DECISION ANALYSIS
Vol. 2, No. 3, September 2005, pp. 148–164
ISSN 1545-8490 | EISSN 1545-8504 | 05 | 0203 | 0148
DOI 10.1287/deca.1050.0048

Graphical Models for Groups: Belief Aggregation and Risk Sharing

David M. Pennock
Yahoo! Research, 74 N. Pasadena Avenue, 3rd Floor, Pasadena, California 91103, pennockd@yahoo-inc.com
Michael P. Wellman
University of Michigan AI Laboratory, 1101 Beal Avenue, Ann Arbor, Michigan 48109-2110, wellman@umich.edu

We investigate the practical value of using graphical models to aid in two fundamental problems of group coordination: (1) belief aggregation and (2) risk sharing. We identify restrictive conditions under which graphical models can be useful in both settings. We show that the output of the logarithmic opinion pool (LogOP) can be represented as a Markov network (MN) or a decomposable Bayesian network (BN), and give an algorithm for doing so. We show that a securities market structured like a decomposable BN can support optimal risk sharing, if all agents have exponential utility and all of their Markov independencies coincide with the market structure. On the other hand, most of our results are negative, taking the form of impossibility theorems. We show that no belief aggregation function can maintain all independencies representable in a BN. Neither can an aggregation computation be decomposed into local computations on graph subsets. We show that computing query outputs of LogOP or the linear opinion pool (LinOP) is NP-hard. Except in fairly restrictive settings, structuring securities markets according to unanimously agreed upon independencies may be of no help in supporting optimal risk sharing because agents' behavioral independencies change as they engage in securities trade.

Key words: belief aggregation; opinion pools; consensus probabilities; risk sharing and allocation; securities markets; Bayesian networks; Markov networks; graphical models; multiagent systems
History: Received on June 15, 2004. Accepted by Eric Horvitz on September 12, 2005 after 1 revision.

1. Introduction

The normative basis of Bayesian decision making, set forth by von Neumann and Morgenstern (1953) and Savage (1954), forms a satisfying theory of rational decision making by an individual. More recent *graphical models* like Bayesian networks (BNs), Markov networks (MNs), influence diagrams, and CP-nets help make this normative framework more practical in two ways: semantically and computationally. Semantically, graphical models make the process of belief and preference elicitation more natural and efficient. Computationally, graphical models allow complex probability distributions to be represented more concisely, and allow posterior beliefs and optimal decisions to be computed more efficiently.

The setting of *group coordination* is also well studied across a number of disciplines, including statistics, decision theory, economics, political science, and artificial intelligence. However, unlike in the individual setting, there is nothing close to a single well-accepted normative basis for group beliefs, group preferences, or group decision making. Common to nearly all formal treatments of group coordination is an unsettling number of paradoxes and impossibility theorems. Within this muddied landscape, we examine the use of graphical models to aid two special-case group coordination problems: belief aggregation and risk sharing. We show that graphical models can be of practical use in both settings under certain restrictive conditions. However, our results are generally negative, surfacing new paradoxes, limitations, and impossibility theorems.

Belief aggregation is the task of forming a group consensus probability distribution by combining the beliefs of the individual members of the group in some fashion. The two most common opinion pool functions are (1) the linear opinion pool (LinOP), or weighted arithmetic average, and (2) the logarithmic opinion pool (LogOP), or weighted geometric average.

The most common form of risk sharing is insurance, where an individual pays a small fixed amount to an insurance company in return for an agreement to receive a larger payment if and only if (iff) some contingency is met in the future, for example the insured's car is stolen. More generically, individuals exchange *securities*, or state-contingent payoffs of the form "X dollars iff event A happens," whenever such an exchange is mutually beneficial. In equilibrium, the allocation of risk is *Pareto optimal*, meaning that there is no further possible exchange that would harm no one and help at least one individual. The standard securities market, called an Arrow-Debreu securities market, is unstructured and "flat," containing one security for every possible state of nature, or equivalently one security for all (exponentially many) possible combinations of event outcomes (Arrow 1964). We show that securities markets can be structured in analogy to BNs, so that hedging and speculating activity can be conducted more naturally at the event level rather than at the state level.

We investigate the value of using a graphical model to represent the aggregate probability distribution of a group of individuals. Under restrictive circumstances, graphical models do provide practical benefits. In particular, if the aggregate distribution is a LogOP, then any shared Markov independencies among the group are reflected in the aggregate distribution, and so can be represented in a relatively concise and natural manner using a BN or MN.

Similarly for risk sharing, we give strict conditions under which a securities market that is structured like a BN can support a Pareto optimal allocation of risk. A structured securities market may require exponentially fewer securities than the standard unstructured Arrow-Debreu securities market, thus greatly reducing the cognitive/computational burden on individuals and the auctioneer.

Unfortunately, under generic conditions, graphical models are of little help in belief aggregation. Even if every individual's beliefs can be represented concisely with a graphical model, it is unlikely that their aggregate distribution can be represented concisely. In particular, if the aggregate distribution is a LinOP, it will in general require a fully connected graphical model to represent it. More generally, we prove that no reasonable aggregation rule (not even LogOP) is able to maintain all agreed upon independencies represented in a set of BNs. Also, no aggregation rule can operate locally within the variable families of a BN. We show that computing a query (some conditional probability) in both the LinOP and LogOP aggregate distributions is NP-hard, even when the analogous query is easy to compute in all of the component BNs.

Most of our risk sharing results are negative as well. Except in very special circumstances, structured securities markets are required to be fully connected—and thus no more compact than Arrow-Debreu markets—in order to support Pareto-optimal risk allocations.

This paper is organized as follows. Section 2 presents the prerequisite background material and introduces our notation. Section 3 describes the use of graphical models for belief aggregation. Section 4 explores the value of graphical models for risk sharing, paralleling several of the positive and negative results from §3. We conclude in §5.

2. Background and Notation

We consider a group of N agents, indexed $i = 1, 2, \ldots, N$, each with a subjective probability distribution $\Pr_i$ over states of the world. Denote the set of all possible states of the world as $\Omega = \{\omega_1, \omega_2, \ldots\}$. The ω are mutually exclusive and exhaustive.

State is often more concisely and naturally characterized as the set of outcomes of *events*. Denote the set of modeled events as $Z = \{A_1, A_2, \ldots, A_M\}$. Underlying M arbitrary events is a state space Ω of size $|\Omega| = 2^M$, consisting of all possible combinations of event outcomes. Conversely, any set of states can be factored into a set of $M = \lceil \lg |\Omega| \rceil$ events. Without further assumption, the two representations are equivalent in both expressivity and size, although the event factorization may be more natural. In most of what follows, the events $\{A_j\}$ are the focus of attention, with Ω the implied joint outcome space. We refer to the $\{A_j\}$ as the *primary events*, so as to distinguish them from the other $2^{2^M} - M$ possible sets of states, each of which is also an event.

2.1. Opinion Pools

An opinion pool is a function that maps a set of probability distributions, $\Pr_1, \Pr_2, \ldots, \Pr_N$, to a single aggregate distribution, denoted $\Pr_0$:

$$\Pr_0 \equiv f(\Pr_1, \Pr_2, \ldots, \Pr_N). \tag{1}$$

A variety of authors have proposed or advocated a corresponding variety of aggregation functions of the form (1). Genest and Zidek (1986) and French (1985) provide comprehensive surveys. We distinguish four different approaches to belief aggregation, categorized according to whether they are justified based on (1) normative grounds, (2) axiomatic grounds, (3) worst-case arguments, or (4) maximum entropy arguments.

The normative approach stays within the standard Bayesian decision-making framework. A single decision maker (real or fictitious, within or outside the group), called the *supra Bayesian*, is responsible for carrying out the aggregation. The supra Bayesian updates its beliefs (which are defined over the joint space of all individuals' beliefs) via Bayes's rule, given the "evidence" of everyone's beliefs. The resulting posterior is taken to be the consensus belief.

The axiomatic approach first posits desiderata that the aggregation function should satisfy, then derives the aggregation function implied by the chosen axioms (Dalkey 1975; Genest 1984a, b, c; Genest and Zidek 1986; Genest et al. 1986; Genest and Wagner 1987; Wagner 1984). The two most common and well-studied aggregation functions are LinOP and LogOP. The LinOP is a weighted arithmetic mean of the members' probabilities,

$$\Pr_0(\omega) = \sum_{i=1}^{N} \alpha_i \Pr_i(\omega), \tag{2}$$

and the LogOP is a normalized, weighted geometric mean,

$$\Pr_0(\omega) \propto \prod_{i=1}^{N} [\Pr_i(\omega)]^{\alpha_i}, \tag{3}$$

where the α_i are called *expert weights*, usually nonnegative numbers that sum to one. The LinOP and LogOP can actually be characterized as two instances of a parameterized family of weighted aggregation functions (Cooke 1991). Note that under certain assumptions, the LinOP and LogOP can each be interpreted as the outcome of a supra Bayesian update.

Both the LinOP and LogOP satisfy the following two seemingly incontrovertible assumptions.

Property 1 (Unanimity (UNAM)). If $\Pr_h(\omega) = \Pr_i(\omega)$ for all agents h and i, and for all states $\omega \in \Omega$, then $\Pr_0 = \Pr_1$.

Property 2 (Nondictatorship (ND)). There is no single agent i such that $\Pr_0(\omega) = \Pr_i(\omega)$ for all $\omega \in \Omega$ regardless of the agents' beliefs.

UNAM states that if everyone's assessments are in complete agreement, then the consensus agrees as well. ND simply ensures that what is inherently a multiagent problem is not reduced to the single-agent case.

UNAM, ND, and the following *marginalization* property together form a necessary and sufficient set of axioms for LinOP (Genest 1984c).

Property 3 (Marginalization Property (MP)). Let E be an arbitrary event, that is, any subset of Ω. Then,

$$f(\Pr_1, \Pr_2, \ldots, \Pr_n)(E) = f(\Pr_1(E), \Pr_2(E), \ldots, \Pr_n(E)).$$

UNAM, ND, and the following *externally Bayesian* property form a necessary and sufficient set of axioms for LogOP (Genest 1984a).

