Perspectives on the Theory and Practice of Belief Functions

Glenn Shafer[1]

The theory of belief functions provides one way to use mathematical probability in subjective judgment. It is a generalization of the Bayesian theory of subjective probability. When we use the Bayesian theory to quantify judgments about a question, we must assign probabilities to the possible answers to that question. The theory of belief functions is more flexible; it allows us to derive degrees of belief for a question from probabilities for a related question. These degrees of belief may or may not have the mathematical properties of probabilities; how much they differ from probabilities will depend on how closely the two questions are related.

Examples of what we would now call belief-function reasoning can be found in the late seventeenth and early eighteenth centuries, well before Bayesian ideas were developed. In 1689, George Hooper gave rules for combining testimony that can be recognized as special cases of Dempster's rule for combining belief functions (Shafer 1986a). Similar rules were formulated by Jakob Bernoulli in his *Ars Conjectandi*, published posthumously in 1713, and by Johann-Heinrich Lambert in his *Neues Organon*, published in 1764 (Shafer 1978).

Examples of belief-function reasoning can also be found in more recent work, by authors who were unaware of the seventeenth and eighteenth century work. For example, Per Olof Ekelöf, a Swedish legal scholar, reinvented Hooper's rules in the early 1960s (Shafer 1986a, Gärdenfors et al. 1983).

In its present generality, however, the theory of belief functions is due to Arthur P. Dempster and myself. Dempster set out the basic ideas of the theory in a series of articles in the 1960s. I developed the theory further and established its terminology and notation in 1976, in *A Mathematical Theory of Evidence*. Because of the role of Dempster and myself, the theory is sometimes called the "Dempster-Shafer theory."

My 1976 monograph is still the most comprehensive source of information on belief functions. Unfortunately, this monograph says little about interpretation, implementation, or computation. It is concerned primarily with mathematical foundations.

During the past fifteen years, a great deal has been learned about the interpretation, implementation, and computation of belief functions, and fresh progress has been made on the mathematical foundations as well. This work is scattered in journals in a wide variety of fields, including statistics, psychology, philosophy, engineering, accounting, and artificial intelligence. No one has drawn the new work together in a comprehensive way.

This article is primarily a summary of my own current views. It pulls together some of the strands in the scattered literature, but it falls short of the comprehensive review that is needed. Though the bibliography is lengthy, it is not comprehensive. My failure to mention particular contributions should not be taken as an indication of disinterest or disagreement.

I begin with some opinions about the place of belief functions in the larger tool-drawer of probabilistic methods and about the place of this whole tool-drawer in artificial intelligence (Sections 1 and 2). These opinions are personal. I expect that few readers will agree with them fully. They may serve, nonetheless, to dispel some misunderstandings about the scope of the theory of belief functions.

[1]Ronald G. Harper Distinguished Professor of Business, School of Business, Summerfield Hall, University of Kansas, Lawrence, Kansas 66045.

After this general introduction, I introduce the basic ideas of belief functions, in a way that should be accessible to readers with no previous familiarity with the topic (Section 3), and I review the many competing mathematical formulations of the theory (Section 4).

I then turn to issues of interpretation and semantics. I explain that belief-function degrees of belief should not be interpreted as lower bounds on unknown probabilities (Section 5). Instead, they should be given a process-oriented semantics; they are the result of deliberately drawing a certain analogy (Section 6).

I then turn to issues closer to implementation. I discuss how the concept of independence is used in belief-function reasoning (Section 7) and how belief functions can and cannot be used in reasoning about frequency distributions (Section 8). I discuss how the computational complexity of belief functions can be dealt with (Section 9). And I discuss the extent to which belief-function reasoning can be automated (Section 10).

Finally, I discuss briefly a number of other topics in belief-function theory—generalizations, decision methods, consensus, infinite frames, weights of evidence, and other mathematical advances (Section 11).

1. *The Place of Belief Functions in Probability*

In my earliest work on belief functions (Shafer 1973, 1976a,b) I took a youthfully ambitious view of their possible role. I saw belief functions as a very general tool for subjective judgment. Almost every item of evidence could be represented, I suggested, by a belief function.

By the early 1980s, however, I was expressing a more sober view (see especially Shafer 1981). Then, as now, I saw belief functions as one among many tools for making subjective judgements using probability. Other such tools include Bayesian probability, Fisherian tests of significance, and Neyman-Pearson confidence intervals.

