent with this notation, we adopt the convention that if \mathbf{x} is a configuration for g , then $(\mathbf{x}, \blacklozenge) = \mathbf{x}$.

2.2. Valuations

Suppose that $h \subseteq \mathcal{X}$. A *payoff valuation* π for h is a function from \mathcal{W}_h to \mathbb{R} , where \mathbb{R} denotes the set of real numbers. The values of payoff valuations are consequences, for example, utilities, profit, and cost. If $h = d \cup r$, where $d \subseteq \mathcal{X}_D$ and $r \subseteq \mathcal{X}_R$, $\mathbf{x} \in \mathcal{W}_d$, and $\mathbf{y} \in \mathcal{W}_r$, then $\pi(\mathbf{x}, \mathbf{y})$ denotes a partial payoff to the decision maker if the decision maker chooses configuration \mathbf{x} and the true configuration of r is \mathbf{y} . The exact meaning of a payoff valuation is given in Section 4. If π is a payoff valuation for h and $X \in h$, then we say that π *bears on* X .

In a valuation network, a payoff valuation is represented by a diamond-shaped node. To permit the identification of all valuations that bear on a variable, we draw undirected edges between the payoff valuation node and all the variable nodes it bears on. In the oil wildcatter's problem, there are two payoff valuations π and κ , as shown in Figure 3. Table III shows the details of these valuations.

Suppose that $h \subseteq \mathcal{X}$ such that $h \cap \mathcal{X}_R \neq \emptyset$. A *potential* ρ for h is a function from \mathcal{W}_h to the unit interval $[0, 1]$. The values of potentials are probabilities. The exact meaning of a potential is given in Section 4.

In a valuation network, a potential is represented by a triangular node. Again, to permit the identification of valuations that bear on a variable, we draw undirected edges between a potential node and all the variables nodes it bears on.

Suppose that $h \subseteq \mathcal{X}$, R is a random variable in h , and ρ is a potential for h . If all the values of ρ represent conditional probabilities for R given configurations of the variables in $h - \{R\}$, then we call ρ a *conditional potential for R given $h - \{R\}$* . In this case, ρ satisfies

$$\sum \{\rho(\mathbf{c}, \mathbf{r}) \mid \mathbf{r} \in \mathcal{W}_R\} = 1 \text{ for all } \mathbf{c} \in \mathcal{W}_{h-\{R\}}. \quad (1)$$

Table III

Payoff Valuations in the Oil Wildcatter's Problem

$\mathcal{W}_{\{D,O\}}$	π	$\mathcal{W}_{\{T\}}$	κ
d dr	-70,000	t	-10,000
d we	50,000	$\sim t$	0
d so	200,000		
$\sim d$ dr	0		
$\sim d$ we	0		
$\sim d$ so	0		

Table IV

Potentials in the Oil Wildcatter's Problem

\mathcal{W}_O	ρ	$\mathcal{W}_{\{T,R,O\}}$	μ^a
dr	0.500	t ns dr	0.600
we	0.300	t ns we	0.300
so	0.200	t ns so	0.100
		t os dr	0.300
		t os we	0.400
		t os so	0.400
		t cs dr	0.100
		t cs we	0.300
		t cs so	0.500
		$\sim t$ nr dr	1.000
		$\sim t$ nr we	1.000
		$\sim t$ nr so	1.000

^a Only nonzero probability values of μ are shown.

In a valuation network, a conditional potential ρ for R given $h - \{R\}$ is represented by making the edge between the potential node ρ and the variable node R directed (pointed toward R). In the oil wildcatter's problem, there are two potentials ρ and μ , as shown in Figure 3. Table IV shows the details of these potentials. Note that μ is a conditional potential for R given $\{T, O\}$, and ρ is a conditional potential for O given \emptyset .

In a valuation network, both variables and valuations are represented by nodes. To help the reader keep these objects separate, we use upper case Roman alphabets to label variables and lower case Greek alphabets to label valuations.

Let \mathcal{H}_D denote the set of subsets of \mathcal{X} for which payoff valuations exist in the valuation-based system. For simplicity of exposition, we assume that each decision variable in \mathcal{X}_D is included in some element of \mathcal{H}_D , i.e.,

$$\cup \mathcal{H}_D \supseteq \mathcal{X}_D. \quad (2)$$

Let \mathcal{H}_R denote the set of subsets of \mathcal{X} for which potentials exist. We assume that each random variable is included in some element of \mathcal{H}_R , i.e.,

$$\cup \mathcal{H}_R \supseteq \mathcal{X}_R. \quad (3)$$

$\mathcal{H} = \mathcal{H}_D \cup \mathcal{H}_R$. From (2) and (3), it is clear that $\cup \mathcal{H} = \mathcal{X}$. Using the language of graph theory, the set \mathcal{H} is called a *hypergraph on \mathcal{X}* , and each of its elements is called a *hyperedge*. We call \mathcal{H}_D the *payoff hypergraph* and \mathcal{H}_R the *potential hypergraph*.

2.3. Precedence Constraints

Besides acts, events, probabilities and payoffs, an important ingredient of problems in decision analysis

is information constraints. Some decisions have to be made before the observation of some uncertain events, and some decisions can be postponed until some events are observed. In the oil wildcatter's problem, for example, the amount of oil is revealed only after the ground is drilled or perhaps it is never revealed; the decision whether to drill or not may be postponed until the seismic test result is revealed.

If a decision-maker expects to be informed of the true value of random variable R before they make a decision D , then we represent this situation by the binary relation $R \rightarrow D$ (read as R precedes D). On the other hand, if a random variable R is only revealed after a decision D is made or perhaps never revealed, then we represent this situation by the binary relation $D \rightarrow R$. It is possible that in some problems, we may have precedence constraints between two decision nodes or between two random variable nodes. For example, if random variable R_2 is only revealed after random variable R_1 is revealed, we represent this by the relation $R_1 \rightarrow R_2$.

In the oil wildcatter's problem, we have the precedence constraints $T \rightarrow R$, $R \rightarrow D$, $D \rightarrow O$. The seismic test result (R) is only revealed after we make the decision to do a seismic test or not (T). The decision to drill or not to drill (D) is only made after observing the test results (R). Finally, the amount of oil (O) is only revealed if we make the decision to drill (D). Obviously, these are not the only constraints. In the oil wildcatter's problem we also have, for example, $T \rightarrow O$. But such constraints can be deduced, and there is no reason to include all such constraints in the representation. On the other hand, the problem can incorrectly be overconstrained, permitting no solution. For example, if $T \rightarrow R$ and $R \rightarrow T$, then this will preclude a solution. Therefore, we do not permit such precedence constraints.

What restrictions do we need to impose on the precedence relation \rightarrow ? We require four conditions. First, we require that the transitive closure of \rightarrow , denoted by $>$, is a partial order on \mathcal{X} . We call this first condition the *partial order condition*. Second, we require that this partial order $>$ is such that for any $D \in \mathcal{X}_D$ and any $R \in \mathcal{X}_R$, either $D > R$ or $R > D$. We call this second condition the *perfect recall condition*. (This terminology is borrowed from Kuhn (1953). In the influence diagram literature, this condition is called the "no-forgetting assumption" (Howard and Matheson 1984).) Third, if R is a random variable such that there exists a conditional potential for R given $h - R$, and if D is a decision variable in h , then we require that $D > R$. Fourth, if D is a decision variable and there exists a potential for h such

that $D \in h$, then we require that $D > R$ for some random variable $R \in h$. We call the third and fourth conditions *consistency conditions*.

Before we explain the reasons for these three conditions, let us explain the terms transitive closure and partial ordering. The *transitive closure* of \rightarrow is defined as:

- i. $X > Y$ whenever $X \rightarrow Y$, and
- ii. $X > Y$ whenever there exists a $Z \in \mathcal{X}$ such that $X > Z$ and $Z > Y$.

A binary relation $>$ is a *partial order* on \mathcal{X} if it satisfies the following two properties:

- i. (Irreflexive): $X > X$ for no $X \in \mathcal{X}$.
- ii. (Transitive): $X > Y$ and $Y > Z$ imply $X > Z$.

The reason for the partial order requirement is obvious. The reason for the perfect recall condition is as follows. Given the meaning of the precedence relation \rightarrow , for any decision variable D and any random variable R , either R is known when decision D has to be made, or not. This translates to either $R > D$ or $D > R$. Finally, the consistency conditions are dictated by the meaning of potentials. If the conditional probability distribution of random variable R depends on the act chosen by the decision maker at node D , then it must be the case that $D > R$.

In summary, a valuation-based system representation of a decision problem consists of a finite set of decision variables \mathcal{X}_D , a finite set of random variables \mathcal{X}_R , a finite frame \mathcal{W}_X for each variable X in $\mathcal{X}_D \cup \mathcal{X}_R$, a finite collection of payoff valuations $\{\pi_1, \dots, \pi_m\}$, a finite collection of potentials $\{\rho_1, \dots, \rho_n\}$, and a precedence relation \rightarrow on $\mathcal{X}_D \cup \mathcal{X}_R$. Thus, a *valuation-based system* (VBS) can be denoted formally by the 6-tuple

$$\{\mathcal{X}_D, \mathcal{X}_R, \{\mathcal{W}_X\}_{X \in \mathcal{X}}, \{\pi_1, \dots, \pi_m\}, \{\rho_1, \dots, \rho_n\} \rightarrow\}$$

representing decision variables, random variables, frames, payoff valuations, potentials, and the precedence relation, respectively.

3. COMBINATION AND MARGINALIZATION

In this section, we define two operations called combination and marginalization. We use these operations to solve valuation-based systems. Precise definitions require extensive notation that will test the patience of the reader. For relief, we use the oil wildcatter's problem to illustrate all definitions. First, we start with some notation we need to define combination and marginalization.

3.1. Projection of Configurations

Projection of configurations simply means dropping extra coordinates; if (w, x, y, z) is a configuration of $\{W, X, Y, Z\}$, for example, then the projection of (w, x, y, z) to $\{W, X\}$ is simply (w, x) , which is a configuration of $\{W, X\}$.

If g and h are sets of variables, $h \subseteq g$, and \mathbf{x} is a configuration of g , then let \mathbf{x}^{lh} denote the projection of \mathbf{x} to h . The projection \mathbf{x}^{lh} is always a configuration of h . If $h = g$ and \mathbf{x} is a configuration of g , then $\mathbf{x}^{lh} = \mathbf{x}$. If $h = \emptyset$, then $\mathbf{x}^{lh} = \diamond$.

3.2. Combination

The definition of combination depends on the type of valuations being combined. Suppose that h and g are subsets of \mathcal{X} , π_i is a payoff valuation for h , and π_j is a payoff valuation of g . Then the *combination* of π_i and π_j , denoted by $\pi_i \otimes \pi_j$, is a payoff valuation for $h \cup g$ defined as

$$(\pi_i \otimes \pi_j)(\mathbf{x}) = \pi_i(\mathbf{x}^{lh}) + \pi_j(\mathbf{x}^{lg}) \quad (4)$$

for all $\mathbf{x} \in \mathcal{W}_{h \cup g}$.