Property 4 (External Bayesianity (EB)). Let E and $F \neq \varnothing$ be arbitrary events. Then,

$$\begin{aligned} f(\Pr_1, \Pr_2, \ldots, \Pr_n)(E \mid F) \\ = f(\Pr_1 \mid F, \Pr_2 \mid F, \ldots, \Pr_n \mid F)(E). \end{aligned}$$

MP and EB require consistency for probabilistic operations performed before and after pooling. MP states that we obtain the same probability for an event E whether we pool the opinions first, and then compute $\Pr_0(E) = \sum_{\omega \in E} \Pr_0(\omega)$, or if we first compute $\Pr_i(E) = \sum_{\omega \in E} \Pr_i(\omega)$ for each agent i, and then pool their opinions only over E. Similarly, EB holds that we obtain the same $\Pr_0(E \mid F)$ whether we combine opinions first and condition on F second, or condition on F first and combine opinions second. Genest (1984b) shows that no f can simultaneously satisfy MP, EB, UNAM, and ND.

Property 5 (Proportional Dependence on States (PDS)):

$$\Pr_0(\omega) \propto f(\Pr_1(\omega), \Pr_2(\omega), \ldots, \Pr_n(\omega)).$$

PDS is sometimes called *independence of irrelevant states*, or termed a *likelihood principle*. It assures that the consensus likelihood ratio between two states does not depend on the agents' assessments of any other "irrelevant" state. The LinOP, LogOP, and most other proposed opinion pools satisfy PDS.

Property 6 (Independence Preservation Property (IPP)). Let E and F be arbitrary events. If $\Pr_i(E \mid F) = \Pr_i(E)$ for all agents i, then $\Pr_0(E \mid F) = \Pr_0(E)$.

IPP requires that all unanimously held independencies are preserved in the consensus. Advocates of IPP reason that identifying the independencies in a model is central to understanding the underlying phenomena, and that complete agreement on this dimension should be embraced. On the other hand, Genest and Wagner (1987) make a compelling case *against* the use of IPP by proving that *no* aggregation function whatsoever can satisfy it along with PDS and ND, when $|\Omega| \geq 5$.

A third approach to belief aggregation focuses on the worst-case behavior of the aggregation rule in the spirit of most computational complexity results. The idea is to find an aggregation rule such that, no matter what experts actually say and how outcomes actually unfold, the accuracy of the consensus is guaranteed to be within some factor of the accuracy of the best individual expert. The setting is framed as a sequential prediction task, where each expert gives probabilities for a series of events, and the consensus rule attempts to learn—usually via a dynamic reweighting procedure—which experts to trust and which to ignore. A number of articles in the computer science theory and machine learning literature derive increasingly tighter bounds on the worst-case accuracy loss of the consensus as compared to the best expert; see for example Cesa-Bianchi et al. (1997).

A fourth approach to pooling opinions is based on maximum entropy inference. The consensus is the unique probability distribution that maximizes Shannon entropy, chosen from among the distributions that are consistent with all available information, including the experts' beliefs, their past performance, and/or dependencies among experts (Levy and Delic 1994, Myung et al. 1996).

The debate over which belief aggregation method is best continues (Benediktsson and Swain 1992, Cooke 1991, Jacobs 1995, Ng and Abramson 1992, Winkler 1986). Several authors (most emphatically Lindley (1985, 1988)) argue that the supra Bayesian approach is superior, as it is grounded in standard normative Bayesian theory (Clemen and Winkler 1993; Morris 1974, 1977; Rosenblueth and Ordaz 1992; West and Crosse 1992; Winkler 1981). Among the axiomatic approaches, many of the proposed axioms seem reasonable, but disagreement persists on which are essential. For example, Lindley (1985) regards the *marginalization property* as an "adhockery," while Cooke (1991) characterizes any consensus function that does *not* respect it as "downright queer." Another point of contention is how best to determine the expert weights. In most cases they are chosen in a somewhat ad hoc manner to encode some measure of confidence, reliability, or importance (Benediktsson and Swain 1992, French 1985, Winkler 1968). Some more formal methods to derive weights have been proposed by making assumptions concerning the form of, or interdependence among, participants' beliefs (Cooke 1991, Degroot and Mortera 1991, Jacobs 1995, Morris 1977), or through iterative self-weighting procedures (Degroot 1974). In the worst-case framework, authors have proven that certain reweighting policies guarantee that the resulting consensus accuracy is not much worse than the best expert's accuracy.

2.2. Graphical Models

A joint probability distribution can often be represented more compactly as a BN. Conciseness is achieved by exploiting conditional independence among the primary events. Let $\mathrm{CI}[A_j, W, X]$ be shorthand for $\Pr(A_j \mid WX) = \Pr(A_j \mid W)$, indicating that A_j is conditionally independent of the set of events X given another set W. Consider the event $A_k \in Z$, with *predecessors* $\mathbf{pred}(A_j) \equiv \{A_1, A_2, \ldots, A_{k-1}\}$, where the predecessor relation is with respect to the index ordering of variables. Suppose that given the outcomes of a subset $\mathbf{pa}(A_k) \subseteq \mathbf{pred}(A_k)$ of its predecessors—called A_k's *parents*—the event A_k is conditionally independent of all other preceding events, or $\mathrm{CI}[A_k, \mathbf{pa}(A_k), \mathbf{pred}(A_k) - \mathbf{pa}(A_k)]$. This structure can be depicted graphically as a *directed acyclic graph* (DAG): Each event is a node in the graph, and there is a directed edge from node A_j to node A_k if and only if A_j is a parent of A_k. We also refer to A_k as the *child* of A_j. A DAG has no directed cycles and thus defines a partial order over its vertices. By construction, the event indices are consistent with this partial ordering; in other words, if A_j is a parent of A_k then $j < k$. We can write the joint probability distribution in a (usually) more compact form:

$$\Pr(A_1 A_2 \cdots A_M) = \prod_{k=1}^{M} \Pr(A_k \mid \mathbf{pa}(A_k)).$$

For each event A_k, we record a *conditional probability table* (CPT), which contains probabilities $\Pr(A_k \mid \mathbf{pa}(A_k))$ for all possible combinations of outcomes of events in $\mathbf{pa}(A_k)$. Thus, it is possible to implicitly represent the full joint with $O(M \cdot 2^{\max\{q(k)\}})$ probabilities, instead of $2^M - 1$, where $q(k) = |\mathbf{pa}(A_k)|$ is the number of parents of A_k.

A MN is another graphical framework for modeling conditional independence and for implicitly describing joint distributions (Whittaker 1990, Darroch et al. 1980). Events are again associated with nodes in a graph, and edges encode probabilistic dependencies. The underlying structure of a MN is an *undirected* graph. Given the outcomes of its direct neighbors, an event A_j is conditionally independent of *every* other event in the network, not just preceding events. The neighbors of an event form a *Markov blanket* around it, "shielding" it from direct influence from the rest of the events (Pearl 1988).

A *Markov independence* is a special type of conditional independence (Darroch et al. 1980, Pearl 1988, Whittaker 1990). The node A_j and the set of nodes $X \subseteq Z - A_j$ are Markov independent, given another set $W \subseteq Z - X - A_j$, if $\mathrm{CI}[A_j, W, X]$ and $A_j \cup W \cup X = Z$. Recall that Z is the set of all modeled events.

The Markov blanket of a node in a BN consists of its direct parents, its direct children, and its children's direct parents (Pearl 1988). Therefore a BN can be converted into a MN by *moralizing* the network, or fully connecting ("marrying") each node's parents, and dropping edge directionality (Lauritzen and Spiegelhalter 1988, Neapolitan 1990). A MN can be converted into a BN by *filling in* or *triangulating* (Kloks 1994) the graph, and adding directionality according to the fill-in ordering (Jensen 1996, Lauritzen and Spiegelhalter 1988, Neapolitan 1990, Pearl 1988). With each transformation, independence information may be lost.

A DAG is *decomposable* if there is an edge between every two nodes that share a common child (Chyu 1991, Darroch et al. 1980, Pearl 1988, Shachter et al. 1991). Trees and complete graphs are two examples of decomposable graphs. Any BN can be made decomposable by reorienting some edges and introducing new edges where needed (Chyu 1991, Shachter et al. 1991). A two-step procedure of moralization plus fill-in (triangulation) will render a BN decomposable. Finding the smallest decomposable representation (finding the optimal fill-in ordering) is NP-hard, and even the smallest decomposable representation can be exponentially larger than the original BN. Still, the decomposable representation can be exponentially more compact than the full joint distribution. The independencies encoded in a decomposable BN are all Markov independencies (Pearl 1988).

3. Graphical Models for Belief Aggregation

Given their efficacy for single-agent belief representation, we consider whether graphical models can also enable the concise representation of aggregate beliefs.

3.1. Property Definitions

Recall the independence preservation property (IPP), defined in §2.1. For an aggregation function f to satisfy IPP, any independencies that are agreed upon by all agents must be maintained within the consensus distribution. Genest and Wagner (1987) prove that *no* aggregation function can simultaneously satisfy IPP, proportional dependence on states (PDS), and nondictatorship (ND). However, one might argue that IPP is overly strong. It requires preservation of, for example, a unanimous independence between the events $E = A_3\bar{A}_7$ and $F = \bar{A}_2 A_4 \vee A_7$. This kind of independence seems of little descriptive value to a modeler, and indeed cannot be represented with a BN. One may be willing to forgo preserving *all* independencies, being content to preserve independencies among the *primary* events, $A_1, A_2, \ldots, A_M$. With this in mind, we define a weaker independence property.

Property 7 (Event Independence Preservation Property (EIPP)). If for $A_j, A_k \in Z$, $\Pr_i(A_j \mid A_k) = \Pr_i(A_j)$ for all agents i, then $\Pr_0(A_j \mid A_k) = \Pr_0(A_j)$.

In §3.2, we see that substituting EIPP for IPP does admit a possibility that is consistent with both PDS and ND, although not a very satisfactory one. In search of a nontrivial possibility, we define an independence condition based on Markov independence.

Property 8 (Markov Independence Preservation Property (MIPP)). If for $A_j \in Z$, $X \subseteq Z - A_j$, and $W = Z - A_j - X$, $\Pr_i(A_j \mid WX) = \Pr_i(A_j \mid W)$ for all agents i, then $\Pr_0(A_j \mid WX) = \Pr_0(A_j \mid W)$.

Finally, we define a property that captures what seems to be a natural assumption within the context of graphical models, advocated independently by other authors (Matzkevich and Abramson 1992). We say that an aggregator satisfies the *family aggregation* (FA) property if it operates locally, within each conditional probability table (CPT) of the consensus structure.

Property 9 (Family Aggregation (FA)).

$$\Pr_0(A_j \mid \mathbf{pa}(A_j)) = f(\Pr_1(A_j \mid \mathbf{pa}(A_j)), \ldots, \Pr_N(A_j \mid \mathbf{pa}(A_j))),$$

where the $\mathbf{pa}(\cdot)$ operator is defined for the consensus BN representing $\Pr_0$.