It was once customary to equate Bayesian methods with subjective probability and to identify Fisherian and Neyman-Pearson methods with a frequentist conception of probability. These were the terms of the debate that raged between Bayesians and frequentists in the 1960s. In recent years, however, the debate has cooled, and both sides have acknowledged that subjectivity is involved in all applications of probability to real problems.

In recent work (Shafer 1990), I have tried to give a foundation to this emerging consensus by advancing the thesis that the mathematical theory of probability is really the theory of an ideal picture in which belief, fair price, and knowledge of the long run are bound together. Probabilities in this ideal picture are long-run frequencies, but they are also degrees of belief, because they are known and nothing else that is relevant is known. This ideal picture seldom occurs in nature, but there are many ways of making it relevant to real problems. In some cases, the ideal picture serves as a standard of comparison—Fisherian tests of significance are often used to compare correlations in real problems to accidental correlations in the ideal picture. In other cases, we simulate the ideal picture (we obtain random numbers from tables or numerical algorithms) and then deliberately entangle this simulation in a real problem (we use the random numbers to randomize experiments or draw random samples). In other cases, we draw an analogy between the state of knowledge in the ideal picture and our evidence in a real problem.

Subjectivity is involved in all these applications of probability, not only because it is involved in the ideal picture, but also because the comparison, entanglement, or analogy that relates the ideal picture to the real problem requires subjective judgment.

Bayesian methods rely on one particular type of analogy that can be drawn to the ideal picture. In this analogy, we compare our evidence about the answer to a question of interest to knowledge of frequencies in the ideal picture. This analogy is strongest when extensive frequency data is available for our problem (in this case, it is customary to talk about "empirical Bayes"). When our evidence does not consist of frequencies, the analogy may or may not be convincing.

Belief functions arise when the Bayesian type of analogy is drawn between the ideal picture and one question, and then the resulting probabilities are examined for their relevance to a related question, which is of greater interest to us but for which we do not have a convincing direct Bayesian analogy. Here, as in the Bayesian case, there is no guarantee that the argument by analogy will be convincing. But if there are a number of questions that are related to the question of interest, then we may have a chance to find a successful analogy between the ideal picture and our evidence for at least one of these questions (Shafer and Tversky 1985).

Which of these many different ways of using probability is most important in practice? Fisherian and Neyman-Pearson applications are by far the most common and most important. Bayesian applications have been steadily increasing in recent decades. There are still few belief-function applications. I expect both the dominance of Fisherian and Neyman-Pearson applications and the steady increase in Bayesian applications to continue. I am also optimistic that belief-function applications will continue to increase as the ideas of the theory become more widely understood.

2. *The Role of Probability in Artificial Intelligence*

What role has probability played in AI, and what role should it play? Does probability provide a general theory for handling uncertainty in AI? In this section, I review the history of probability in AI, and I argue that we should be modest and realistic about the future prospects for probability in AI.

Historically, probability has not played a strong role in the theory of AI. Beginning with the inception of AI as a field in the 1950s, the theory of AI has been dominated by logic. On the other hand, there has always been a role for probability in AI applications. Standard probabilistic and statistical methods have long been used in AI work in pattern recognition and learning (Nilsson 1965, Duda and Hart 1973), and recently probabilistic ideas and techniques have been used in expert systems (Spiegelhalter 1986).

The drive to make logic a foundation for AI has faltered in recent years. At the beginning of the 1980s, logicians in AI were still optimistic about extending logic to deal with commonsense or uncertain reasoning, and the term "non-monotonic logic" was coined as a general name for such extensions. As the 1980s have drawn to a close, however, the promise of non-monotonic logic has faded, and some of its most prominent original proponents have questioned the fundamental idea of using logic for practical reasoning (see especially McDermott 1987).

Does the faltering role of logic mean that it should be replaced by probability as a general foundation for AI? A number of probabilists, including Cheeseman (1988) and Pearl (1988) believe that it should. These probabilists point out that uncertainty is involved in the vast majority of AI problems, and they repeat general arguments by Savage (1954), Cox (1961) and others to the effect that probability is always appropriate for dealing with uncertainty. In my opinion, this line of thinking is misguided. Common sense tells us that there are many situations involving uncertainty in which the theory of probability is not useful. The arguments advanced by Savage and Cox should be seen as clarifications of subjective probability, not as demonstrations that it is always appropriate (Shafer 1986b, 1988). The cause of probability in AI will be better served, in the long run, by a realistic and modest assessment of its limitations and potential.