Suppose that h and g are subsets of \mathcal{X} , π_i is a payoff valuation of h , and ρ_j is a potential for g . Then the *combination* of π_i and ρ_j , denoted by $\pi_i \otimes \rho_j$, is a payoff valuation for $h \cup g$ defined as

$$(\pi_i \otimes \rho_j)(\mathbf{x}) = \pi_i(\mathbf{x}^{lh})\rho_j(\mathbf{x}^{lg}) \quad (5)$$

for all $\mathbf{x} \in \mathcal{W}_{h \cup g}$. And the *combination* of ρ_j and π_i , denoted by $\rho_j \otimes \pi_i$, is a payoff valuation for $h \cup g$ defined as

$$\rho_j \otimes \pi_i = \pi_i \otimes \rho_j. \quad (6)$$

Suppose that h and g are subsets of \mathcal{X} , ρ_i is a potential for h , and ρ_j is a potential for g . Then the *combination* of ρ_i and ρ_j , denoted by $\rho_i \otimes \rho_j$, is a potential for $h \cup g$ defined as

$$(\rho_i \otimes \rho_j)(\mathbf{x}) = \rho_i(\mathbf{x}^{lh})\rho_j(\mathbf{x}^{lg}) \quad (7)$$

for all $\mathbf{x} \in \mathcal{W}_{h \cup g}$.

First, note that the combination of two payoff valuations is a payoff valuation, the combination of two potentials is a potential, and the combination of a payoff valuation and a potential is a payoff valuation. This is consistent with the units of the values of the valuations because values of potentials are probabilities, which are dimensionless quantities.

Second, note that a combination of two payoff valuations consists of pointwise addition. This assumes that the joint payoff function factors additively as in the oil wildcatter's problem. (If it factored multiplicatively, we would have defined this combi-

nation as pointwise multiplication.) Combination of two potentials consists of pointwise multiplication. Probability theory mandates this. Finally, combination of a payoff valuation and a potential consists of pointwise multiplication. This is one of the operations in computing expected payoffs. (The other operation, summation, is part of the definition of marginalization given below.)

Third, note that combination is commutative. Also, the combination for payoff valuations is associative. Thus, if $\{\pi_1, \dots, \pi_k\}$ is a set of payoff valuations, we write $\otimes\{\pi_1, \dots, \pi_k\}$ to mean the combination of valuations in $\{\pi_1, \dots, \pi_k\}$ in some sequence. Similarly, the combination of potentials is associative, and if $\{\rho_1, \dots, \rho_k\}$ is a set of potentials, we write $\otimes\{\rho_1, \dots, \rho_k\}$ to mean the combination of potentials in $\{\rho_1, \dots, \rho_k\}$ in some sequence. In general, the combination for a mixture of payoff valuations and potentials is not associative, e.g., $(\rho_1 \otimes \pi_1) \otimes \pi_2 \neq \rho_1 \otimes (\pi_1 \otimes \pi_2)$. In such cases, what sequence should we use? We define combination such that the payoff valuations are combined before combining the potentials. Formally, suppose that π_1, \dots, π_k are payoff valuations, and ρ_1, \dots, ρ_j are potentials. Then let $\otimes\{\pi_1, \dots, \pi_k, \rho_1, \dots, \rho_j\}$ denote $(\otimes\{\pi_1, \dots, \pi_k\}) \otimes (\otimes\{\rho_1, \dots, \rho_j\})$. Finally, note that if we have only one payoff valuation and several potentials, then combination is associative. We can assume that we have only one payoff valuation by combining all payoff valuations before we combine the potentials.

Tables V and VI illustrate the combination operation using valuations from the oil wildcatter's problem.

3.3. Vacuous Potentials

Suppose that ρ is a potential for g . We say that ρ is *vacuous* if $\rho(\mathbf{r}) = 1$ for all $\mathbf{r} \in \mathcal{W}_g$. It follows from the definition of combination, that if μ is a potential for g and ρ is a vacuous potential for g , then $\rho \otimes \mu = \mu$.

3.4. Marginalization

Suppose that h and g are subsets of variables, and g is a subset of h . Marginalization is an operation when we reduce a valuation for h to a valuation for g by eliminating variables in $h - g$. Unlike combination, the definition of marginalization does not depend on the type of valuation being marginalized. But the definition of marginalization depends on the type of variables being eliminated. If the variable being eliminated is random, marginalization is achieved by summing the valuation over the frame of the eliminated variable. In this case, the valuation being marginalized may be either a payoff valuation or a potential. If the

Table V
Combination of Payoff Valuations π and κ ,
and Potentials ρ and μ

$\mathcal{W}_{\{T,D,O\}}$			π	κ	$\pi \otimes \kappa$
t	d	dr	-70,000	-10,000	-80,000
t	d	we	50,000	-10,000	40,000
t	d	so	200,000	-10,000	190,000
t	$\sim d$	dr	0	-10,000	-10,000
t	$\sim d$	we	0	-10,000	-10,000
t	$\sim d$	so	0	-10,000	-10,000
$\sim t$	d	dr	-70,000	0	-70,000
$\sim t$	d	we	50,000	0	50,000
$\sim t$	d	so	200,000	0	200,000
$\sim t$	$\sim d$	dr	0	0	0
$\sim t$	$\sim d$	we	0	0	0
$\sim t$	$\sim d$	so	0	0	0
$\mathcal{W}_{\{T,R,O\}}$			ρ	μ	$\rho \otimes \mu$
t	ns	dr	0.500	0.600	0.300
t	ns	we	0.300	0.300	0.090
t	ns	so	0.200	0.100	0.020
t	os	dr	0.500	0.300	0.150
t	os	we	0.300	0.400	0.120
t	os	so	0.200	0.400	0.080
t	cs	dr	0.500	0.100	0.050
t	cs	we	0.300	0.300	0.090
t	cs	so	0.200	0.500	0.100
$\sim t$	nr	dr	0.500	1.000	0.500
$\sim t$	nr	we	0.300	1.000	0.300
$\sim t$	nr	so	0.200	1.000	0.200

variable being eliminated is a decision variable, marginalization is achieved by maximization (or minimization depending on the nature of the values of the payoff valuations) over the frame of the eliminated variable. In this case, the valuation being marginalized must be a payoff valuation. A formal definition of marginalization is as follows.

Suppose that h is a subset of \mathcal{X} containing random variable R , and α is a valuation for h . The *marginal of α for $h - \{R\}$* , denoted by $\alpha^{\downarrow(h-\{R\})}$, is a valuation for $h - \{R\}$ defined as

$$\alpha^{\downarrow(h-\{R\})}(\mathbf{c}) = \sum \{\alpha(\mathbf{c}, \mathbf{r}) \mid \mathbf{r} \in \mathcal{W}_R\} \quad (8)$$

for all $\mathbf{c} \in \mathcal{W}_{h-\{R\}}$.

Suppose that h is a subset of \mathcal{X} containing decision variable D , and α is a payoff valuation for h . The *marginal of α for $h - \{D\}$* , denoted by $\alpha^{\downarrow(h-\{D\})}$, is a valuation for $h - \{D\}$ defined as

$$\alpha^{\downarrow(h-\{D\})}(\mathbf{c}) = \text{MAX}\{\alpha(\mathbf{c}, \mathbf{d}) \mid \mathbf{d} \in \mathcal{W}_D\} \quad (9)$$

for all $\mathbf{c} \in \mathcal{W}_{h-\{D\}}$.

Some observations: If α is a payoff valuation, then (8) corresponds to the second operation in “averaging out” a random variable. (The first operation is included in the combination operation.) On the other hand, if α is a potential, then (8) simply represents the familiar marginalization operation in probability theory. Table VI shows an example of marginalizing a random variable.

Condition 9 assumes that the nature of the values of a payoff valuation is such that the decision-maker’s objective is to maximize it. This is true if the payoff values represent utility, profits, probability of success, etc. On the other hand, in some problems, the decision-maker may wish to minimize the payoff values, when they represent disutility, cost, probability of failure, etc. In this case, we need to substitute MIN for MAX in (9). Condition 9 corresponds to “folding-back.” Table VII shows an example of marginalizing a decision variable.

We now state three lemmas regarding the marginalization operation. Lemma 1 states that in marginalizing two decision variables out of a valuation, the order in which the variables are eliminated does not affect the final result. Lemma 2 states a similar result for marginalizing two random variables out of a valuation. Lemma 3 states that in marginalizing a decision variable and a random variable out of a valuation, the order in which the two variables are eliminated may make a difference.

Lemma 1. Suppose that h is a subset of \mathcal{X} containing decision variables D_1 and D_2 , and α is a payoff valuation for h . Then

$$(\alpha^{\downarrow(h-\{D_1\})})^{\downarrow(h-\{D_1,D_2\})}(\mathbf{c}) = (\alpha^{\downarrow(h-\{D_2\})})^{\downarrow(h-\{D_1,D_2\})}(\mathbf{c})$$

for all $\mathbf{c} \in \mathcal{W}_{h-\{D_1,D_2\}}$.

Proof. The proof follows trivially from the definition of marginalization.

Lemma 2. Suppose that h is a subset of \mathcal{X} containing random variables R_1 and R_2 , and α is a valuation for h . Then

$$(\alpha^{\downarrow(h-\{R_1\})})^{\downarrow(h-\{R_1,R_2\})}(\mathbf{c}) = (\alpha^{\downarrow(h-\{R_2\})})^{\downarrow(h-\{R_1,R_2\})}(\mathbf{c})$$

for all $\mathbf{c} \in \mathcal{W}_{h-\{R_1,R_2\}}$.

Proof. The proof follows trivially from the definition of marginalization.