FA seems especially reasonable when all agents' BNs have identical topologies.

3.2. Combining Bayesian Networks: Examples and Impossibility

3.2.1. Event Independence Preservation.

Example 1 (EIPP and the LinOP). Suppose that two agents agree that two primary events, A_1 and A_2, are independent, as pictured in Figure 1(a), but disagree on the associated marginal probabilities.

For concreteness, let the first agent hold beliefs $\Pr_1(A_1) = \Pr_1(A_2) = 0.5$, and the second $\Pr_2(A_1) = 0.8$ and $\Pr_2(A_2) = 0.6$. Thus,

$$\Pr_1(A_1A_2) = 0.25 \qquad \Pr_2(A_1A_2) = 0.48$$
$$\Pr_1(A_1\bar{A}_2) = 0.25 \qquad \Pr_2(A_1\bar{A}_2) = 0.32$$
$$\Pr_1(\bar{A}_1A_2) = 0.25 \qquad \Pr_2(\bar{A}_1A_2) = 0.12$$
$$\Pr_1(\bar{A}_1\bar{A}_2) = 0.25 \qquad \Pr_2(\bar{A}_1\bar{A}_2) = 0.08.$$

Now if we apply the LinOP (2) with, say, equal weights of $w_1 = w_2 = 0.5$, we get:

$$\Pr_0(A_1A_2) = 0.365$$
$$\Pr_0(A_1\bar{A}_2) = 0.41$$

Figure 1 **Independence Preservation Behavior of (a) LinOP and (b)–(d) LogOP**

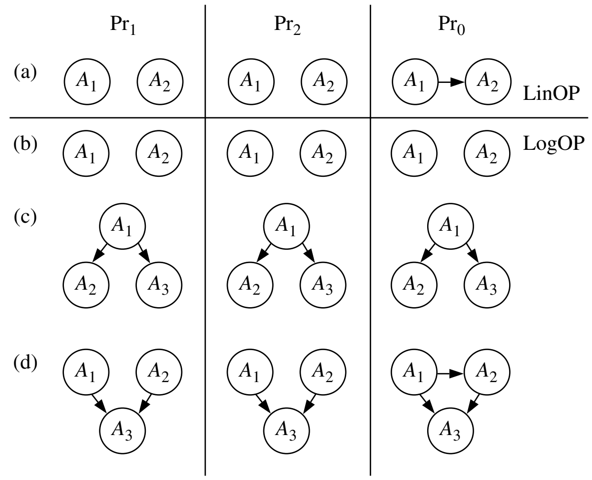

Note. If two agents' beliefs $\Pr_1$ and $\Pr_2$ have the dependency structures shown, then the consensus $\Pr_0$ will in general have the dependency structure depicted in column three.

$$\Pr_0(\bar{A}_1A_2) = 0.185$$
$$\Pr_0(\bar{A}_1\bar{A}_2) = 0.165.$$

In particular, $\Pr_0(A_1)\Pr_0(A_2) \neq \Pr_0(A_1A_2)$, and so the two events are *not* independent in the consensus.[1] Even though the precondition of the EIPP is met, the postcondition is not: A BN representation of the derived consensus would have to include an edge between A_1 and A_2, as depicted in the third column of Figure 1(a).

Example 2 (EIPP and the LogOP). Suppose that two agents' beliefs over two primary events are as described in Example 1. If we apply the LogOP with equal weights, we get:

$$\Pr_0(A_1A_2) = 0.367007$$
$$\Pr_0(A_1\bar{A}_2) = 0.29966$$
$$\Pr_0(\bar{A}_1A_2) = 0.183503$$
$$\Pr_0(\bar{A}_1\bar{A}_2) = 0.14983.$$

In this case, $\Pr_0(A_1)\Pr_0(A_2) = \Pr_0(A_1A_2)$, and the two events remain independent, as shown in Figure 1(b). This is not a numerical coincidence; in fact, independence between only two events is always maintained

[1] As early as Yule (1903) it was recognized that averaging two distributions may mask a commonly held independence.

by the LogOP (Genest and Wagner 1987). Now suppose that among three primary events, both agents agree that A_3 is independent of A_2 given A_1. That is, both agents agree that dependencies conform to a tree structure, with A_1 the parent of both A_2 and A_3, as depicted in Figure 1(c). Then once again, the LogOP will maintain this structure. LogOP does not, however, maintain all BN structures. For example, suppose that, among three primary events, the two agents agree that A_1 and A_2 are mutually independent, and that A_3 depends on both A_1 and A_2. That is, both agents agree on the polytree structure in Figure 1(d). In this case, when we compute the consensus with the LogOP, A_1 and A_2 will in general become mutually dependent, the EIPP is not satisfied, and a consensus BN will require an arc between the two nodes.

Having seen that both the LinOP and the LogOP violate the EIPP, we seek a more general characterization of the class of functions that obey it. We begin by adapting a result originally proved with respect to the IPP (Genest and Wagner 1987, Lemma 3.2 in) to the weaker condition EIPP.

Lemma 1 (Adapted from Genest and Wagner 1987). *If f obeys EIPP and PDS, then there exist constants $\alpha_1, \alpha_2, \ldots, \alpha_N$, and c such that*

$$\Pr_0(\omega_j) = \sum_{i=1}^{N} \alpha_i \Pr_i(\omega_j) + c. \tag{4}$$

Proof. Consider three events A_1, A_2, and A_3, with agents' beliefs described as follows:

$$\begin{aligned} \Pr_i(A_1A_2A_3) = \Pr_i(A_1A_2\bar{A}_3) &= \frac{(1-z_i)^2}{4(1+z_i)} \\ \Pr_i(A_1\bar{A}_2A_3) = \Pr_i(A_1\bar{A}_2\bar{A}_3) &= \frac{1-z_i}{4} \\ \Pr_i(\bar{A}_1\bar{A}_2A_3) &= x_i \\ \Pr_i(\bar{A}_1\bar{A}_2\bar{A}_3) &= y_i, \end{aligned} \tag{5}$$

where $z_i = x_i + y_i$ for all i. In this case, all agents agree that A_1 and A_2 are independent and, as long as $z_i < 1$, these equations describe a legal probability distribution. Because f obeys PDS, there must be some function g such that,

$$\Pr_0(\bar{A}_1\bar{A}_2A_3) = \frac{g(x_1, x_2, \ldots, x_N)}{\sum_{k=1}^{8} g(\Pr_1(\omega_k), \ldots, \Pr_N(\omega_k))},$$

and similarly for $\Pr_0(\bar{A}_1\bar{A}_2\bar{A}_3)$. Now imagine a second situation exactly as in (5), except with $\Pr_i(\bar{A}_1\bar{A}_2A_3) = x'_i$ and $\Pr_i(\bar{A}_1\bar{A}_2\bar{A}_3) = y'_i$. Genest and Wagner show that, as long as $x_i + y_i = x'_i + y'_i < 1$, then

$$\begin{aligned} &g(x_1, x_2, \ldots, x_N) + g(y_1, y_2, \ldots, y_N) \\ &\quad = g(x'_1, x'_2, \ldots, x'_N) + g(y'_1, y'_2, \ldots, y'_N). \end{aligned} \tag{6}$$

From here, they show that *because x_i and y_i can be chosen arbitrarily* (as long as their sum is less than one), then f must have the linear form specified in (4). □

Genest and Wagner go on to show, without further assumption, that f must be a dictatorship. However, that proof does *not* carry through under the weaker condition EIPP. This can be seen via a simple counterexample. Let f always ignore the agents' opinions, and simply assign a uniform distribution over all $\omega \in \Omega$. In this case, the consensus distribution holds that *all* primary events A_j are independent, and thus any agreed upon independencies are trivially maintained. One might wonder whether EIPP admits any other, more appealing, aggregation functions. The following proposition essentially establishes that it does not.

Proposition 2. *No aggregation function f can simultaneously satisfy EIPP, PDS, UNAM, and ND.*

Proof. With the addition of UNAM, it is clear that c must be zero in (4), and thus f must have the form of a standard LinOP (2). From Example 1, we know that the LinOP does not maintain independence even between just two events. The fact that the LinOP cannot satisfy both IPP and ND is proved formally by several authors (Genest 1984c, Lehrer and Wagner 1983, Wagner 1984). Their proofs are applicable to EIPP as well because they hold even when $|\Omega| = 4$, in which case EIPP and IPP coincide. □

3.2.2. Family Aggregation.

Example 3 (Family Aggregation). Consider two agents, each with a BN consisting of two primary events, with A_1 the parent of A_2 and with beliefs as follows:

$$\begin{aligned} \Pr_1(A_1) &= 0.2 & \Pr_2(A_1) &= 0.8 \\ \Pr_1(A_2 \mid A_1) &= 0.4 & \Pr_2(A_2 \mid A_1) &= 0.8 \\ \Pr_1(A_2 \mid \bar{A}_1) &= 0.6 & \Pr_2(A_2 \mid \bar{A}_1) &= 0.3. \end{aligned}$$

We compute each consensus CPT as an average of the corresponding individual CPTs. That is, $\Pr_0(A_1) = (0.2 + 0.8)/2 = 0.5$, $\Pr_0(A_2 \mid A_1) = (0.4 + 0.8)/2 = 0.6$, etc. This results in the following consensus joint distribution:

$$\Pr_0(A_1A_2) = 0.3$$
$$\Pr_0(A_1\bar{A}_2) = 0.2$$
$$\Pr_0(\bar{A}_1A_2) = 0.225$$
$$\Pr_0(\bar{A}_1\bar{A}_2) = 0.275.$$

Next suppose that both agents reverse their edge between the two events, such that A_2 is the parent of A_1, but that their joint distributions remain unchanged. Now the agents' CPTs are:

$$\Pr_1(A_2) = 0.56 \qquad \Pr_2(A_2) = 0.7$$
$$\Pr_1(A_1 \mid A_2) = 0.142857 \qquad \Pr_2(A_1 \mid A_2) = 0.914286$$
$$\Pr_1(A_1 \mid \bar{A}_2) = 0.272727 \qquad \Pr_2(A_1 \mid \bar{A}_2) = 0.533333$$

and if we average locally within each CPT, we get a different consensus distribution:

$$\Pr_0(A_1A_2) = 0.333$$
$$\Pr_0(A_1\bar{A}_2) = 0.149121$$
$$\Pr_0(\bar{A}_1A_2) = 0.297$$
$$\Pr_0(\bar{A}_1\bar{A}_2) = 0.220878.$$

Thus, averaging only within each family of the BN violates the form of the opinion pool itself (1), which insists that the consensus joint distribution depend only on the underlying joint distributions of the agents involved.