The mere fact that there is uncertainty in a problem does not mean that the theory of probability is useful in the problem. As I pointed out in the preceding section, the theory of probability is really about a special ideal picture, not about all situations involving uncertainty. Bringing the theory of probability to bear on a particular problem means relating it in some way to the ideal picture. This requires imagination, and success is not guaranteed. Even if there is a way to bring the ideal picture to bear on a problem, doing so may not be worth the trouble. In many AI applications, it is more sensible to generate one plausible guess or plan than to weigh

probabilities. In others, it is more efficient to give more reliable sources of information priority in a hierarchical system of control than to weigh the probabilities of each report from these sources (Brooks 1986, Cohen 1987).

Since the ideal picture of probability involves frequencies, probability is most easily applied when relevant frequencies are available. Artificial intelligence is no exception to the rule that frequentist applications far outnumber subjective applications. The applications to pattern recognition and learning are predominantly frequentist, as are the most successful current applications of probability to expert systems. Pearl (1988) has emphasized the potential of subjective judgment in causal models, arguing that causal models are persuasive even when frequencies are replaced by subjective guesses. But this persuasiveness is not always matched by reliability. Numerous studies have shown that simple actuarial rules of thumb, when they are available, are more reliable than subjective clinical judgment (Dawes, Faust, and Meehl 1989).

Finding ways of using probability—or alternatives to probability—when directly relevant frequencies are not available should be seen as a difficult task. It is not a task that can be disposed of by general normative arguments such as those of Savage. It is not a task that can be disposed of by finding the right alternative calculus, such as the theory of belief functions or Zadeh's theory of possibility measures (Zadeh 1978). It is a task that must be dealt with in the context of each application.

3. *The Basic Ideas of Belief Functions*

The theory of belief functions is based on two ideas: the idea of obtaining degrees of belief for one question from subjective probabilities for a related question, and Dempster's rule for combining such degrees of belief when they are based on independent items of evidence.

The simplest way to illustrate these ideas is to go back to the topic addressed by George Hooper in 1689, the reliability of testimony. The belief-function approach to testimony is to derive degrees of belief for statements made by witnesses from subjective probabilities for the reliability of these witnesses.

Suppose that Betty tells me a tree limb fell on my car. My subjective probability that Betty is reliable is 0.9; my subjective probability that she is unreliable is 0.1. Since they are probabilities, these numbers add to one. But Betty's statement, which must true if she is reliable, is not necessarily false if she is unreliable. So I say that her testimony alone justifies a 0.9 degree of belief that a limb fell on my car, but only a zero degree of belief (not a 0.1 degree of belief) that no limb fell on my car. This zero does not mean that I am sure that no limb fell on my car, as a zero probability would; it merely means that Betty's testimony gives me no reason to believe that no limb fell on my car. The 0.9 and the zero together constitute a belief function.

Dempster's rule is based on the standard idea of probabilistic independence, applied to the questions for which we have subjective probabilities. I can use the rule to combine evidence from two witnesses if I consider the first witness's reliability subjectively independent (before I take account of what the witnesses say) of the second's reliability. (This means that finding out whether one witness is reliable would not change my subjective probability for whether the other is reliable.) The rule uses this subjective independence to determine joint probabilities for the various possibilities as to which of the two are reliable.

After using independence to compute joint probabilities for who is reliable, I must check whether some possibilities are ruled out by what the witnesses say. (If Betty says a tree limb fell on my car, and Sally says nothing fell on my car, then they cannot both be reliable.) If so, I normalize the probabilities of the remaining possibilities so they add to one. This is an example of probabilistic conditioning, and it may destroy the initial independence. (After I notice that Betty and Sally have contradicted each other, their reliabilities are no longer subjectively independent for me. Now finding out that one is reliable would tell me that the other is not.) After the normalization, I determine what each possibility for the reliabilities implies about the

4

truth of what the witnesses said, and I use the normalized probabilities to get new degrees of belief.