Lemma 3. Suppose that h is a subset of \mathcal{X} containing decision variable D and random variable R , and α is

Table VI
Computation of $(\pi \otimes \kappa) \otimes \rho \otimes \mu$ and $((\pi \otimes \kappa) \otimes \rho \otimes \mu)^{\downarrow\{T,R,D\}}$

$\mathcal{W}_{\{T,R,D,O\}}$				$\pi \otimes \kappa$	$\rho \otimes \mu$	$(\pi \otimes \kappa) \otimes \rho \otimes \mu$	$((\pi \otimes \kappa) \otimes \rho \otimes \mu)^{\downarrow\{T,R,D\}}$
t	ns	d	dr	-80,000	0.300	-24,000	-16,600
t	ns	d	we	40,000	0.090	3,600	
t	ns	d	so	190,000	0.020	3,800	
t	ns	$\sim d$	dr	-10,000	0.300	-3,000	-4,100
t	ns	$\sim d$	we	-10,000	0.090	-900	
t	ns	$\sim d$	so	-10,000	0.020	-200	
t	os	d	dr	-80,000	0.150	-12,000	8,000
t	os	d	we	40,000	0.120	4,800	
t	os	d	so	190,000	0.080	15,200	
t	os	$\sim d$	dr	-10,000	0.150	-1,500	-3,500
t	os	$\sim d$	we	-10,000	0.120	-1,200	
t	os	$\sim d$	so	-10,000	0.080	-800	
t	cs	d	dr	-80,000	0.050	-4,000	18,600
t	cs	d	we	40,000	0.090	3,600	
t	cs	d	so	190,000	0.100	19,000	
t	cs	$\sim d$	dr	-10,000	0.050	-500	-2,400
t	cs	$\sim d$	we	-10,000	0.090	-900	
t	cs	$\sim d$	so	-10,000	0.100	-1,000	
$\sim t$	nr	d	dr	-70,000	0.500	-35,000	20,000
$\sim t$	nr	d	we	50,000	0.300	15,000	
$\sim t$	nr	d	so	200,000	0.200	40,000	
$\sim t$	nr	$\sim d$	dr	0	0.500	0	0
$\sim t$	nr	$\sim d$	we	0	0.300	0	
$\sim t$	nr	$\sim d$	so	0	0.200	0	

Table VII
Computation of $\tau^{\downarrow\{T,R\}}$, Ψ_D , $\tau^{\downarrow\{T\}}$, $\tau^{\downarrow\emptyset}$, and Ψ_T^a

$\mathcal{W}_{\{T,R,D\}}$			$\tau^{\downarrow\{T,R,D\}}$	$\tau^{\downarrow\{T,R\}}$	Ψ_D	$\tau^{\downarrow\{T\}}$	$\tau^{\downarrow\emptyset}(\blacklozenge)$	$\Psi_T(\blacklozenge)$
t	ns	d	-16,600	-4,100	$\sim d$	22,500	22,500	t
t	ns	$\sim d$	-4,100					
t	os	d	8,000	8,000	d			
t	os	$\sim d$	-3,500					
t	cs	d	18,600	18,600	d			
t	cs	$\sim d$	-2,400					
$\sim t$	nr	d	20,000	20,000	d	20,000		
$\sim t$	nr	$\sim d$	0					

^a τ denotes the joint valuation $(\pi \otimes \kappa) \otimes \rho \otimes \mu$.

a payoff valuation of h . Then

$$(\alpha^{\downarrow(h-\{D\})})^{\downarrow(h-\{R,D\})}(\mathbf{c}) \geq \alpha^{\downarrow(h-\{R\})}(\mathbf{c})^{\downarrow(h-\{R,D\})}(\mathbf{c})$$

for all $\mathbf{c} \in \mathcal{W}_{h-\{R,D\}}$.

Proof. The proof follows trivially from the definition of marginalization.

It is clear from Lemma 3, that in marginalizing more than one variable, the order of elimination of

the variables may make a difference. As we will see shortly, we need to marginalize all variables out of the joint valuation. What sequence should we use? This is where the precedence constraints come into play. We define marginalization such that variable Y is marginalized before X whenever $X > Y$.

Suppose h and g are nonempty subsets of \mathcal{X} such that g is a proper subset of h , α is a valuation for h , and $>$ is a partial order on \mathcal{X} satisfying the perfect recall condition. The *marginal of α for g with respect*

to the partial order $>$, denoted by $\alpha^{\downarrow g}$, is a valuation for g defined as

$$\alpha^{\downarrow g} = (((\alpha^{\downarrow(h-\{X_1\})})^{\downarrow(h-\{X_1, X_2\})}) \dots)^{\downarrow(h-\{X_1, X_2, \dots, X_k\})}, \quad (10)$$

where $h - g = \{X_1, \dots, X_k\}$, and $X_1 X_2 \dots X_k$ is a sequence of variables in $h - g$ such that with respect to the partial order $>$, X_1 is a minimal element of $h - g$, X_2 is a minimal element of $h - g - \{X_1\}$, etc.

The marginalization sequence $X_1 X_2 \dots X_k$ may not be unique because $>$ is only a partial order. But, since $>$ satisfies the perfect recall condition, it is clear from Lemmas 1 and 2, that the definition of $\alpha^{\downarrow g}$ in (10) is well defined.

3.5. Solution for a Variable

As we will see shortly, the main objective in solving a decision problem is computing an optimal strategy. Computing an optimal strategy is a matter of book-keeping. Each time we eliminate a decision variable from a payoff valuation using maximization, we store a table of optimal values of the decision variables where the maximums are achieved. We can think of this table as a function. We call this function “a solution” for the decision variable. Formally, we define a solution as follows.

Suppose h is a subset of variables such that decision variable $D \in h$, and π is payoff valuation for h . A function $\Psi_D: \mathcal{W}_{h-\{D\}} \rightarrow \mathcal{W}_D$ is called a *solution for D* (with respect to π) if

$$\pi^{\downarrow(h-\{D\})}(\mathbf{c}) = \pi(\mathbf{c}, \Psi_D(\mathbf{c})) \quad (11)$$

for all $\mathbf{c} \in \mathcal{W}_{h-\{D\}}$.

Table VII shows the solution for D with respect to $((\pi \otimes \kappa) \otimes \rho \otimes \mu)^{\downarrow\{T, R, D\}}$, and the solution for T with respect to $((\pi \otimes \kappa) \otimes \rho \otimes \mu)^{\downarrow\{T\}}$ in the oil wildcatter's problem.

3.6. Strategy

The main task in a decision problem is to compute an optimal strategy. What constitutes a strategy? Intuitively, a strategy is a choice of an act for each decision variable D as a function of configurations of random variables R such that $R > D$. Let $\text{Pr}(D) = \{R \in \mathcal{R}_R \mid R > D\}$. We call $\text{Pr}(D)$ the *predecessors of D* . Thus, a *strategy* σ is a collection of functions $\{\xi_D\}_{D \in \mathcal{D}_D}$ where $\xi_D: \mathcal{W}_{\text{Pr}(D)} \rightarrow \mathcal{W}_D$.

Suppose that $\sigma = \{\xi_D\}_{D \in \mathcal{D}_D}$ is a strategy, and \mathbf{y} is a configuration of \mathcal{R}_R . Then σ and \mathbf{y} together determine a unique configuration of \mathcal{R}_D . Let $\mathbf{a}_{\sigma, \mathbf{y}}$ denote this unique configuration of \mathcal{R}_D . By definition

$$\mathbf{a}_{\sigma, \mathbf{y}}^{\downarrow\{D\}} = \xi_D(\mathbf{y}^{\downarrow\text{Pr}(D)}) \quad \text{for all } D \in \mathcal{D}_D.$$

In the next section, we will formally define the decision problem. Before we can do this, we will have to explain the semantics of a VBS representation.

4. WELL DEFINED VBS REPRESENTATIONS AND SEMANTICS

In this section, we describe when a VBS representation of a decision problem is well defined. We also describe the semantics of a well defined VBS representation. Suppose that

$$\Delta = \{\mathcal{R}_D, \mathcal{R}_R, \{\mathcal{W}_X\}_{X \in \mathcal{X}}, \{\pi_1, \dots, \pi_m\}, \{\rho_1, \dots, \rho_n\}, \rightarrow\}$$

is a VBS representation of a decision problem. How can we tell if Δ is well defined? And, assuming that Δ is well defined, what does Δ mean? We will answer these two related questions in terms of a canonical decision problem.

4.1. Canonical Decision Problem

A *canonical decision problem* Δ_C consists of a single decision variable D with a finite frame \mathcal{W}_D , a single random variable R with a finite frame \mathcal{W}_R , a single payoff valuation π for $\{D, R\}$, a single conditional potential ρ for R given $\{D\}$, and a precedence relation \rightarrow defined by $D \rightarrow R$. Figure 4 shows a valuation network and a decision tree representation of the canonical decision problem.

The meaning of the canonical decision problem is as follows. The elements of \mathcal{W}_D are acts, and the elements of \mathcal{W}_R are states of nature. The conditional potential ρ is a family of probability distributions for R , one for each act $\mathbf{d} \in \mathcal{W}_D$. In other words, the probability distribution of random variable R is conditioned on the act \mathbf{d} chosen by the decision maker. The probability $\rho(\mathbf{d}, \mathbf{r})$ can be interpreted as the conditional probability of $R = \mathbf{r}$ given that $D = \mathbf{d}$. Using

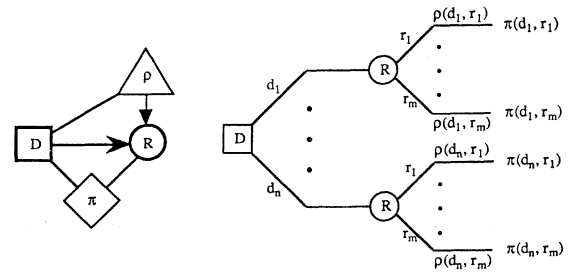


Figure 4. A VBS and a decision tree representation of the canonical decision problem.

our marginalization notation, (1) for a conditional potential can be written as

$$\rho^{\downarrow D} \text{ is the vacuous potential for } D, \text{ i.e., } \rho^{\downarrow D}(\mathbf{d}) = 1 \quad \text{for all } \mathbf{d} \in \mathcal{W}_D. \quad (12)$$

The payoff valuation π is a conditional payoff function—if the decision maker chooses act \mathbf{d} and the state of nature \mathbf{r} prevails, then the payoff to the decision maker is $\pi(\mathbf{d}, \mathbf{r})$. The precedence relation \rightarrow states that the true state of nature is revealed to the decision maker only after the decision maker has chosen an act.

Solving a canonical decision problem using the criterion of maximizing expected payoff is easy. The expected payoff associated with act \mathbf{d} is

$$\sum \{(\pi \otimes \rho)(\mathbf{d}, \mathbf{r}) \mid \mathbf{r} \in \mathcal{W}_R\} = (\pi \otimes \rho)^{\downarrow D}(\mathbf{d}).$$

The maximum expected payoff (associated with an optimal act, say \mathbf{d}^*) is

$$\begin{aligned} \text{MAX}\{(\pi \otimes \rho)^{\downarrow D}(\mathbf{d}) \mid \mathbf{d} \in \mathcal{W}_D\} \\ = ((\pi \otimes \rho)^{\downarrow D})^{\downarrow \Diamond}(\Diamond) = (\pi \otimes \rho)^{\downarrow \Diamond}(\Diamond). \end{aligned}$$

Finally, act \mathbf{d}^* is optimal if and only if

$$(\pi \otimes \rho)^{\downarrow D}(\mathbf{d}^*) = (\pi \otimes \rho)^{\downarrow \Diamond}(\Diamond).$$

Consider the decision problem

$$\Delta = \{\mathcal{D}, \mathcal{R}, \{\mathcal{W}_X\}_{X \in \mathcal{X}}, \{\pi_1, \dots, \pi_m\}, \{\rho_1, \dots, \rho_n\}, \rightarrow\}.$$

We will explain the meaning of Δ by reducing it to an equivalent canonical decision problem $\Delta_C = \{\{D\}, \{R\}, \{\mathcal{W}_D, \mathcal{W}_R\}, \{\pi\}, \{\rho\}, \rightarrow\}$. To define Δ_C , we need to define $\mathcal{W}_D, \mathcal{W}_R, \pi$, and ρ . Define \mathcal{W}_D such that for each distinct strategy σ of Δ , there is a corresponding act \mathbf{d}_σ in \mathcal{W}_D . Define \mathcal{W}_R such that for each distinct configuration \mathbf{y} of \mathcal{R} in Δ , there is a corresponding configuration \mathbf{r}_y in \mathcal{W}_R .