We now show that this inconsistency is not confined solely to the averaging aggregator.

Proposition 3. *No aggregation function f can simultaneously satisfy FA, UNAM, and ND.*

Proof. Let the first event in the consensus BN be A_{j_1} the second A_{j_2}, …, and the last A_{j_M}. The FA property requires both of the following:

$$\Pr_0(A_{j_1}) = f(\Pr_1(A_{j_1}), \Pr_2(A_{j_1}), \ldots, \Pr_N(A_{j_1})) \qquad (7)$$

$$\Pr_0(A_{j_M} \mid Z - A_{j_M}) = f(\Pr_1(A_{j_M} \mid Z - A_{j_M}), \ldots, \Pr_N(A_{j_M} \mid Z - A_{j_M})). \qquad (8)$$

By the definition of an opinion pool (1), the consensus belief depends only on the agents' underlying joint distributions, and not on the particular ordering of events in each BN. Thus, we must arrive at the same consensus distribution as long as $\{j_1, j_2, \ldots, j_M\}$ is some permutation of $\{1, 2, \ldots, M\}$. Consider two permutations, one where $j_1 = 1$ and one where $j_M = 1$. Then (7) and (8) become:

$$\Pr_0(A_1) = f(\Pr_1(A_1), \Pr_2(A_1), \ldots, \Pr_N(A_1)) \qquad (9)$$

$$\Pr_0(A_1 \mid Z - A_1) = f(\Pr_1(A_1 \mid Z - A_1), \ldots, \Pr_N(A_1 \mid Z - A_1)). \qquad (10)$$

Dalkey (1975) proves that no function can simultaneously satisfy (9), (10), UNAM, and ND. Alternatively, the two equations essentially require that f satisfy both MP and EB, defined in §2.1, which Genest (1984b) shows are incompatible with UNAM and ND. □

3.3. The LogOP and Consensus Markov Networks

The results in §3.2 suggest that insisting on general event independence preservation has rather severe consequences. In this section, we see that preserving Markov independencies is in fact compatible with PDS, UNAM, and ND. Let A_j be a primary event, and $W \subseteq Z - A_j$ and $X = Z - W - A_j$ be sets of events. Then A_j is Markov independent of X given W if $\Pr(A_j \mid WX) = \Pr(A_j \mid W)$.

Proposition 4. *The LogOP satisfies MIPP.*

Proof. Because the LogOP is defined in terms of atomic states ω, we make use of the following two identities:

$$\Pr_0(A \mid WX) \equiv \frac{\Pr_0(AWX)}{\Pr_0(AWX) + \Pr_0(\bar{A}WX)}$$

$$\Pr_0(A \mid W) \equiv \frac{\sum_X \Pr_0(AWX)}{\sum_X \Pr_0(AWX) + \sum_X \Pr_0(\bar{A}WX)}$$

where $\sum_X$ represents a sum over all possible combinations of outcomes of events in the set X. Then we have that,

$$\Pr_0(A \mid WX) = \frac{\prod_{i=1}^N [\Pr_i(AWX)]^{\alpha_i}}{\prod_{i=1}^N [\Pr_i(AWX)]^{\alpha_i} + \prod_{i=1}^N [\Pr_i(\bar{A}WX)]^{\alpha_i}}$$

$$= \frac{\prod\left[\frac{\Pr_i(AW)\Pr_i(WX)}{\Pr_i(W)}\right]^{\alpha_i}}{\prod\left[\frac{\Pr_i(AW)\Pr_i(WX)}{\Pr_i(W)}\right]^{\alpha_i}+\prod\left[\frac{\Pr_i(\bar{A}W)\Pr_i(WX)}{\Pr_i(W)}\right]^{\alpha_i}}$$
$$= \frac{\prod[\Pr_i(AW)]^{\alpha_i}}{\prod[\Pr_i(AW)]^{\alpha_i}+\prod[\Pr_i(\bar{A}W)]^{\alpha_i}}$$
$$= \frac{\prod[\Pr_i(AW)]^{\alpha_i}}{\prod[\Pr_i(AW)]^{\alpha_i}+\prod[\Pr_i(\bar{A}W)]^{\alpha_i}}\cdot\frac{\sum_X\prod[\Pr_i(WX)]^{\alpha_i}}{\sum_X\prod[\Pr_i(WX)]^{\alpha_i}}$$
$$= \frac{\sum_X\prod[\Pr_i(AW)\Pr_i(WX)]^{\alpha_i}}{\sum_X\prod[\Pr_i(AW)\Pr_i(WX)]^{\alpha_i}+\sum_X\prod[\Pr_i(\bar{A}W)\Pr_i(WX)]^{\alpha_i}}$$
$$= \frac{\sum_X\prod\left[\frac{\Pr_i(AW)\Pr_i(WX)}{\Pr_i(W)}\right]^{\alpha_i}}{\sum_X\prod\left[\frac{\Pr_i(AW)\Pr_i(WX)}{\Pr_i(W)}\right]^{\alpha_i}+\sum_X\prod\left[\frac{\Pr_i(\bar{A}W)\Pr_i(WX)}{\Pr_i(W)}\right]^{\alpha_i}}$$
$$= \frac{\sum_X\prod[\Pr_i(AWX)]^{\alpha_i}}{\sum_X\prod[\Pr_i(AWX)]^{\alpha_i}+\sum_X\prod[\Pr_i(\bar{A}WX)]^{\alpha_i}}$$
$$= \frac{\sum_X\Pr_0(AWX)}{\sum_X\Pr_0(AWX)+\sum_X\Pr_0(\bar{A}WX)}$$
$$= \Pr_0(A\,|\,W). \quad \square$$

Suppose that each agent's belief is given as a MN, and we wish to generate a consensus MN structure that can encode the results of the LogOP. As discussed in §2.2, graph connectivity in a MN represents probabilistic dependence, and the neighborhood relation represents direct influence. For each node A_j, the set of its neighbors plays the role of W in Proposition 4, and all other nodes constitute the set X. The proposition ensures that, if all agents agree on a common MN structure, then the consensus distribution derived by the LogOP will respect the same structure. When agents are not in complete agreement on the structure, then the consensus can be represented as a MN defined by the union of all the individual MNs. In other words, there is an edge between A_j and A_k in the consensus MN if and only if there is an edge between those two nodes in at least one of the agents' MNs.

Given a collection of BNs, generating a consensus BN structure that is consistent with the LogOP is also relatively straightforward. We first convert each BN into a MN by moralizing the graphs, or fully connecting each node's parents and dropping edge directionality (Lauritzen and Spiegelhalter 1988, Neapolitan 1990). Next, we compute the union of the individual MNs, and finally we convert the resulting consensus MN back into a BN by filling in or triangulating the network, reintroducing directionality according to the fill-in order (Jensen 1996, Lauritzen and Spiegelhalter 1988, Neapolitan 1990, Pearl 1988).[2]

We have outlined how to derive consensus MN or BN *structures*. What about computing the associated probabilities? In §3.4, we give an algorithm for computing the probabilities in a consensus BN that is polynomial in the size of its CPTs. Note that, even when all agents agree on a BN structure, the size of the final representation may grow exponentially during fill in, and computing the union of the intermediate MNs when agents disagree will only exacerbate this problem. Nevertheless, even a decomposable representation can be exponentially smaller than the full joint distribution, and the most popular algorithms for exact Bayesian inference do operate on decomposable (i.e., moralized and triangulated) models in practice.

3.4. Computing LogOP and LinOP

Because the LinOP (2) and LogOP (3) are defined over atomic states, computing, for example, the consensus marginal probability of a single event involves, in the worst case, a summation over 2^{M-1} terms. Moreover, even computing the LogOP consensus for a single state requires a normalization factor that is itself a sum over all 2^M states. In this section, we see that if each agent's belief is represented as a BN, the LinOP and LogOP consensus for any probabilistic query can be computed more efficiently. In particular, for the LogOP, we can compute all of the CPTs of a consensus BN with time complexity $O(NM^2\cdot 2^{\max\{q(j)\}})$, where $q(j)$ is the number of parents of A_j in the consensus structure. Any particular LinOP query can be computed with time complexity $O(NM\cdot 2^q)$, where q is the maximum treewidth of any of the individual BNs.

3.4.1. LogOP. We focus first on the task of generating a LogOP-consistent consensus BN. We compute its structure as described in §3.3. Consider computing the CPT at A_j, that is, $\Pr_0(A_j\,|\,\mathbf{pa}(A_j))$ for all combinations of outcomes of events in $\mathbf{pa}(A_j)$. From Proposition 3, we know that simply combining each agent's

[2] We do not claim that these consensus structures are minimal, or even that LogOP is the preferred aggregation method. Our goal is more to guide a modeler's decision process by delineating what representations are consistent under what circumstances.

assessment of this conditional probability will not succeed in general. However, we *can* compute the *last* CPT, $\Pr_0(A_M \mid \mathbf{pa}(A_M))$, in terms of only the $\Pr_i(A_M \mid \mathbf{pa}(A_M))$, by computing the LogOP over the single event A_M:

$$\Pr_0(A_M \mid \mathbf{pa}(A_M)) = \frac{\prod_{i=1}^{N} [\Pr_i(A_M \mid \mathbf{pa}(A_M))]^{\alpha_i}}{\prod [\Pr_i(A_M \mid \mathbf{pa}(A_M))]^{\alpha_i} + \prod [\Pr_i(\bar{A}_M \mid \mathbf{pa}(A_M))]^{\alpha_i}}. \quad (11)$$

Because the LogOP satisfies EB, if we condition on *all* other events $Z - A_M$ in the network, then the LogOP over just A_M will return the same result as if we had computed the LogOP over all events, and then conditioned on $Z - A_M$. Equation 11 also reflects the fact that $\Pr_0(A_M \mid \mathbf{pa}(A_M)) = \Pr_0(A_M \mid Z - A_M)$ and $\Pr_i(A_M \mid \mathbf{pa}(A_M)) = \Pr_i(A_M \mid Z - A_M)$, by the semantics of the BNs.