Before we define payoff valuation π for $\{D, R\}$, we need some notation. Consider the joint payoff valuation $\pi_1 \otimes \dots \otimes \pi_m$ in Δ . By (2), the domain of this valuation includes all of \mathcal{W}_D . Typically, the domain of this valuation will also include some (or all) random variables. Let p denote the subset of random variables included in the domain of the joint payoff valuation, i.e., $p \subseteq \mathcal{R}$ such that $\pi_1 \otimes \dots \otimes \pi_m$ is a payoff valuation for $\mathcal{W}_D \cup p$. Define payoff valuation π for $\{D, R\}$ such that

$$\pi(\mathbf{d}_\sigma, \mathbf{r}_y) = (\pi_1 \otimes \dots \otimes \pi_m)(\mathbf{a}_{\sigma, y}, \mathbf{y}^{\downarrow p}) \quad (13)$$

for all strategy σ of Δ , and for all configurations $\mathbf{y} \in \mathcal{W}_R$. Remember that $\mathbf{a}_{\sigma, y}$ is the unique configuration of \mathcal{X}_D determined by σ and \mathbf{y} .

Consider the joint potential $\rho_1 \otimes \dots \otimes \rho_n$. By (3), this potential includes all random variables in its domain. Let q denote the subset of decision variables included in the domain of the joint potential, i.e., $q \subseteq \mathcal{X}_D$ such that $\rho_1 \otimes \dots \otimes \rho_n$ is a potential for $q \cup \mathcal{X}_R$. Note that q could be empty. Define potential ρ for $\{D, R\}$ such that

$$\rho(\mathbf{d}_\sigma, \mathbf{r}_y) = (\rho_1 \otimes \dots \otimes \rho_n)(\mathbf{a}_{\sigma, y}^{\downarrow q}, \mathbf{y}) \quad (14)$$

for all strategy σ , and for all configurations $\mathbf{y} \in \mathcal{W}_R$; Δ_C , as defined above in (13) and (14), is a canonical decision problem only if ρ , defined in (14), is a conditional potential satisfying condition (1). This motivates the following definition: Δ is a well defined VBS representation of a decision problem if and only if

$$\sum \{(\rho_1 \otimes \dots \otimes \rho_n)(\mathbf{x}, \mathbf{y}) \mid \mathbf{y} \in \mathcal{W}_R\} = 1 \quad (15)$$

for every $\mathbf{x} \in \mathcal{W}_q$, or equivalently, if and only if $(\rho_1 \otimes \dots \otimes \rho_n)^{\downarrow q}$ is the vacuous potential for q . We are assuming, of course, that (1) for conditional potentials, (2) for the payoff hypergraph \mathcal{H}_D , (3) for the potential hypergraph \mathcal{H}_R , and the four conditions for the precedence relation \rightarrow , are true.

In summary, a VBS representation Δ of a decision problem is well defined if: 1) for conditional potentials, 2) for the payoff hypergraph, 3) for the potential hypergraph, the four conditions for the precedence relation, and (15) for the potentials, are all satisfied. Furthermore, in a well defined VBS representation Δ of a decision problem, the payoff valuations $\{\pi_1, \dots, \pi_m\}$ represent the factors of a joint payoff function π , and the potentials $\{\rho_1, \dots, \rho_n\}$ represent the factors of a family of probability distributions ρ .

It is easy to verify that the VBS representation of the oil wildcatter's problem described in Section 2 is well defined. Note that the domain of the joint potential $\rho \otimes \mu$ (shown in Table V) is $\{T, R, O\}$. Since T is the only decision variable in this set, because T has no predecessors, and the frame for T has two configurations, there are exactly two distinct strategy components ξ_T . Thus, $\rho \otimes \mu$ represents a family of two distinct joint distributions for $\{R, O\}$, one conditioned on the decision to perform a seismic test ($\xi_T(\Diamond) = t$) and another conditioned on the decision to not perform a seismic test ($\xi_T(\Diamond) = \sim t$).

4.2. The Decision Problem

Suppose that

$$\Delta = \{\mathcal{D}, \mathcal{R}, \{\mathcal{W}_X\}_{X \in \mathcal{X}}, \{\pi_1, \dots, \pi_m\}, \{\rho_1, \dots, \rho_n\}, \rightarrow\}$$

is a well defined decision problem. Let $\Delta_C = \{\{D\}, \{R\}, \{\mathcal{W}_D, \mathcal{W}_R\}, \{\pi\}, \{\rho\}, \rightarrow\}$ represent an equivalent

canonical decision problem. In the canonical decision problem Δ_C , the two computations that are of interest are the computation of the maximum expected value $(\pi \otimes \rho)^{\downarrow \Diamond}(\Diamond)$, and the computation of an optimal act \mathbf{d}_{σ^*} such that $(\pi \otimes \rho)^{\downarrow D}(\mathbf{d}_{\sigma^*}) = (\pi \otimes \rho)^{\downarrow \Diamond}(\Diamond)$. Since we know the mapping between Δ and Δ_C , we can now formally define the questions posed in a decision problem Δ . There are two computations of interest.

First, we would like to compute the maximum expected value of the payoffs. The maximum expected payoff is given by $(\otimes(\pi_1, \dots, \pi_m, \rho_1, \dots, \rho_n))^{\downarrow \Diamond}(\Diamond)$.

Second, we would like to compute an optimal strategy σ^* that gives the maximum expected value $(\otimes(\pi_1, \dots, \pi_m, \rho_1, \dots, \rho_n))^{\downarrow \Diamond}(\Diamond)$. A strategy σ^* of Δ is *optimal* if

$$(\pi \otimes \rho)^{\downarrow D}(\mathbf{d}_{\sigma^*}) = (\otimes(\pi_1, \dots, \pi_m, \rho_1, \dots, \rho_n))^{\downarrow \Diamond}(\Diamond),$$

where π , ρ , and D refer to the equivalent canonical decision problem Δ_C .

The computation of an optimal strategy in the oil wildcatter's problem is shown in Table VII. The maximum expected value is \$22,500. An optimal strategy can be constructed from the information in Ψ_T and Ψ_D as follows. From Ψ_T , it can be seen that the oil wildcatter should do the seismic test (t). From Ψ_D , it can be seen that if the test result is ns , then the optimal decision is not to drill ($\sim d$). However, if the test result is either os or cs , then the optimal decision is to drill (d). Note that the solution of VBS involves no divisions. Thus, the computation of the conditional distribution of O given R (during the preprocessing of probabilities) in the decision tree methodology is unnecessary.

5. A FUSION ALGORITHM FOR SOLVING VBS USING LOCAL COMPUTATION

In this section, we describe a method for solving VBS using local computation. The solution for the oil wildcatter's problem shown in Tables V, VI, and VII involves combination on the space $\mathcal{W}_{\mathcal{X}}$. While this is possible for small problems, it is computationally not feasible for problems with many variables. Given the structure of the oil wildcatter's problem, it is not possible to avoid the combination operation on the space of all four variables, T , R , D , and O . However, in some problems, it may be possible to avoid such global computations.

The basic idea of the method is to successively delete all variables from the VBS. The sequence in which variables are deleted must respect the precedence con-

straints in the sense that if $X > Y$, then Y must be deleted from X . Since $>$ is only a partial order, a problem may allow several deletion sequences. Any allowable deletion sequence may be used. All allowable deletion sequences lead to the same answers. However, different deletion sequences may involve different computational costs. We will comment on good deletion sequences at the end of this section.

When we delete a variable, we have to do a "fusion" operation on valuations that bear on the variable. Before we describe the fusion operation, we need to define a division operation for potentials.

Division. Suppose that α is a potential for g , and $h \subseteq g$. Then we define $\alpha/\alpha^{\downarrow h}$, called α divided by $\alpha^{\downarrow h}$, to be a potential for g defined as

$$(\alpha/\alpha^{\downarrow h})(\mathbf{r}) = \alpha(\mathbf{r})/\alpha^{\downarrow h}(\mathbf{r}^{\downarrow h}) \quad (16)$$

for all $\mathbf{r} \in \mathcal{W}_g$. If $\alpha(\mathbf{r}) = \alpha^{\downarrow h}(\mathbf{r}^{\downarrow h}) = 0$, then we consider $(\alpha/\alpha^{\downarrow h})(\mathbf{r}) = 0$. In all other respects, the right-hand side of (16) should be interpreted as the usual division of two real numbers. Since $\alpha^{\downarrow h}(\mathbf{r}^{\downarrow h}) \geq \alpha(\mathbf{r})$ for all $\mathbf{r} \in \mathcal{W}_g$, $\alpha/\alpha^{\downarrow h}$ is a well defined potential. If α is a potential for g such that $\alpha^{\downarrow h}$ is a vacuous potential for h , then $\alpha/\alpha^{\downarrow h} = \alpha$. For example, if α is a conditional potential for R given $g - \{R\}$, then it follows from the definition of a conditional potential that $\alpha^{\downarrow(g - \{R\})}$ is a vacuous potential. Thus, if α is a conditional potential for R given $g - \{R\}$, then $\alpha/\alpha^{\downarrow(g - \{R\})} = \alpha$.

Fusion. The fusion operation depends on the type of variable being eliminated and the nature of valuations that bear on the variable. Consider a set of valuations consisting of j payoff valuations π_1, \dots, π_j and k potentials ρ_1, \dots, ρ_k . Suppose that π_i is a payoff valuation of h_i , and ρ_i is a potential for g_i . Let $\text{Fus}_X\{\pi_1, \dots, \pi_j, \rho_1, \dots, \rho_k\}$ denote the collection of valuations after fusing the valuations in the set $\{\pi_1, \dots, \pi_j, \rho_1, \dots, \rho_k\}$ with respect to variable X .

Case 1. Suppose that D is a decision variable, and none of the potentials bear on D . Then $\text{Fus}_D\{\pi_1, \dots, \pi_j, \rho_1, \dots, \rho_k\}$ is a collection of valuations defined as

$$\begin{aligned} & \text{Fus}_D\{\pi_1, \dots, \pi_j, \rho_1, \dots, \rho_k\} \\ &= \{\pi^{\downarrow(h - \{D\})} \mid \pi \in h_i \mid D \notin h_i\} \cup \{\rho_1, \dots, \rho_k\}, \end{aligned} \quad (17)$$

where $\pi = \otimes\{\pi_i \mid D \in h_i\}$, and $h = \cup\{h_i \mid D \in h_i\}$.