We can compute the remainder of the CPTs in reverse index order. Assume that the CPTs $\Pr_0(A_k \mid \mathbf{pa}(A_k))$ have been calculated for all $k > j$, and that next we need to calculate $\Pr_0(A_j \mid \mathbf{pa}(A_j))$. To simplify the discussion, let A_j have exactly two children, A_k and A_l, with $j < k < l$; the analysis generalizes easily to more children (or one child). Because the BN is decomposable, its topology is a tree of cliques (Chyu 1991, Pearl 1988, Shachter et al. 1991), and A_k and A_l can either be in the same clique or in separate cliques, as depicted in Figure 2. Note that decomposability also ensures that A_j's neighbors, $A_l \cup A_k \cup \mathbf{pa}(A_j)$, constitute its Markov blanket. We can query each of the agent's BNs for the probabilities $\Pr_i(A_j \mid A_l \cup A_k \cup \mathbf{pa}(A_j))$ using a standard BN inference algorithm. From these, we can compute the corresponding consensus probability as a LogOP only over A_j, as before:

$$\Pr_0(A_j \mid A_l \cup A_k \cup \mathbf{pa}(A_j)) \propto \prod_{i=1}^{N} [\Pr_i(A_j \mid A_l \cup A_k \cup \mathbf{pa}(A_j))]^{\alpha_i}. \quad (12)$$

Figure 2 **Two Potential Sections of a Decomposable BN**

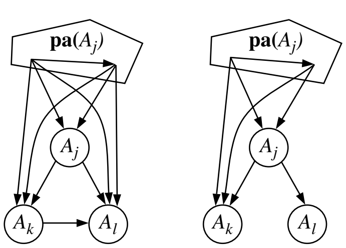

Note. A_j's children can be either in the same clique or in separate cliques.

We now need only eliminate the conditioning on A_l and A_k. By Bayes's rule, we have that

$$\frac{\Pr_0(A_j \mid A_l \cup A_k \cup \mathbf{pa}(A_j))}{\Pr_0(\bar{A}_j \mid A_l \cup A_k \cup \mathbf{pa}(A_j))} = \frac{\Pr_0(A_l \cup A_k \mid A_j \cup \mathbf{pa}(A_j))}{\Pr_0(A_l \cup A_k \mid \bar{A}_j \cup \mathbf{pa}(A_j))} \cdot \frac{\Pr_0(A_j \mid \mathbf{pa}(A_j))}{\Pr_0(\bar{A}_j \mid \mathbf{pa}(A_j))}$$
$$= \frac{\Pr_0(A_l \mid A_k \cup A_j \cup \mathbf{pa}(A_j))}{\Pr_0(A_l \mid A_k \cup \bar{A}_j \cup \mathbf{pa}(A_j))} \cdot \frac{\Pr_0(A_k \mid A_j \cup \mathbf{pa}(A_j))}{\Pr_0(A_k \mid \bar{A}_j \cup \mathbf{pa}(A_j))} \cdot \frac{\Pr_0(A_j \mid \mathbf{pa}(A_j))}{\Pr_0(\bar{A}_j \mid \mathbf{pa}(A_j))}.$$

Because the BN is decomposable, and regardless of whether A_k and A_l are in the same or different cliques, $\Pr_0(A_l \mid A_k \cup A_j \cup \mathbf{pa}(A_j)) = \Pr_0(A_l \mid \mathbf{pa}(A_l))$, and $\Pr_0(A_k \mid A_j \cup \mathbf{pa}(A_j)) = \Pr_0(A_k \mid \mathbf{pa}(A_k))$, both of which have already been computed. Therefore we can calculate the CPT at A_j as follows:

$$\frac{\Pr_0(A_j \mid \mathbf{pa}(A_j))}{\Pr_0(\bar{A}_j \mid \mathbf{pa}(A_j))} = \frac{\Pr_0(A_j \mid A_l \cup A_k \cup \mathbf{pa}(A_j))}{\Pr_0(\bar{A}_j \mid A_l \cup A_k \cup \mathbf{pa}(A_j))} \cdot \frac{\Pr_0(A_l \mid \widetilde{\mathbf{pa}}(A_l))}{\Pr_0(A_l \mid \mathbf{pa}(A_l))} \cdot \frac{\Pr_0(A_k \mid \widetilde{\mathbf{pa}}(A_k))}{\Pr_0(A_k \mid \mathbf{pa}(A_k))}, \quad (13)$$

where $\widetilde{\mathbf{pa}}(A_k)$ and $\widetilde{\mathbf{pa}}(A_l)$ contain $\bar{A}_j$, and $\mathbf{pa}(A_k)$ and $\mathbf{pa}(A_l)$ contain A_j. Once we compute the likelihood ratio on the LHS of (13), the desired probabilities are uniquely determined because $\Pr_0(A_j \mid \mathbf{pa}(A_j)) + \Pr_0(\bar{A}_j \mid \mathbf{pa}(A_j)) = 1$. The pseudocode for the full algorithm is given in Figure 3.

Next we characterize the computational complexity of LogOP. We assume that the input distributions are represented as BNs. We show that computing the LogOP of a query is NP-hard even when computing the same query in each BN is easy.

Proposition 5. *Given input distributions* $\Pr_1, \Pr_2, \ldots$ *represented as BNs, computing* $\Pr_0(A_j \mid \mathbf{pa}(A_j))$ *consistent with LogOP is NP-hard.*

Figure 3 Algorithm for Computing the CPTs of a LogOP-Consistent Consensus BN

LOGOP-CONSENSUS-BN($\Pr_1, \Pr_2, \ldots, \Pr_N$)

INPUT: N Bayesian networks: $\Pr_1, \Pr_2, \ldots, \Pr_N$

OUTPUT: LogOP-consistent consensus BN: $\Pr_0$

1. Structure of $\Pr_0 = \text{TRIANGULATE}[\bigcup_{i=1}^{N} \text{MORALIZE}[\Pr_i]]$
2. $\Pr_0(A_M \mid \mathbf{pa}(A_M)) \propto \prod_{i=1}^{N} [\Pr_i(A_M \mid \mathbf{pa}(A_M))]^{\alpha_i}$
3. **for** $j = M-1$ **downto** 1
4. $\Pr_0(A_j \mid A_l \cup A_k \cup \mathbf{pa}(A_j)) \propto \prod_{i=1}^{N} [\Pr_i(A_j \mid A_l \cup A_k \cup \mathbf{pa}(A_j))]^{\alpha_i}$
5. $\dfrac{\Pr_0(A_j \mid \mathbf{pa}(A_j))}{\Pr_0(\bar{A}_j \mid \mathbf{pa}(A_j))} = \dfrac{\Pr_0(A_j \mid A_l \cup A_k \cup \mathbf{pa}(A_j))}{\Pr_0(\bar{A}_j \mid A_l \cup A_k \cup \mathbf{pa}(A_j))} \cdot \dfrac{\Pr_0(A_l \mid \widetilde{\mathbf{pa}}(A_l))}{\Pr_0(A_l \mid \mathbf{pa}(A_l))} \cdot \dfrac{\Pr_0(A_k \mid \widetilde{\mathbf{pa}}(A_k))}{\Pr_0(A_k \mid \mathbf{pa}(A_k))}$

Proof. (Sketch.) We prove that even computing $\Pr_0(A_1)$ is NP-hard. First note that the $\Pr_i(A_1)$ are directly given in the input BN encodings. Suppose that $N = 2$. Let $\Pr_1$ be an arbitrary BN and let $\Pr_2$ encode $\Pr_2(A_M) = 1$ and $\Pr_2(\omega) = 1/2^{M-1}$ for all $\omega \in Z - A_M$ (that is, a uniform distribution over the events in $Z - A_M$). We have shown that, if $\Pr_0(A_1)$ were computable in polynomial time, then $\Pr_1(A_1 \mid A_M)$ could be inferred in polynomial time. Computing the latter query is NP-hard (Cooper 1990), and so the former must be as well. □

3.4.2. LinOP. A consensus BN consistent with the LinOP would in general be fully connected. However, if all agents' beliefs are given as BNs, we can retain their separation and still compute LinOP queries more efficiently. We exploit the fact that the LinOP obeys MP, and thus that the LinOP of any compound, marginal event can be computed as a LinOP over only that event. For example,

$$\Pr_0(A_2\bar{A}_5A_9) = \sum_{i=1}^{N} \alpha_i \Pr_i(A_2\bar{A}_5A_9),$$

where the terms on the right-hand side are calculated using a standard algorithm for Bayesian inference. Any conditional probability can be computed as the division of two compound, marginal probabilities.

Finally, we characterize the computational complexity of LinOP when all input models are BNs. Clearly, computing an arbitrary query $\Pr_0(E \mid F)$ is NP-hard because Bayesian inference itself is NP-hard (Cooper 1990). Proposition 6 establishes that, even when only computing the LinOP of a CPT entry, the problem remains intractable.

Proposition 6. *Given input distributions* $\Pr_1, \Pr_2, \ldots$ *represented as BNs, computing* $\Pr_0(A_j \mid \mathbf{pa}(A_j))$ *consistent with LinOP is NP-hard.*

Proof. (Sketch.) Suppose that $N = 2$. Let $\Pr_1$ be an arbitrary BN and let $\Pr_2$ encode a uniform distribution—that is, $\Pr_2(\omega) = 1/2^M$ for all $\omega \in \Omega$. We have shown that, if $\Pr_0(A_M \mid \mathbf{pa}(A_M))$ were computable in polynomial time, then $\Pr_1(A_M)$ could be inferred in polynomial time. Computing the later query is NP-hard (Cooper 1990), and so the former must be as well. □

3.5. Related Work

Faria and Smith (1996) examine a group decision-making situation where agents agree on a common decomposable BN structure and have identical preferences. They define a weaker form of EB, called *conditional external Bayesianity* (CEB), which requires EB to hold only for CPT entries, and only when evidence updates are based on *cutting* likelihood functions—those that can be factored according to the model structure. They show that a generalized LogOP, called a conditional LogOP, is the only pooling function that satisfies both CEB and UNAM. The conditional modified LogOP preserves the agreed upon structure and allows expert weights to vary across families in the structure. The authors also present an associated procedure for iteratively revising weights that reflects the relative alignment of the experts' predictions with actual observed outcomes.

Ng and Abramson (1994) describe an architecture called the *probabilistic multiknowledge-base system*, which consists of a collection of BNs, each encoding the knowledge of a single expert. The BNs are kept separate and probabilities are combined *at run time* with a variable-weight variant of the LinOP. The authors address a variety of engineering issues, including the elicitation and propagation of expert confidence information, and build a working prototype to diagnose pathologies of the lymph system. Xiang (1996) describes conditions under which *multiply sectioned Bayesian networks*, originally developed for single-agent reasoning, can represent the combined beliefs of multiple agents. The main assumption is that, whenever two agents' BNs contain some of the same events, they must agree on the joint distribution over these common events. Bonduelle (1987) prescribes both normative and behavioral techniques for a decision maker (DM) to identify and reconcile differences of opinion among experts. When those opinions are expressed as graphical models, he suggests that the DM first choose a consensus topology, and then calculate aggregate probabilities. Jacobs (1995) compares the LinOP and supra Bayesian approaches as methods for combining the multiple feature analyzers found in real and artificial neural systems.

Matzkevich and Abramson (1992) give an algorithm for explicitly combining two BN DAGs into a single DAG, or *fusing* the two topologies. The algorithm transfers one arc at a time from the second DAG to the first, possibly reversing the arc in order to remain consistent with the current partial ordering. Reversing arcs may add new arcs to the second DAG (Shachter 1988), which would in turn need to be transferred. In a second paper, Matzkevich and Abramson (1993) show that the task of minimizing the number of arcs in their combined DAG is NP-hard, as are several other related tasks. They argue that, intuitively, the consensus model should capture independencies agreed on by at least $c \leq n$ of the agents; in particular, when $c = n$ and the orderings are mutually consistent, the consensus DAG should be a union of the individual DAGs. In both of these papers, and in Bonduelle's work, it is essentially assumed that the EIPP, or a stronger version thereof, should hold.

Though Matzkevich and Abramson (1992, 1993) make no commitment on how to combine probabilities, they do give an example (Matzkevich and Abramson 1992) where the LinOP is applied *locally*, or separately within each CPT, thus satisfying the FA property. Although such a constraint on aggregation may seem natural, we see in §3.2 that it actually has very severe implications.

4. Graphical Models for Risk Sharing

We now turn to the group coordination problem of risk sharing. We show that several of the impossibility and possibility results obtained for belief aggregation in the previous section have implications for risk sharing.

We assume that risk is shared, or allocated, among agents via a process of exchanging *securities*, or financial contingent payoffs of the form "\$1 iff A," denoted $\langle A\rangle$. We assume that agents' utility functions $u(\mu)$ are expressed in terms of money or dollars μ; utility for securities is then expressed as expected utility for dollars.

A *complete* securities market contains at least $|\Omega|-1$ linearly independent[3] securities, to span the entire state space Ω. The canonical complete securities market is the Arrow-Debreu market, which contains one security $\langle\omega\rangle$ for every possible state $\omega \in \Omega$. Only a complete market can guarantee Pareto-optimal risk sharing. Nonetheless, we describe some special cases when fewer securities—potentially exponentially fewer than $|\Omega|$ securities—can suffice for Pareto optimality. We call such a market *operationally complete*.

We describe how securities markets can be structured like BN graphical models, allowing agents to trade securities defined in terms of primary events and conditionals, rather than in terms of base states. Paralleling the previous section, we show that in general, the structure of the market must be fully connected to support Pareto optimality. However in one interesting special case—when all agents have constant risk aversion—a market structured according to unanimous Markov independencies can be operationally complete and thus Pareto optimal.

4.1. Equilibrium in a Securities Market

Agents trade securities until *equilibrium*, or until all mutually beneficial exchanges are exhausted. Denote the price of security $\langle A\rangle$ as $p^{\langle A\rangle}$. Equilibrium prices in a securities market form a coherent probability distribution over events (Huang and Litzenberger 1988, Varian 1987). For example, $p^{\langle A_1\rangle} = p^{\langle A_1A_2\rangle} + p^{\langle A_1\bar{A}_2\rangle}$. In a strong sense, the market process yields an aggregate probability distribution (encoded as prices) that is a function of the agents' beliefs and utilities. All the limitative theorems on aggregation functions discussed in §§2 and 3 apply.

In special cases, the market aggregation function corresponds to known belief aggregation rules (Pennock and Wellman 1997, 2001; Wolfers and Zitzewitz 2005). For example, if all agents have generalized logarithmic utility, the price equilibrium is a LinOP (weighted average) of individual beliefs. If all agents have constant absolute risk aversion (CARA), or negative exponential utility, the price equilibrium

[3] By linear independence, we mean that the securities' payoff-triggering events are linearly independent.

is a LogOP (weighted geometric average) of individual beliefs. In these two cases, weights are normalized measures of risk tolerance.

4.2. Structured Markets: An Analogy to Bayesian Networks

To structure a market like a BN, we make use of *conditional securities*. A conditional security $\langle A_1 | A_2 \rangle$ pays off \$1 contingent on A_1 but *conditional* on A_2. That is, if A_2 occurs, then the security pays out exactly as $\langle A_1 \rangle$; on the other hand, if $\bar{A}_2$ occurs, then the bet is called off and any price paid for the security is refunded (de Finetti 1974). Stated another way, buying one unit of $\langle A_1 | A_2 \rangle$ for price p results in a payoff of $1-p$ if $A_1 A_2$, $-p$ if $\bar{A}_1 A_2$, and 0 if $\bar{A}_2$. The Arrow-Debreu complete market consists of one security paying out in each state of nature. In general, though, any set of securities (possibly including conditionals) with a payoff-by-state matrix of rank $|\Omega|-1$ is complete.

Securities markets can be structured according to the directed acyclic graph D of any BN. Simply introduce one conditional security $\langle A_j | \mathbf{pa}(A_j) \rangle$ for every conditional probability $\Pr(A_j | \mathbf{pa}(A_j))$ in the BN. For each event A_j with $q(j) = |\mathbf{pa}(A_j)|$ parents, this adds $2^{q(j)}$ securities, one for each possible combination of outcomes of events in $\mathbf{pa}(A_j)$. Call such a market *D-structured*. Imagine for the moment that D is fully connected (that is, no independencies are represented). Then a D-structured market contains $\sum_{j=1}^{M} 2^{j-1} = 2^M - 1 = |\Omega| - 1$ linearly independent securities, and is thus complete.

The main advantage of a structured market realized when D is less than fully connected, and thus the market contains less than $|\Omega|-1$ securities. What can be said in this case? Under fairly strict circumstances explored below, the smaller market may suffice for operational completeness.

4.3. Compact Securities Markets

Intuitively, we would like to structure a securities market according to independence relationships among primary events, and thus perhaps use fewer securities to span the space of possible outcomes. Assuming that all agents agree on a set of independence relationships, one might expect that a market structured according to these independencies would be operationally complete. However, in general, such a market is not operationally complete. The reason stems from the distinction between "true" subjective beliefs and observed (behaviorally inferred) beliefs.

True subjective beliefs are privately held and not directly observable. Beliefs that are elicited based on an agent's behavior—for example, the amount paid for securities—may differ from true beliefs depending on the agent's risk tolerance, risk exposure, and prior stakes. Observable beliefs are sometimes called *risk-neutral beliefs* because they are the beliefs that a behaviorally identical risk-neutral agent would have (Kadane and Winkler 1988, Nau and McCardle 1991, Nau 1995). We denote an agents' risk-neutral beliefs as $\Pr_i^{\mathrm{RN}}(\cdot)$. The distinction between true beliefs and risk-neutral beliefs carries over to independence relations. Even if an agent believes that two events are (conditionally) independent, if the agent has risk exposure to either event, that agent may behave outwardly as if the events were (conditionally) dependent. For example, an agent with risk exposure to A_1 may be willing to pay $p^{\langle A_1 \rangle}$ for $\langle A_1 \rangle$, $p^{\langle A_2 \rangle}$ for $\langle A_2 \rangle$, and $p^{\langle A_1 A_2 \rangle} < p^{\langle A_1 \rangle} p^{\langle A_2 \rangle}$ for $\langle A_1 A_2 \rangle$, even if the agent believes A_1 and A_2 are independent. Paralleling our notation for true conditional independence, let $\mathrm{CI}_i^{\mathrm{RN}}[A_j, W, X]$ denote the risk-neutral conditional independence $\Pr_i^{\mathrm{RN}}(A_j | WX) = \Pr_i^{\mathrm{RN}}(A_j | W)$.

If a securities market is D-structured, and in equilibrium all agents' risk-neutral independencies conform to D, then the market is operationally complete (Pennock and Wellman 2000). However, this condition for operational completeness seems of little practical value because risk-neutral independencies change as agents buy or sell securities on the path to equilibrium.

In the remainder of this section, we explore a condition on agents that supports operational completeness under unanimity of true independencies.

4.4. Markov Independency Markets

Suppose that all agents have *exponential utility*: $u_i(\mu) = -e^{-c_i \mu}$, also called *constant absolute risk aversion* (CARA), where c_i is agent i's coefficient of risk aversion, or $1/c_i$ its risk tolerance.

Define an *independency market*, or an *I-market*, as a *D-structured* market such that all agents' true distributions agree with the independencies in D. An I-market is *decomposable* if D is decomposable—every node's parents are fully connected.

Let $Z = \{A_1, \ldots, A_M\}$ be the set of all events, $A_j \in Z$ a particular event, and $W \subseteq Z - A_j$ and $X = Z - W - A_j$ subsets of events. We are interested in whether agent i's Markov independencies $\mathrm{CI}_i[A_j, W, X]$ are reflected as risk-neutral independencies $\mathrm{CI}_i^{\mathrm{RN}}[A_j, W, X]$, and are thus observable. For brevity, we drop the subscript i when only one agent is under consideration.

Proposition 7.

$$\mathrm{CI}[A_j, W, X] \;\&\; \left(\frac{u'(\Upsilon^{\langle \bar{A}_j W X\rangle})}{u'(\Upsilon^{\langle A_j W X\rangle})} = \frac{u'(\Upsilon^{\langle \bar{A}_j W \widetilde{X}\rangle})}{u'(\Upsilon^{\langle A_j W \widetilde{X}\rangle})}\right) \Rightarrow \mathrm{CI}^{\mathrm{RN}}[A_j, W, X], \tag{14}$$

where the second precondition must hold for all possible joint outcomes of the events in W, and all pairs $(X, \widetilde{X})$ of different joint outcomes of events in X.