In this case, after fusion, the set of valuations is changed as follows. All payoff valuations that bear on D are combined, and the resulting payoff valuation is marginalized such that D is eliminated from its domain. The payoff valuations that do not bear on D , and all potentials, remain unchanged.

Case 2. Suppose that R is a random variable, and none of the j payoff valuations bear on R . In this case, $\text{Fus}_R\{\pi_1, \dots, \pi_j, \rho_1, \dots, \rho_k\}$ is defined as

$$\text{Fus}_R\{\pi_1, \dots, \pi_j, \rho_1, \dots, \rho_k\} = \{\pi_1, \dots, \pi_j\} \cup \{\rho_i \mid R \notin g_i\} \cup \{\rho^{\downarrow(g-\{R\})}\}, \quad (18)$$

where $\rho = \otimes\{\rho_i \mid R \in g_i\}$, and $g = \cup\{g_i \mid R \in g_i\}$.

In this case, after fusion, the set of valuations is changed as follows. The payoff valuations remain unchanged (since they do not bear on R). The potentials that do not bear on R remain unchanged. The potentials that bear on R are combined and the resulting potential is marginalized such that R is eliminated from its domain.

Case 3. Suppose that R is a random variable, and all j payoff valuations bear on R . In this case, $\text{Fus}_R\{\pi_1, \dots, \pi_j, \rho_1, \dots, \rho_k\}$ is defined as

$$\text{Fus}_R\{\pi_1, \dots, \pi_j, \rho_1, \dots, \rho_k\} = \{(\pi \otimes \rho)^{\downarrow((h \cup g) - \{R\})}\} \cup \{\rho_i \mid R \notin g_i\}, \quad (19)$$

where

$$\pi = \otimes\{\pi_i \mid R \in h_i\} = \otimes\{\pi_1, \dots, \pi_j\},$$

$$\rho = \otimes\{\rho_i \mid R \in g_i\}, \quad h = \cup\{h_i \mid R \in h_i\},$$

and

$$g = \cup\{g_i \mid R \in g_i\}.$$

In this case, after fusion, the set of valuations is changed as follows. All payoff valuations and those potentials that depend on R are combined, and the resulting payoff valuation is marginalized such that R is eliminated from its domain. The potentials that do not bear on R remain unchanged.

Case 4. Suppose that R is a random variable, there is a payoff valuation that bears on R , and there is a payoff valuation that does not bear on R . In this case, $\text{Fus}_R\{\pi_1, \dots, \pi_j, \rho_1, \dots, \rho_k\}$ is defined as

$$\begin{aligned} \text{Fus}_R\{\pi_1, \dots, \pi_j, \rho_1, \dots, \rho_k\} &= \{\pi_i \mid R \notin h_i\} \cup \{[\pi \otimes (\rho / \rho^{\downarrow(g-\{R\})})]^{\downarrow((h \cup g) - \{R\})}\} \\ &\quad \cup \{\rho_i \mid R \notin g_i\} \cup \{\rho^{\downarrow(g-\{R\})}\}, \end{aligned} \quad (20)$$

where

$$\pi = \otimes\{\pi_i \mid R \in h_i\}, \quad \rho = \otimes\{\rho_i \mid R \in g_i\},$$

$$h = \cup\{h_i \mid R \in h_i\},$$

and

$$g = \cup\{g_i \mid R \in g_i\}.$$

In this case, after fusion, the set of valuations is changed as follows. The payoff valuations that do not bear on R , and the potentials that do not bear on R , remain unchanged. A new potential, $\rho^{\downarrow(g-\{R\})}$, is created. Finally, we combine all potentials that bear on R , divide the resulting potential by the new potential that was created, combine the resulting potential with the payoff valuations that bear on R , and marginalize the resulting payoff valuation such that R is eliminated from its domain.

In Markov decision processes, it is often the case that the new potential created in Case 4 is vacuous. If $\rho^{\downarrow(g-\{R\})}$ is a vacuous potential for $g - \{R\}$, then (20) can be simplified as

$$\begin{aligned} \text{Fus}_R\{\pi_1, \dots, \pi_j, \rho_1, \dots, \rho_k\} &= \{\pi_i \mid R \notin h_i\} \cup \{(\pi \otimes \rho)^{\downarrow((h \cup g) - \{R\})}\} \\ &\quad \cup \{\rho_i \mid R \notin g_i\}, \end{aligned} \quad (21)$$

where

$$\pi = \otimes\{\pi_i \mid R \in h_i\}, \quad \rho = \otimes\{\rho_i \mid R \in g_i\},$$

$$h = \cup\{h_i \mid R \in h_i\},$$

and

$$g = \cup\{g_i \mid R \in g_i\}.$$

We are now ready to state the main theorem.

Theorem 1. Suppose that

$$\Delta = \{\mathcal{X}_D, \mathcal{X}_R, \{\mathcal{W}_X\}_{X \in \mathcal{X}}, \{\pi_1, \dots, \pi_m\}, \{\rho_1, \dots, \rho_n\}, \rightarrow\}$$

is a well defined decision problem. Suppose that $X_1 X_2 \dots X_k$ is a sequence of variables in $\mathcal{X} = \mathcal{X}_D \cup \mathcal{X}_R$ such that with respect to the partial order $>$, X_1 is a minimal element of \mathcal{X} , X_2 is a minimal element of $\mathcal{X} - \{X_1\}$, etc. Then $\{(\otimes\{\pi_1, \dots, \pi_m, \rho_1, \dots, \rho_n\})^{\downarrow \otimes}\} = \text{Fus}_{X_k}\{\dots \text{Fus}_{X_2}\{\text{Fus}_{X_1}\{\pi_1, \dots, \pi_m, \rho_1, \dots, \rho_n\}\}\}$.

Let us illustrate the statement of Theorem 1 for the oil wildcatter's problem. In this problem, we have payoff valuation κ for $\{T\}$, payoff valuation π for $\{D, O\}$, potential ρ for $\{O\}$, and potential μ for $\{T, R, O\}$. Also, as per the precedence constraints, $T > R > D > O$. First, it is easy to see that after fusing with respect to O (using (20)), the set of valuations is

$$\{\kappa, (\pi \otimes ((\rho \otimes \mu) / (\rho \otimes \mu)^{\downarrow\{T,R\}}))^{\downarrow\{T,R,D\}}, (\rho \otimes \mu)^{\downarrow\{T,R\}}\}.$$

Second, after fusing with respect to D (using (17)), the set of valuations is

$$\{\kappa, (\pi \otimes ((\rho \otimes \mu) / (\rho \otimes \mu)^{\downarrow\{T,R\}}))^{\downarrow\{T,R\}}, (\rho \otimes \mu)^{\downarrow\{T,R\}}\}.$$

Third, after fusing with respect to R (using (21)) since $(\rho \otimes \mu)^{\downarrow\{T\}}$ is a vacuous valuation), the set of valuations is

$$\{\kappa, ((\pi \otimes ((\rho \otimes \mu)/(\rho \otimes \mu)^{\downarrow\{T,R\}}))^{\downarrow\{T,R\}} \otimes (\rho \otimes \mu)^{\downarrow\{T,R\}})^{\downarrow\{T\}}\}.$$

Note that the second payoff valuation

$$((\pi \otimes ((\rho \otimes \mu)/(\rho \otimes \mu)^{\downarrow\{T,R\}}))^{\downarrow\{T,R\}} \otimes (\rho \otimes \mu)^{\downarrow\{T,R\}})^{\downarrow\{T\}}$$

simplifies to

$$((\pi \otimes (\rho \otimes \mu))^{\downarrow\{T,R\}})^{\downarrow\{T\}} = (\pi \otimes \rho \otimes \mu)^{\downarrow\{T\}}.$$

Finally, after fusing with respect to T (using (17)), the set of valuation reduces to the singleton set

$$\{(\kappa \otimes (\pi \otimes \rho \otimes \mu)^{\downarrow\{T\}})^{\downarrow\emptyset}\}.$$

As per Theorem 1,

$$((\kappa \otimes \pi) \otimes \rho \otimes \mu)^{\downarrow\emptyset} = (\kappa \otimes (\pi \otimes \rho \otimes \mu)^{\downarrow\{T\}})^{\downarrow\emptyset}.$$

In the oil wildcatter's example, notice that the division in the first fusion operation is neutralized by the multiplication in the third fusion operation. A natural question that arises is: Can the division in (20) always be avoided? One way to avoid the division in (20) is to avoid circumstances where (20) applies. One way to avoid the circumstance where (20) applies is to make sure we have only one payoff valuation. If there is only one payoff valuation in the valuation-based system, then whenever we fuse valuations with respect to a random variable, either Case 2 or 3 will apply, never Case 4. If we have more than one payoff valuation, then we can always combine these to get one valuation before we start the fusion algorithm. Of course, this means that we will be unable to take advantage of a factorization of the joint utility function. In the oil wildcatter's problem, if we combine κ and π before we start the fusion algorithm, then the computations in the fusion algorithm are exactly the same as the computations in solving the problem globally (as was done in Section 4).

Another way to avoid divisions is to have a factorization of the joint probability distributions such that the added potential is always vacuous. (In this case, we can always use (21) instead of (20) and avoid divisions.) The added potential is always vacuous if we have only a conditional potential for each random variable such that the variables on which the probabilities are conditioned always precede the random variable. This typically happens in Markov decision problems. An example of such a problem follows.

5.1. A Finite Markov Decision Problem

Consider a valuation-based system as shown in Figure 5. There are three random variables, R_1 , R_2 ,




Figure 5. A finite Markov decision problem.

and R_3 , and two decision variables, D_1 , and D_2 . The information constraints are as follows. The decision maker first chooses an act from \mathcal{W}_{D_1} , then from \mathcal{W}_{D_2} . When the decision maker has to choose an act from frame \mathcal{W}_{D_i} , she knows only the true value of all random variables in the set $\{R_j | j \leq i\}$. The probability distribution of R_j only depends on R_{j-1} and D_{j-1} . Thus, we are given a conditional potential ρ_1 for R_1 given \emptyset representing the prior probability of R_1 , a conditional potential ρ_2 for R_2 given $\{R_1, D_1\}$, and a conditional potential ρ_3 for R_3 given $\{R_2, D_2\}$. Finally, the joint payoff function factorizes additively into three factors π_1 for $\{R_1, D_1\}$, π_2 for $\{R_2, D_2\}$, and π_3 for $\{R_3\}$.

Figure 6 shows the solution of the Markov decision problem using the fusion algorithm. The first valuation network in the figure shows the result after deletion of random variable R_3 and the resulting fusion. Note that since ρ_3 is a conditional potential for R_3 given R_2 and D_2 , by definition of a conditional potential, $\rho^{\downarrow\{R_2, D_2\}}$ is a vacuous potential. Therefore, (21) applies. The second valuation network shows the result after deletion of decision node D_2 and the resulting fusion using (17). Step 3 is similar to Step 1, and Step 4 is similar to Step 2. Finally, Step 5 is similar to Step 1. Note that the fusion method for solving this problem involves no divisions.

In general, if we wish to take advantage of an additive factorization of the joint payoff valuation, and if we have arbitrary potentials, then divisions may be inescapable. The following example demonstrates this.

Consider a valuation-based system as shown in Figure 7 with two random variables, R_1 and R_2 , and one decision variable D . We are given one potential ρ for $\{R_1, R_2\}$ representing the joint probability distribution for $\{R_1, R_2\}$. The joint payoff valuation factors additively into two payoff valuations, π_1 for $\{R_1\}$, and π_2 for $\{D, R_2\}$. The precedence constraints are: $R_1 \rightarrow D \rightarrow R_2$. Figure 7 shows the details of the valuations and the valuation network representation. A global solution of this problem involves computing

$$(((\pi_1 \otimes \pi_2) \otimes \rho)^{\downarrow\{R_1, D\}})^{\downarrow\{R_1\}})^{\downarrow\emptyset}$$




Figure 6. The fusion algorithm for the VBS in Figure 5.

which, of course, involves no divisions. If we, however, apply the fusion algorithm to the set of valuations $\{\pi_1, \pi_2, \rho\}$, then after deletion of all three variables, the valuation in the resulting singleton set is

$$([\pi_1 \otimes [\pi_2 \otimes (\rho/\rho^{\downarrow\{R_1\}})]^{\downarrow\{R_1\}}] \otimes \rho^{\downarrow\{R_1\}})^{\downarrow\emptyset}.$$

Note that since combination is not associative, the division in this computation is unavoidable, i.e.,

$$([\pi_1 \otimes [\pi_2 \otimes (\rho/\rho^{\downarrow\{R_1\}})]^{\downarrow\{R_1\}}] \otimes \rho^{\downarrow\{R_1\}})^{\downarrow\emptyset} \neq (\pi_1 \otimes [\pi_2 \otimes \rho])^{\downarrow\{R_1\}}.$$

The fusion operations in (17), (18), (19), and (21), say with respect to X , can simply be redescribed as

$$\text{Fus}_X\{\alpha_1, \dots, \alpha_k\} = \{\alpha^{\downarrow(h-\{X\})}\} \cup \{\alpha_i \mid X \notin h_i\}, \quad (22)$$

where $\alpha = \otimes \{\alpha_i \mid X \in h_i\}$ and $h = \cup \{h_i \mid X \in h_i\}$. In words, after fusion the set of valuations is changed as follows. All valuations that bear on X are combined, and the resulting valuation is marginalized such that

X is eliminated from its domain. The valuations that do not bear on X remain unchanged.

The fusion operation in (20) is the only one that involves division and is designed especially to take care of the nonassociativity of the combination operation (see the proof of Theorem 1 in Section 8). One implication is that if the combination is associative, then the fusion operation can be simplified as in (22) and no divisions are necessary (see Shenoy 1990b for a proof of this assertion). The combination operation is associative if, for example, the utility function in a decision problem factorizes multiplicatively. In this case, for computational purposes, it is not necessary to distinguish between payoff valuations and potentials.

5.2. Deletion Sequences

Since $>$ is only a partial order, in general, we may have many sequences of variables that satisfy the condition stated in Theorem 1. (We call such sequences *deletion sequences*.) If so, which deletion sequence should one use? First, note that all deletion sequences lead to the same final result. This is implied in the statement of the theorem. Second, different deletion sequences may involve different computational efforts. For example, consider the VBS shown in Figure 8. In this example, deletion sequence $R_2 R_1 D$ involves less computational effort than $R_1 R_2 D$ as the former involves combinations on the frame of only two variables, whereas the latter involves combinations on the frame of all three variables. Finding an optimal deletion sequence is a secondary optimization problem that has been shown to be NP-complete (Arnborg, Corneil and Proskurowski 1987). However, there are several heuristics for finding good deletion sequences (Kong 1986, Mellouli 1987, Zhang 1988).




Figure 7. A decision problem requiring divisions in the fusion algorithm.




Figure 8. A VBS with two deletion sequences, R_2R_1D and R_1R_2D .

One such heuristic is called one-step-look-ahead (Kong). This heuristic tells us which variable to delete next from among those that qualify. As per this heuristic, the variable that should be deleted next is one that leads to combination over the smallest frame. For example, in the VBS of Figure 8, two variables qualify for the first deletion, R_1 and R_2 . This heuristic would pick R_2 over R_1 because deletion of R_1 involves combination over the frame of $\{D, R_1, R_2\}$, whereas deletion of R_2 only involves combination over the frame of $\{R_1, R_2\}$. Thus, this heuristic would choose deletion sequence R_2R_1D .

Consider the VBS in Figure 9. This VBS is essentially the same as the VBS of Figure 8 except that the precedence constraint $R_2 \rightarrow R_1$ has been added (say R_1 is only observed after R_2 is observed). In this VBS, the partial order $>$ only allows deletion sequence R_1R_2D . It is clear from Lemma 2, however, that we would get the same answer if we had dropped the constraint $R_2 > R_1$ from the partial order $>$. One advantage of dropping this constraint is that we can solve the problem more efficiently. Let $>'$ denote the order obtained from $>$ by dropping pairs of the type $R_i > R_j$ and by dropping pairs of the type $D_i > D_j$, where R_i, R_j are random variables and D_i, D_j are decision variables. From Lemmas 1 and 2, it is clear that we can substitute $>'$ for $>$ in the statement of Theorem 1 and the theorem remains valid.




Figure 9. The solution of this VBS can be found more efficiently by dropping the pair $R_2 > R_1$ from the partial order $>$.

6. COMPARISON WITH INFLUENCE DIAGRAMS

In this section, we briefly compare our VBS representation and solution of decision problems with the influence diagram representation and solution. We assume that the reader is familiar with the influence diagram representation and solution methodology (Olmsted 1983, Howard and Matheson 1984, Shachter 1986, Tatman 1986, Ezawa 1986, Howard 1990, Tatman and Shachter 1990).

An influence diagram representation of a decision problem consists of decision nodes, random variable nodes, subvalue nodes, supvalue nodes, directed edges, a conditional potential for each random variable node, and a payoff valuation for each subvalue node. Figure 10 shows an influence diagram representation of the oil wildcatter's problem.

In influence diagrams, decision variables are shown as rectangular nodes, random variables are shown as circular nodes, and payoff valuations are shown as diamond-shaped nodes.

In influence diagrams, potentials are not depicted explicitly. It is implicit that each random variable node has a conditional potential given its direct predecessors. (Node X is a direct predecessor of node Y if there is a directed edge (X, Y) in the diagram.) For example, the influence diagram for the oil wildcatter's problem shown in Figure 10 assumes a conditional potential for O given \emptyset , and a conditional potential for R given $\{T, O\}$. One advantage of this representation is that all conditional independences among random variables can easily be read from an influence diagram (Pearl, Geiger and Verma 1990). VBSs, on the other hand, emphasize factorizations of the joint probability distributions. Although factorizations have not been as well studied as conditional independences, the two are equivalent. For example, it is well known from probability theory that random variables X and Y are independent with respect to joint probability P if and only if $P = Q \otimes R$, where Q is a potential for $\{X\}$, and R is a potential for $\{Y\}$ (e.g., see Pearl, p. 83).

Payoff valuations are shown explicitly in influence diagrams as subvalue nodes. The directed edges that point to subvalue nodes indicate the domain of the payoff valuations. The direction of these edges have no special significance. In the influence diagram for the oil wildcatter's problem, there are two subvalue nodes corresponding to payoff valuations π for $\{D, O\}$ and κ for $\{T\}$. In the influence diagram representation of the oil wildcatter's problem, there is a third value node τ connected to π and κ ; τ represents the combination $\pi \otimes \kappa$ and is called a *supvalue node*. In

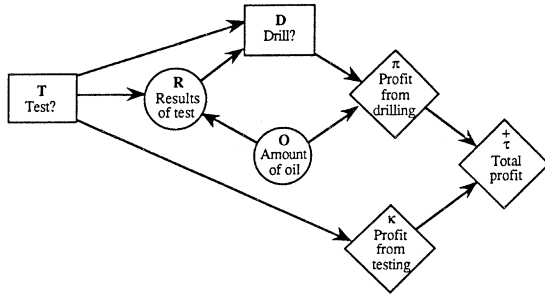


Figure 10. An influence diagram representation of the oil wildcatter's problem.

valuation networks, combination is implicit and is not shown explicitly.

In influence diagrams, information constraints are encoded in the form of directed edges pointing to decision nodes. The semantics of these directed edges are as follows. For any decision node D and any random variable node R , the true value of R is known to the decision maker at the point in time when they have to choose an act from the frame D if and only if there is a directed edge in the diagram from R to D . In the influence diagram representation of the oil wildcatter's problem, there are no directed edges pointing to T . This implies that nothing is known when the decision whether to perform a seismic test or not is made. Also, there are two directed edges (T, D) and (R, D) . This means that at the point in time when the decision whether to drill or not has to be made, the oil wildcatter's knows the decision regarding performance of the seismic test and also knows the test results, but not the amount of oil. The oil wildcatter does not know the amount of oil (at the time the decision to drill is made) because there is no directed edge (O, D) in the diagram.

Given an influence diagram representation, the translation to a VBS representation is clear. This means that the fusion algorithm described in Section 5 can be applied also to influence diagrams. On the other hand, if we have a well defined VBS representation such that all potentials are conditional potentials, then the translation to an influence diagram representation is also clear. However, if we have a VBS representation with potentials that are not conditional potentials, then there is no *direct* equivalent influence diagram representation. (By direct, we mean without any preprocessing.) For example, Figure 11 shows a valuation-based system with two random variables and one decision variable. The potential ρ for $\{R_1, R_2\}$ is the joint probability distribution of

R_1 and R_2 . The joint utility function factors multiplicatively into two factors, π_1 for $\{D, R_1\}$ and π_2 for $\{R_1, R_2\}$. There is no direct influence diagram representation of this problem. An influence diagram representation of this problem involves, for example, the computation of the marginal of R_1 , $\rho^{\downarrow R_1}$, and the computation of the conditional of R_2 given $\{R_1\}$, $\rho/\rho^{\downarrow R_1}$. Note that these computations are unnecessary for an efficient solution of the problem given by $([\pi_1 \otimes (\pi_2 \otimes \rho)^{\downarrow R_1}]^{\downarrow D})^{\downarrow \emptyset}$.