Proof.

$$\frac{u'(\Upsilon^{\langle \bar{A}_j W X\rangle})}{u'(\Upsilon^{\langle A_j W X\rangle})} = \frac{u'(\Upsilon^{\langle \bar{A}_j W \widetilde{X}\rangle})}{u'(\Upsilon^{\langle A_j W \widetilde{X}\rangle})}$$

$$\Pr(A_j W) + \Pr(\bar{A}_j W)\frac{u'(\Upsilon^{\langle \bar{A}_j W X\rangle})}{u'(\Upsilon^{\langle A_j W X\rangle})} = \Pr(A_j W) + \Pr(\bar{A}_j W)\frac{u'(\Upsilon^{\langle \bar{A}_j W \widetilde{X}\rangle})}{u'(\Upsilon^{\langle A_j W \widetilde{X}\rangle})}$$

$$\frac{\frac{\Pr(A_j W)\Pr(WX)}{\Pr(W)} u'(\Upsilon^{\langle A_j W X\rangle})}{\frac{\Pr(A_j W)\Pr(WX)}{\Pr(W)} u'(\Upsilon^{\langle A_j W X\rangle}) + \frac{\Pr(\bar{A}_j W)\Pr(WX)}{\Pr(W)} u'(\Upsilon^{\langle \bar{A}_j W X\rangle})} = \frac{\frac{\Pr(A_j W)\Pr(W\widetilde{X})}{\Pr(W)} u'(\Upsilon^{\langle A_j W \widetilde{X}\rangle})}{\frac{\Pr(A_j W)\Pr(W\widetilde{X})}{\Pr(W)} u'(\Upsilon^{\langle A_j W \widetilde{X}\rangle}) + \frac{\Pr(\bar{A}_j W)\Pr(W\widetilde{X})}{\Pr(W)} u'(\Upsilon^{\langle \bar{A}_j W \widetilde{X}\rangle})}$$

$$\frac{\Pr(A_j W X) u'(\Upsilon^{\langle A_j W X\rangle})}{\Pr(A_j W X) u'(\Upsilon^{\langle A_j W X\rangle}) + \Pr(\bar{A}_j W X) u'(\Upsilon^{\langle \bar{A}_j W X\rangle})} = \frac{\Pr(A_j W \widetilde{X}) u'(\Upsilon^{\langle A_j W \widetilde{X}\rangle})}{\Pr(A_j W \widetilde{X}) u'(\Upsilon^{\langle A_j W \widetilde{X}\rangle}) + \Pr(\bar{A}_j W \widetilde{X}) u'(\Upsilon^{\langle \bar{A}_j W \widetilde{X}\rangle})}$$

$$\frac{\Pr^{\mathrm{RN}}(A_j W X)}{\Pr^{\mathrm{RN}}(A_j W X) + \Pr^{\mathrm{RN}}(\bar{A}_j W X)} = \frac{\Pr^{\mathrm{RN}}(A_j W \widetilde{X})}{\Pr^{\mathrm{RN}}(A_j W \widetilde{X}) + \Pr^{\mathrm{RN}}(\bar{A}_j W \widetilde{X})}$$

$$\Pr^{\mathrm{RN}}(A_j \mid WX) = \Pr^{\mathrm{RN}}(A_j \mid W\widetilde{X}). \quad \square$$

The second precondition in (14) is true if the agent is risk neutral and holds approximately if utility is state-independent and stakes are small. However, this approximation is not realistic for an agent engaged in trading securities because a central role of the market is precisely to enable the transfer of wealth across states.

Let $\Upsilon^{\langle A_j W\rangle}$ be the agent's payoff from all securities that depend only on events in $A_j \cup W$. Examples are $\langle A_j\rangle$, $\langle A_j W\rangle$, and $\langle A_j \mid W\rangle$, which return the same dollar amount regardless of the realizations of events in $X = Z - W - A_j$. Similarly, let $\Upsilon^{\langle WX\rangle}$ be the payoff from securities that do not depend on A_j.

Suppose that the agent exhibits CARA, and that its payoffs are separable according to $\Upsilon^{\langle A_j W X\rangle} = \Upsilon^{\langle A_j W\rangle} + \Upsilon^{\langle WX\rangle} - \Upsilon^{\langle W\rangle}$. Separability essentially means that any of the agent's securities (or prior stakes) whose payoff depends on A_j cannot also depend on events in X. In this case,

$$\begin{aligned}\frac{u'(\Upsilon^{\langle \bar{A}_j W X\rangle})}{u'(\Upsilon^{\langle A_j W X\rangle})} &= \frac{u'(\Upsilon^{\langle \bar{A}_j W\rangle} + \Upsilon^{\langle WX\rangle} - \Upsilon^{\langle W\rangle})}{u'(\Upsilon^{\langle A_j W\rangle} + \Upsilon^{\langle WX\rangle} - \Upsilon^{\langle W\rangle})} \\ &= \frac{ce^{-c\Upsilon^{\langle \bar{A}_j W\rangle}} e^{-c\Upsilon^{\langle WX\rangle}} e^{c\Upsilon^{\langle W\rangle}}}{ce^{-c\Upsilon^{\langle A_j W\rangle}} e^{-c\Upsilon^{\langle WX\rangle}} e^{c\Upsilon^{\langle W\rangle}}} \\ &= \frac{ce^{-c\Upsilon^{\langle \bar{A}_j W\rangle}} e^{-c\Upsilon^{\langle W\widetilde{X}\rangle}} e^{c\Upsilon^{\langle W\rangle}}}{ce^{-c\Upsilon^{\langle A_j W\rangle}} e^{-c\Upsilon^{\langle W\widetilde{X}\rangle}} e^{c\Upsilon^{\langle W\rangle}}} \\ &= \frac{u'(\Upsilon^{\langle \bar{A}_j W\rangle} + \Upsilon^{\langle W\widetilde{X}\rangle} - \Upsilon^{\langle W\rangle})}{u'(\Upsilon^{\langle A_j W\rangle} + \Upsilon^{\langle W\widetilde{X}\rangle} - \Upsilon^{\langle W\rangle})} \\ &= \frac{u'(\Upsilon^{\langle \bar{A}_j W \widetilde{X}\rangle})}{u'(\Upsilon^{\langle A_j W \widetilde{X}\rangle})}.\end{aligned}$$

Thus the constraint on utility in (14) is satisfied, and any Markov independencies are observable.

We are now in a position to derive the main result of this section.

Proposition 8. *When all agents have CARA, every decomposable I-market is operationally complete.*

Proof. Let W_j be the set of direct parents and direct children of event A_j, and X_j all other events. From decomposability and I-marketness, we can infer that:

1. $\mathrm{CI}_i[A_j, W_j, X_j]$ for all agents i and events j.
2. None of the securities $\langle A_j \mid \mathbf{pa}(A_j)\rangle$ that are contingent on A_j depend on X_j.

3. None of the securities $\langle A_k \mid \mathbf{pa}(A_k)\rangle$ such that $A_j \in \mathbf{pa}(A_k)$ that are conditional on A_j depend on X_j.

Items 2 and 3 ensure separability of payoffs from the available securities (we assume that any prior stakes are also separable). Then, invoking Proposition 7, $\mathrm{CI}_i^{\mathrm{RN}}[A_j, W_j, X_j]$ for all agents i and events j. As a result, all the independencies in D hold for every $\Pr_i^{\mathrm{RN}}$, regardless of allocations or prices, including those at equilibrium. Pennock and Wellman (2000) prove that a D-structured market conforming to unanimous risk-neutral independencies is operationally complete. □

5. Conclusion

Positive progress in research in group coordination, though certainly not lacking, is circumscribed by controversy and impossibilities. Contributions in this paper fall on both sides of the impossibility fence.

5.1. For the Pessimist...

A subset of results in this paper further confine and confound the search for reasonable belief aggregation and risk allocation procedures, by extending the impossibility theorems to new domains and by raising new concerns.

A series of theorems (Lehrer and Wagner 1983, Wagner 1984) culminating in that of Genest and Wagner (1987) show that very weak and reasonable constraints on an aggregation function are enough to rule out independence preservation (i.e., the retention of all agreed upon independencies within the aggregate distribution). However, these theorems apply to functions that preserve all possible independencies between any events—even those not representable in a graphical model. A potential loophole remained that some reasonable function might preserve the independencies among primary events in a graphical model. Indeed, in §3, we see that the same conditions sufficient to rule out general independence preservation are *not* sufficient to rule out this weaker form. However, we show that, with the additional (uncontroversial) assumption of unanimity, the impossibility returns.

This result resurfaces in the study of structured securities markets in §4. The intuitive inclination to structure the market according to agreed upon independencies proved fatally flawed. Prices in a securities market are essentially the output of an aggregation function—the one defined by market equilibrium—and are thus subject to all the general limitative theorems.

We derive a second impossibility theorem in §3 regarding the combination of BNs. A natural policy—that other authors have advocated or assumed—is to confine the aggregation locally, within each conditional probability table of the BN. We prove that any such local aggregation function necessarily fails to satisfy either unanimity or nondictatorship, two seemingly incontrovertible assumptions.

Other results demonstrate that desirable operations, while not impossible, are instead (worst-case) intractable. Someone interested in computing the LinOP of several probability distributions, each represented as a BN, would not want to construct a consensus BN, as it would in general be fully connected. This suggests keeping the individual BNs separate, and computing the LinOP of any desired query at runtime. In Proposition 6, we prove that performing this computation is NP-hard, even if answering the same query is easy within each individual BN. Similarly, computing LogOP is in general NP-hard.

Even the positive results in §3 describing BN representations of LogOP are shaded by potential computational barriers. The consensus network structure must be made decomposable, a process that can increase the size of the representation exponentially. Similar computational concerns arise in §4 when the structure of the securities market is required to be decomposable.

5.2. For the Optimist...

On the other hand, some results in this paper can be characterized as possibility results. Each identifies a weakening that circumvents an impossibility theorem.

Weakening IPP to EIPP does expose a new, albeit thoroughly unreasonable, "aggregation" solution where $\Pr_0$ is always uniform, regardless of the inputs. A more interesting possibility result arises by weakening IPP to require only the preservation of Markov independencies. Section 3 demonstrates that the LogOP does in fact maintain all agreed upon Markov independencies. This suggests that, if the preservation of independence structure is important—and many authors argue that it is

(Laddaga 1977, Raiffa 1968)—then the LogOP may be the most viable option. Markov independencies play an important role in the theory of graphical models and are precisely the type representable in MNs and decomposable BNs. We describe procedures for constructing MN and BN structures consistent with the LogOP consensus. We also delineate an algorithm for computing all of the conditional probability tables in a LogOP-consensus BN. This structured representation is potentially exponentially smaller than the standard representation.