6.1. Solving Influence Diagrams

Here we only compare the arc reversal method for solving influence diagrams with the fusion algorithm for solving VBSs. We refer the reader to Olmsted, (1983), Shachter (1986), Ezawa (1986), Tatman (1986), and Tatman and Shachter (1990) for details about the arc reversal method.

Assume that each random variable node has a conditional potential stored at its location, and assume that each subvalue node has a payoff valuation stored at its location. These valuations get modified in the solution process. The process of solving influence diagrams involves arc reversals and node removals. Nodes are successively removed until only one value node remains.

Removal of a random variable node involves averaging the payoff valuations that depend on the random variable node using the potential stored at the random variable node. Compared to VBSs, this operation corresponds to fusing the valuations that bear on the random variable node.

Removal of a decision node involves first combining the payoff valuations that depend on the decision node and then maximizing the resulting payoff valuation over the frame of the decision node. The resulting valuation is stored at a value node. Compared to VBSs, this operation corresponds to fusing the payoff valuations that bear on the decision node.




Figure 11. A VBS that has no direct influence diagram representation (left); an equivalent influence diagram representation after preprocessing (right).

Removal of value nodes corresponds to combining the valuations stored at the value nodes and storing the result at another value node.

As in VBSs, the sequence of removals of random variable nodes and decision nodes must respect the precedence constraints expressed by the directed edges pointing to decision nodes. Also, a random variable node can be removed only if it is not a predecessor for any other random variable nodes in the diagram. If a random variable node is a predecessor for other random variable nodes, then these arcs have to be reversed before the random variable node can be removed.

In VBSs, the fusion algorithm combines all valuations that contain the variable being marginalized. In influence diagrams, this condition is achieved by arc reversals. Arc reversals ensure that all valuations containing the node being removed are combined prior to removing the node. However, the arc reversal operation also involves unnecessary divisions and multiplications to ensure that the resulting potentials at all random variable nodes are conditional potentials. It is this aspect of the arc reversal operation that makes the influence diagram solution less computationally efficient than the solution method of VBSs.

If we disregard the unnecessary divisions and multiplications, the solution method of influence diagrams is as efficient as that of VBSs. The computational efficiency of the influence diagram solution method depends on the node removal sequence. The computational efficiency of the VBS solution method also depends on the deletion sequence. If the same sequence is used in both cases, we get approximately the same computational efficiency. In Section 5, we described the one-step-look-ahead heuristic due to Kong for picking the sequence of variables. In the case of influence diagrams, Olmsted, and Ezawa describe some heuristics for selecting the sequence of node removals.

7. CONCLUSIONS

The main objective of this paper is to propose a new method to represent and solve decision problems. The VBS representation and solution described here is a hybrid of valuation-based systems for probability propagation (Shenoy 1991c) and valuation-based systems for optimization (Shenoy 1991b).

There are several advantages of the VBS representation and solution of decision problems. First, like influence diagrams, a valuation network representation is compact when compared to decision trees. A valuation network graphically depicts the qualitative

structure of the decision problem and de-emphasizes the quantitative details of the problem. However, both VBSs and influence diagrams are appropriate only for symmetric decision problems. For asymmetric decision problems, decision tree representation is more flexible.

Second, like influence diagrams, the VBS representation separates the formulation of the problem from its solution.

Third, in symmetric decision problems, the solution procedure of VBSs is more efficient than that of decision trees. This assumes that the computational procedure of decision trees includes the preprocessing of probabilities. The solution procedure of decision trees includes unnecessary divisions. The unnecessary divisions take place during preprocessing of probabilities.

Fourth, the VBS representation is more powerful than influence diagram representation. Whereas influence diagram representation is only capable of directly representing conditional potentials, VBS representation is capable of directly representing arbitrary potentials.

Fifth, the solution method of VBSs involves minimal divisions. In comparison, the influence diagram solution method involves unnecessary divisions (in every arc reversal operation). These unnecessary divisions are the same as those in the decision tree solution process. In influence diagrams, these unnecessary operations are performed for semantical considerations. The influence diagram solution process has the property that the diagram resulting from the deletion of a random variable node is again an influence diagram. This means that the resulting potentials in the reduced influence diagram are conditional potentials. It is this demand for conditional potentials at each stage that results in the unnecessary divisions and multiplications.

Sixth, the semantics of VBSs are different from the semantics of influence diagrams. Whereas influence diagrams are based on the semantics of conditional independence, VBSs are based on the semantics of factorization.

Seventh, if a decision problem has no random variables, it reduces to an optimization problem. And, the solution technique of VBSs reduces to dynamic programming (Shenoy 1991b).

Eighth, in cases where a decision problem has no decision variables, we may be interested in finding marginals of the joint distribution for each random variable. In such problems, the solution technique described in this paper reduces to the technique for finding marginals (Shenoy 1991c). This technique also

can revise marginals in light of new observations. We represent each new observation by a potential and then use the fusion algorithm to compute the desired marginals.

7.1. Limitations of VBSs

Like influence diagrams, VBSs are appropriate only if the decision problem is symmetric or almost symmetric. In asymmetric decision problems, decision tree representation is more flexible than VBSs and influence diagrams, and decision tree solutions may be more efficient than VBSs and influence diagrams. This is because, like influence diagrams, VBSs make an asymmetric decision problem symmetric by adding dummy acts and events, and in the process enlarge the space of the problem.

In problems with many variables, like the solution procedure of influence diagrams, the fusion algorithm is tractable only if the space on which combinations are performed stay small. The space on which combinations are performed depends on the sizes of the valuations and also on the precedence constraints. We need strong independence conditions to keep the sizes of the potentials small. We also need strong assumptions on the utility function to decompose it into small payoff valuations. In the worst case, of course, solving a decision problem is NP-hard (Cooper 1987).

8. PROOFS

In this section we give a proof for Theorem 1. First we state and prove a lemma needed to prove Theorem 1.

Lemma 4. *Suppose that*

$$\Delta = \{\mathcal{X}_D, \mathcal{X}_R, \{\mathcal{X}_X\}_{X \in \mathcal{X}}, \{\pi_1, \dots, \pi_m\}, \{\rho_1, \dots, \rho_n\}, \rightarrow\}$$

is a well defined decision problem. Suppose that X is a minimal variable in $\mathcal{X} = \mathcal{X}_D \cup \mathcal{X}_R$ with respect to the partial order $>$, where $>$ is the transitive closure of \rightarrow . Then

$$(\otimes \{\pi_1, \dots, \pi_m, \rho_1, \dots, \rho_n\})^{\downarrow(\mathcal{X}-\{X\})} = \otimes \text{Fus}_X\{\pi_1, \dots, \pi_m, \rho_1, \dots, \rho_n\}.$$

Proof. We prove this result in four mutually exclusive and exhaustive cases, each corresponding to the definition of the fusion operation in (17), (18), (19) and (20), respectively. Suppose that π_i is a payoff valuation for h_i and ρ_i is a potential for g_i .

Case 1. Suppose that X is a decision variable. In this case, there cannot be a potential in the VBS Δ that bears on X because if there were one, then, since Δ is well defined, by the consistency conditions for the precedence relation \rightarrow , X cannot be a minimal variable. This would contradict the hypothesis that X is minimal. Without loss of generality, assume that π_1, \dots, π_k are the only payoff valuations that bear on X . Let $\pi = \pi_1 \otimes \dots \otimes \pi_k$, let $h = h_1 \cup \dots \cup h_k$, and let $\mathbf{c} \in \mathcal{W}_{\mathcal{X}-\{X\}}$. Then

$$\begin{aligned} & (\otimes \{\pi_1, \dots, \pi_m, \rho_1, \dots, \rho_n\})^{\downarrow(\mathcal{X}-\{X\})}(\mathbf{c}) \\ &= \text{Max}\{[\pi_1(\mathbf{c}^{\downarrow h_1}, \mathbf{x}) + \dots + \pi_k(\mathbf{c}^{\downarrow h_k}, \mathbf{x}) \\ & \quad + \pi_{k+1}(\mathbf{c}^{\downarrow h_{k+1}}) + \dots + \pi_m(\mathbf{c}^{\downarrow h_m})] \\ & \quad \cdot [\rho_1(\mathbf{c}^{\downarrow g_1}) \dots \rho_n(\mathbf{c}^{\downarrow g_n})] \mid \mathbf{x} \in \mathcal{W}_X\} \\ &= \text{Max}\{\pi_1(\mathbf{c}^{\downarrow h_1}, \mathbf{x}) + \dots + \pi_k(\mathbf{c}^{\downarrow h_k}, \mathbf{x}) \\ & \quad \cdot \pi_{k+1}(\mathbf{c}^{\downarrow h_{k+1}}) + \dots + \pi_m(\mathbf{c}^{\downarrow h_m}) \mid \mathbf{x} \in \mathcal{W}_X\} \\ & \quad \cdot [\rho_1(\mathbf{c}^{\downarrow g_1}) \dots \rho_n(\mathbf{c}^{\downarrow g_n})] \\ &= [\text{Max}\{\pi_1(\mathbf{c}^{\downarrow h_1}, \mathbf{x}) + \dots + \pi_k(\mathbf{c}^{\downarrow h_k}, \mathbf{x}) \mid \mathbf{x} \in \mathcal{W}_X\} \\ & \quad + \pi_{k+1}(\mathbf{c}^{\downarrow h_{k+1}}) + \dots + \pi_m(\mathbf{c}^{\downarrow h_m})] \\ & \quad \cdot [\rho_1(\mathbf{c}^{\downarrow g_1}) \dots \rho_n(\mathbf{c}^{\downarrow g_n})] \\ &= [\pi^{\downarrow h-\{X\}}(\mathbf{c}^{\downarrow h}) + \pi_{k+1}(\mathbf{c}^{\downarrow h_{k+1}}) + \dots + \pi_m(\mathbf{c}^{\downarrow h_m})] \\ & \quad \cdot [\rho_1(\mathbf{c}^{\downarrow g_1}) \dots \rho_n(\mathbf{c}^{\downarrow g_n})] \\ &= \otimes \text{Fus}_X\{\pi_1, \dots, \pi_m, \rho_1, \dots, \rho_n\}(\mathbf{c}). \end{aligned}$$