The preservation of Markov independencies has a direct corollary in the investigation of structured securities markets in §4. For a certain class of agents, true Markov independencies are always observable as risk-neutral independencies. Thus, if all agents are of this type, all beliefs agree with the independencies encoded in the market structure, and this structure is decomposable, then the market is operationally complete.

Acknowledgments

The authors thank Didier Dubois, C. Lee Giles, Robin Hanson, Eric Horvitz, Jeffrey MacKie-Mason, Robert Nau, Charles Plott, Stephen Pollock, Ross Shachter, James Smith, Mike West, Fredrik Ygge, the members of the Decision Machine Research Group at the University of Michigan, the anonymous reviewers of this manuscript, and the anonymous reviewers of the two conference versions of this work (Pennock and Wellman 1999, 2000). They give special thanks to Pedrito Uriah Maynard-Zhang for extremely in-depth comments, corrections, analysis, and insights. Parts of this investigation were conducted while the first author was at Microsoft Research and at NEC Laboratories America.

References

Arrow, K. J. 1964. The role of securities in the optimal allocation of risk-bearing. *Rev. Econom. Stud.* **31**(2) 91–96.

Benediktsson, J. A., P. H. Swain. 1992. Consensus theoretic classification methods. *IEEE Trans. Systems, Man, Cybernetics* **22**(4) 688–704.

Bonduelle, Y. 1987. Aggregating expert opinions by resolving sources of disagreement. Unpublished doctoral thesis, Stanford University, Stanford, CA.

Cesa-Bianchi, N., Y. Freund, D. Haussler, D. P. Helmbold, R. E. Schapire, M. K. Warmuth. 1997. How to use expert advice. *J. ACM* **44**(3) 427–485.

Chyu, C. C. 1991. Decomposable probabilistic influence diagrams. *Probab. Engrg. Inform. Sci.* **5** 229–243.

Clemen, R. T., R. L. Winkler. 1993. Aggregation of point estimates: A flexible modeling approach. *Management Sci.* **39**(4) 501–515.

Cooke, R. M. 1991. *Experts in Uncertainty: Opinion and Subjective Probability in Science*. Environmental Ethics and Science Policy Series, Oxford University Press, New York.

Cooper, G. 1990. The computational complexity of probabilistic inference using Bayes belief networks. *Artificial Intelligence* **42** 393–405.

Dalkey, N. C. 1975. Toward a theory of group estimation. H. A. Linstone, M. Turoff, eds. *The Delphi Method: Techniques and Applications*. Addison-Wesley, Reading, MA, 236–261.

Darroch, J. N., S. L. Lauritzen, T. P. Speed. 1980. Markov fields and log-linear interaction models for contingency tables. *Ann. Statist.* **8**(3) 522–539.

de Finetti, B. 1974. *Theory of Probability: A Critical Introductory Treatment*, Vol. 1. Wiley, New York.

Degroot, M. H. 1974. Reaching a consensus. *J. Amer. Statist. Association* **69**(345) 118–121.

Degroot, M. H., J. Mortera. 1991. Optimal linear opinion pools. *Management Sci.* **37**(5) 546–558.

Faria, A. E., J. Q. Smith. 1996. Conditional external Bayesianity in decomposable influence diagrams. *Bayesian Statist.* **5** 551–560.

French, S. 1985. Group consensus probability distributions: A critical survey. *Bayesian Statist.* **2** 183–202.

Genest, C. 1984a. A characterization theorem for externally Bayesian groups. *Ann. Statist.* **12**(3) 1100–1105.

Genest, C. 1984b. A conflict between two axioms for combining subjective distributions. *J. Roy. Statist. Soc.* **46**(3) 403–405.

Genest, C. 1984c. Pooling operators with the marginalization property. *Canadian J. Statist.* **12**(2) 153–163.

Genest, C., C. G. Wagner. 1987. Further evidence against independence preservation in expert judgement synthesis. *Aequationes Mathematicae* **32**(1) 74–86.

Genest, C., J. V. Zidek. 1986. Combining probability distributions: A critique and an annotated bibliography. *Statist. Sci.* **1**(1) 114–148.

Genest, C., K. J. McConway, M. J. Schervish. 1986. Characterization of externally Bayesian pooling operators. *Ann. Statist.* **14**(2) 487–501.

Huang, C.-F., R. H. Litzenberger. 1988. *Foundations for Financial Economics*. Elsevier, Amsterdam, The Netherlands.

Jacobs, R. A. 1995. Methods for combining experts' probability assessments. *Neural Comput.* **7**(5) 867–888.

Jensen, F. V. 1996. *An Introduction to Bayesian Networks*. Springer, New York.

Kadane, J. B., R. L. Winkler. 1988. Separating probability elicitation from utilities. *J. Amer. Statist. Association* **83**(402) 357–363.

Kloks, T. 1994. *Treewidth: Computations and Approximations*. Springer-Verlag, Berlin, Germany.

Laddaga, R. 1977. Lehrer and the consensus proposal. *Synthese* **36** 473–477.

Lauritzen, S. L., D. J. Spiegelhalter. 1988. Local computations with probabilities on graphical structures and their application to expert systems. *J. Roy. Statist. Soc. Ser. B* **50** 157–224.

Lehrer, K., C. G. Wagner. 1983. Probability amalgamation and the independence issue: A reply to Laddaga. *Synthese* **55**(3) 339–346.

Levy, W. B., H. Delic. 1994. Maximum entropy aggregation of individual opinions. *IEEE Trans. Systems, Man, Cybernetics* **24**(4) 606–613.

Lindley, D. V. 1985. Reconciliation of discrete probability distributions. J. M. Bernando, M. H. Degroot, D. V. Lindley, A. F. M. Smith, eds. *Bayesian Statistics 2*. North-Holland, Amsterdam, The Netherlands, 375–390.

Lindley, D. V. 1988. The use of probability statements. C. A. Clarotti, D. V. Lindley, eds. *Accelerated Life Testing and Experts' Opinions in Reliability*. Elsevier, Amsterdam, The Netherlands, 25–57.

Matzkevich, I., B. Abramson. 1992. The topological fusion of Bayes nets. *8th Conf. Uncertainty Artificial Intelligence*, 152–158.

Matzkevich, I., B. Abramson. 1993. Some complexity considerations in the combination of belief networks. *9th Conf. Uncertainty Artificial Intelligence*, 152–158.

Morris, P. A. 1974. Decision analysis expert use. *Management Sci.* **20**(9) 1233–1241.

Morris, P. A. 1977. Combining expert judgments: A Bayesian approach. *Management Sci.* **23**(7) 679–693.

Myung, I. J., S. Ramamoorti, A. D. Bailey. 1996. Maximum entropy aggregation of expert predictions. *Management Sci.* **42**(10) 1420–1436.

Nau, R. F. 1995. Coherent decision analysis with inseparable probabilities and utilities. *J. Risk Uncertainty* **10**(1) 71–91.

Nau, R. F., K. F. McCardle. 1991. Arbitrage, rationality, and equilibrium. *Theory Decision* **31** 199–240.

Neapolitan, R. E. 1990. *Probabilistic Reasoning in Expert Systems: Theory and Algorithms*. Wiley, New York.

Ng, K.-C., B. Abramson. 1992. Consensus diagnosis—A simulation study. *IEEE Trans. Systems, Man, Cybernetics* **22**(5) 916–928.

Ng, K.-C., B. Abramson. 1994. Probabilistic multi-knowledge-base systems. *Appl. Intelligence* **4**(2) 219–236.

Pearl, J. 1988. *Probabilistic Reasoning in Intelligent Systems*. Morgan Kaufmann, San Francisco, CA.

Pennock, D. M., M. P. Wellman. 1997. Representing aggregate belief through the competitive equilibrium of a securities market. *13th Conf. Uncertainty Artificial Intelligence*, Providence, RI, 392–400.

Pennock, D. M., M. P. Wellman. 1999. Graphical representations of consensus belief. *15th Conf. Uncertainty Artificial Intelligence*, San Francisco, CA, 531–540.

Pennock, D. M., M. P. Wellman. 2000. Compact securities markets for Pareto optimal reallocation of risk. *16th Conf. Uncertainty Artificial Intelligence*, San Francisco, CA, 481–488.

Pennock, D. M., M. P. Wellman. 2001. A market framework for pooling opinions. Technical report 2001-081, NEC Research Institute, Princeton, NJ.

Raiffa, H. 1968. *Decision Analysis: Introductory Lectures on Choices under Uncertainty*. Addison-Wesley, Reading, MA.

Rosenblueth, E., M. Ordaz. 1992. Combination of expert opinions. *J. Scientific Indust. Res.* **51** 572–580.

Savage, L. J. 1954. *The Foundations of Statistics*. Wiley, New York.

Shachter, R. D. 1988. Probabilistic inference and influence diagrams. *Oper. Res.* **36**(4) 589–604.

Shachter, R. D., S. K. Andersen, K. L. Poh. 1991. Directed reduction algorithms and decomposable graphs. *Uncertainty in Artificial Intelligence*, Vol. 6. North Holland, Amsterdam, The Netherlands, 197–208.

Varian, H. R. 1987. The arbitrage principle in financial economics. *J. Econom. Perspectives* **1**(2) 55–72.

von Neumann, J., O. Morgenstern. 1953. *Theory of Games and Economic Behavior*. Princeton University Press, Princeton, NJ.

Wagner, C. 1984. Aggregating subjective probabilities: Some limitative theorems. *Notre Dame J. Formal Logic* **25**(3) 233–240.

West, M., J. Crosse. 1992. Modeling probabilistic agent opinion. *J. Roy. Statist. Soc. Ser. B. Methodological* **54**(1) 285–299.

Whittaker, J. 1990. *Graphical Models in Applied Multivariate Statistics*. Wiley, Chichester, England.

Winkler, R. L. 1968. The consensus of subjective probability distributions. *Management Sci.* **15**(2) B61–B75.

Winkler, R. L. 1981. Combining probability distributions from dependent sources. *Management Sci.* **27**(4) 479–488.

Winkler, R. L. 1986. Expert resolution. *Management Sci.* **32**(3) 298–303.

Wolfers, J., E. Zitzewitz. 2005. Interpreting prediction market prices as probabilities. Technical report, http://bpp.wharton.upenn.edu/jwolfers/Papers/InterpretingPredictionMarketPrices.pdf.

Xiang, Y. 1996. A probabilistic framework for cooperative multi-agent distributed interpretation and optimization of communication. *Artificial Intelligence* **87**(1–2) 295–342.

Yule, G. U. 1903. Notes on the theory of association of attributes in statistics. *Biometrika* **2** 121–134.