Case 2. Suppose that X is a random variable, and none of the m payoff valuations bear on X . Without loss of generality, assume that ρ_1, \dots, ρ_k are the only potentials that bear on X . Let $\rho = \rho_1 \otimes \dots \otimes \rho_k$, let $g = g_1 \cup \dots \cup g_k$, and let $\mathbf{c} \in \mathcal{W}_{\mathcal{X}-\{X\}}$. Then

$$\begin{aligned} & (\otimes \{\pi_1, \dots, \pi_m, \rho_1, \dots, \rho_n\})^{\downarrow(\mathcal{X}-\{X\})}(\mathbf{c}) \\ &= \sum \{[\pi_1(\mathbf{c}^{\downarrow h_1}) + \dots + \pi_m(\mathbf{c}^{\downarrow h_m})] \\ & \quad \cdot [\rho_1(\mathbf{c}^{\downarrow g_1}, \mathbf{x}) \dots \rho_k(\mathbf{c}^{\downarrow g_k}, \mathbf{x}) \rho_{k+1}(\mathbf{c}^{\downarrow g_{k+1}}) \\ & \quad \dots \rho_n(\mathbf{c}^{\downarrow g_n})] \mid \mathbf{x} \in \mathcal{W}_X\} \\ &= [\pi_1(\mathbf{c}^{\downarrow h_1}) + \dots + \pi_m(\mathbf{c}^{\downarrow h_m})][\rho_{k+1}(\mathbf{c}^{\downarrow g_{k+1}}) \dots \rho_n(\mathbf{c}^{\downarrow g_n})] \\ & \quad \cdot [\sum \{(\rho_1(\mathbf{c}^{\downarrow g_1}, \mathbf{x}) \dots \rho_k(\mathbf{c}^{\downarrow g_k}, \mathbf{x}) \mid \mathbf{x} \in \mathcal{W}_X\}] \\ &= [\pi_1(\mathbf{c}^{\downarrow h_1}) + \dots + \pi_m(\mathbf{c}^{\downarrow h_m})][\rho_{k+1}(\mathbf{c}^{\downarrow g_{k+1}}) \\ & \quad \dots \rho_n(\mathbf{c}^{\downarrow g_n}) \rho^{\downarrow g-\{X\}}(\mathbf{c}^{\downarrow g})] \\ &= \otimes \text{Fus}_X\{\pi_1, \dots, \pi_m, \rho_1, \dots, \rho_n\}(\mathbf{c}). \end{aligned}$$

Case 3. Suppose that X is a random variable, and all of the m payoff valuations bear on X . Without loss of

This work was supported in part by the National Science Foundation under grant IRI-8902444. I am grateful for comments from and conversations with Alice Agogino, Dan Geiger, Steffen Lauritzen, Pierre Ndilikilikiesha, Anthony Neugebauer, Geoff Schemmel, Leen-Kiat Soh, Glenn Shafer, Philippe

Smets, Po-Lung Yu, and Lianwen Zhang. This paper has also profited from the comments of two anonymous referees, and the comments of the area editor, Peter Farquhar.

REFERENCES

- ARNBORG, S., D. G. CORNEIL AND A. PROSKUROWSKI. 1987. Complexity of Finding Embeddings in a k -tree. *SIAM J. Alg. Discrete Meth.* **8**, 277–284.
- BELLMAN, R. E. 1957. *Dynamic Programming*. Princeton University Press, Princeton, N.J.
- BERTELE, U., AND F. BRIOSCHI. 1972. *Nonserial Dynamic Programming*. Academic Press, New York.
- COOPER, G. 1987. Probabilistic Inference Using Belief Networks is NP-Hard. Technical Report KSL-87-27, Medical Computer Science, Stanford University, Stanford, Calif.
- DARROCH, J. N., S. L. LAURITZEN AND T. P. SPEED. 1980. Markov Fields and Log-Linear Interaction Models for Contingency Tables. *Ann. Statist.* **8**, 522–539.
- DUBOIS, D., AND H. PRADE. 1990. Inference in Possibilistic Hypergraphs. In *Proceedings of the Third International Conference on Information Processing and Management of Uncertainty in Knowledge-Based Systems (IPMU-90)*, Paris, France, 228–230.
- EDWARDS, D., AND S. KREINER. 1983. The Analysis of Contingency Tables by Graphical Models. *Biometrika* **70**, 553–565.
- EZAWA, K. J. 1986. Efficient Evaluation of Influence Diagrams. Ph.D. Dissertation, Department of Engineering-Economic Systems, Stanford University, Stanford, Calif.
- HOWARD, R. A. 1990. From Influence to Relevance to Knowledge. In *Influence Diagrams, Belief Nets and Decision Analysis*. R. M. Oliver and J. Q. Smith (eds.). John Wiley, Chichester, U.K., 3–23.
- HOWARD, R. A., AND J. E. MATHESON. 1984. Influence Diagrams. In *The Principles and Applications of Decision Analysis*, R. A. Howard and J. E. Matheson (eds.). Strategic Decisions Group, Menlo Park, Calif., 719–762.
- JENSEN, F. V., S. L. LAURITZEN AND K. G. OLESEN. 1990. Bayesian Updating in Causal Probabilistic Networks by Local Computation. *Comput. Statist. Quart.* **4**, 269–282.
- JENSEN, F. V., K. G. OLESEN AND S. K. ANDERSEN. 1990. An Algebra of Bayesian Belief Universes for Knowledge-Based Systems. *Networks* **20**, 637–659.
- KIIVERI, H., T. P. SPEED AND J. B. CARLIN. 1984. Recursive Causal Models. *J. Australian Math. Soc. Ser. A* **36**, 30–52.
- KONG, A. 1986. Multivariate Belief Functions and Graphical Models. Ph.D. Dissertation, Department of Statistics, Harvard University, Cambridge, Mass.
- KUHN, H. W. 1953. Extensive Games and the Problem of Information. In *Contributions to the Theory of Games*, Vol. 2, H. W. Kuhn and A. W. Tucker (eds.). Princeton University Press, Princeton, N.J., 193–216.
- LAURITZEN, S. L., AND D. J. SPIEGELHALTER. 1988. Local Computations With Probabilities on Graphical Structures and Their Application to Expert Systems (With Discussion). *J. Roy. Statist. Soc. Ser. B* **50**, 157–224.
- MELLOULI, K. 1987. On the Propagation of Beliefs in Networks Using the Dempster-Shafer Theory of Evidence. Ph.D. Dissertation, School of Business, University of Kansas, Lawrence, Kansas.
- MILLER, A. C., M. W. MERKHOFFER, R. A. HOWARD, J. E. MATHESON AND T. R. RICE. 1976. Development of Decision Aids for Decision Analysis. Technical Report, Stanford Research Institute, Menlo Park, Calif.
- OLIVER, R. M., AND J. Q. SMITH (eds.). 1990. *Influence Diagrams, Belief Networks, and Decision Analysis*. John Wiley, New York.
- OLMSTED, S. M. 1983. On Representing and Solving Decision Problems. Ph.D. Dissertation. Department of Engineering-Economic Systems, Stanford University, Stanford, Calif.
- PEARL, J. 1988. *Probabilistic Reasoning in Intelligent Systems*. Morgan Kaufmann, San Mateo, Calif.
- PEARL, J., D. GEIGER AND T. VERMA. 1990. The Logic of Influence Diagrams. In *Influence Diagrams, Belief Nets and Decision Analysis*. R. M. Oliver, and J. Q. Smith (eds.). John Wiley, Chichester, U.K., 67–88.
- RAIFFA, H. 1968. *Decision Analysis*. Addison-Wesley, Reading, Mass.
- RAIFFA, H., AND R. SCHLAIFER. 1961. *Applied Statistical Decision Theory*. MIT Press, Cambridge, Mass.
- SAVAGE, L. J. 1950. *The Foundations of Statistics*. John Wiley, New York.
- SHACHTER, R. D. 1986. Evaluating Influence Diagrams. *Opns. Res.* **34**, 871–882.
- SHACHTER, R. D. 1988. DAVID: Influence Diagram Processing System for the Macintosh. In *Uncertainty in Artificial Intelligence*, Vol. 2, J. F. Lemmer, and L. N. Kanal (eds.). North-Holland, Amsterdam, 191–196.
- SHAFFER, G., AND P. P. SHENOY. 1988. *Local Computation in Hypertrees*. Working Paper No. 201, School of Business, University of Kansas, Lawrence, Kansas.
- SHAFFER, G., AND P. P. SHENOY. 1990. Probability Propagation. *Ann. Math. Artif. Intell.* **2**, 327–352.
- SHENOY, P. P. 1989. A Valuation-Based Language for Expert Systems. *Int. J. Approx. Reasoning* **3**, 383–411.
- SHENOY, P. P. 1990a. Valuation-Based Systems for Propositional Logic. In *Methodologies for Intelligent Systems*, Z. W. Ras, M. Zemankova and M. L. Emrich (eds.). North-Holland, Amsterdam, Vol. 5, 305–312.

- SHENOY, P. P. 1990b. A New Method for Representing and Solving Bayesian Decision Problems. Working Paper No. 223, School of Business, University of Kansas, Lawrence, Kansas.
- SHENOY, P. P. 1991a. On Spohn's Rule for Revision of Beliefs. *Int. J. Approx. Reasoning* **5**, 149–181.
- SHENOY, P. P. 1991b. Valuation-Based Systems for Discrete Optimization. In *Uncertainty in Artificial Intelligence*, P. P. Bonissone, M. Henrion, L. N. Kanal and J. F. Lemmer (eds.). Vol. 6, North-Holland, Amsterdam, 385–400.
- SHENOY, P. P. 1991c. Valuation-Based Systems: A Framework for Managing Uncertainty in Expert Systems. Working Paper No. 226, School of Business, University of Kansas, Lawrence, Kansas.
- SHENOY, P. P., AND G. SHAFER. 1988. Constraint Propagation. Working Paper No. 208, School of Business, University of Kansas, Lawrence, Kansas.
- SHENOY, P. P., AND G. SHAFER. 1990. Axioms for Probability and Belief-Function Propagation. In *Uncertainty in Artificial Intelligence*, Vol. 4, R. D. Shachter, T. S. Levitt, J. F. Lemmer and L. N. Kanal (eds.). North-Holland, Amsterdam, 169–198. Reprinted in G. Shafer and J. Pearl (eds.) 1990. *Readings in Uncertain Reasoning*, Morgan Kaufmann, San Mateo, Calif., 575–610.
- TATMAN, J. A. 1986. Decision Processes in Influence Diagrams: Formulation and Analysis. Ph.D. Dissertation, Department of Engineering-Economic Systems, Stanford University, Stanford, Calif.
- TATMAN, J. A., AND R. D. SHACHTER. 1990. Dynamic Programming and Influence Diagrams. *IEEE Trans. Syst. Man Cybernet.* **20**, 365–379.
- VON NEUMANN, J., AND O. MORGENSTERN. 1953. *Theory of Games and Economic Behavior*, 3rd ed. John Wiley, New York.
- WERMUTH, N., AND S. L. LAURITZEN. 1983. Graphical and Recursive Models for Contingency Tables. *Biometrika* **70**, 537–552.
- ZERMELO, E. 1913. Über eine Anwendung der Mengenlehre auf die Theorie des Schachspiels. In *Proceedings of the Fifth International Congress of Mathematics*, **2**, Cambridge, U.K., 501–504.
- ZHANG, L. 1988. Studies on Finding Hypertree Covers of Hypergraphs. Working Paper No. 198, School of Business, University of Kansas, Lawrence, Kansas.