The net effect of Dempster's rule is that concordant items of evidence reinforce each other, conflicting items of evidence erode each other, and a chain of reasoning is weaker than its weakest link. To illustrate this, consider two independent witnesses, say Betty and Sally. Suppose the reliabilities of Betty and Sally are p_1 and p_2, respectively; Betty's testimony gives us a degree of belief p_1 in what she says, and degree of belief zero in its denial, while Sally's testimony gives us a degree of belief p_2 in what she says, and degree of belief zero in its denial. Then we can derive the following formulas:

If Betty and Sally say exactly the same thing, our degree of belief in what they say will be $1-(1-p_1)(1-p_2)$.

If they make different but consistent assertions, our degree of belief in both assertions being true will be $p_1 p_2$.

If they make contradictory assertions, our degree of belief in Betty's assertion will be $\frac{p_1(1-p_2)}{1-p_1 p_2}$, and our degree of belief in Sally's assertion will be $\frac{p_2(1-p_1)}{1-p_1 p_2}$.

These formulas are derived in Chapter 4 of Shafer (1976a). They represent only the simplest examples of Dempster's rule. When we combine more complex belief functions, Dempster's rule becomes too complex to be represented informatively through simple formulas.

4. The Mathematical Formalism of Belief Functions

The basic ideas of belief functions can be formalized mathematically in a variety of ways. In my 1976 monograph, I defined belief functions axiomatically, and I defined Dempster's rule by a formula. Other approaches include multivalued mappings, compatibility relations, random subsets, and inner measures. This section is concerned with these alternative approaches and with the relations among them.

The choice among these alternatives should be seen as a matter of convenience in mathematical exposition and investigation. Since the alternatives are mathematically equivalent, it makes no fundamental difference which we take as the starting point in such mathematical work.

In applications, the starting point does make a great deal of difference. But for applications, we need more than a mathematical definition or a set of axioms as a starting point. We need a metaphor that can serve as a guide in relating a practical problem to the theory and as a guide in assessing numbers to represent the strength of evidence in the practical problem. I will consider the problem of providing such metaphors in Section 6. The alternatives I discuss in this section are not such metaphors; they are merely mathematical formulations.

In the simple example of Betty's testimony, we started with two questions:

Q_1: Is Betty reliable?
Q_2: Did a tree limb fall on my car?

We had probabilities for Q_1, and we derived degrees of belief for Q_2. This process required no formal notation, because both Q_1 and Q_2 had only two possible answers: yes and no. In more complex examples, we will have questions Q_1 and Q_2 with many possible answers. To talk about such examples in general, we need a notation for each question's set of possible answers, a notation for the probabilities for Q_1 and the degrees of belief for Q_2, and a way of representing the constraints that an answer to Q_1 may put on the answer to Q_2.

We assume that each question comes with an exhaustive list of mutually exclusive answers. We know that exactly one of these answers is correct, but we do not know which one. We call

such a set of answers a *frame*. Let S be the frame for Q_1, the question for which we have probabilities, and let T be the frame for Q_2, the question of interest.

Let us write P(s) for the probability of the element s of S. Given these probabilities, and given a subset A of T, we want to derive a degree of belief Bel(A)—our degree of belief that A contains the correct answer to Q_2.

An answer s to Q_1 may rule out a whole set of answers to Q_2. If A is a set of answers to Q_2, and s rules out all the answers in A's complement, T-A, then s tells us that the answer to Q_2 is somewhere in A. Thus the probability P(s) should contribute to our belief in A. Our total degree of belief in A, Bel(A), should be the total probability for all s that rule out all answers in T-A. How shall we put this into symbols?

Here is where paths diverge. Multivalued mappings, compatibility relations, random subsets, and inner measures all provide different ways of specifying mathematically the answers to Q_2 ruled out by an answer to Q_1 and hence different ways of explaining the relation between the probabilities on S and the degrees of belief on T.

Multivalued Mappings. Let us first consider multivalued mappings, which were used by Dempster in his early articles—e.g., Dempster (1967a).

Let us write $\Gamma(s)$ for the subset of T consisting of the answers to Q_2 that are *not* ruled out by s. In this notation, s tells us that the answer to Q_2 is somewhere in A whenever

$$\Gamma(s) \subseteq A.$$

The degree of belief Bel(A) will be the total probability for all s that satisfy this condition. In symbols:

$$Bel(A) = P\{s \mid \Gamma(s) \subseteq A\}. \qquad (1)$$

The mapping Γ is a *multivalued mapping* from S to T. Formally, a *belief function* is any function Bel given by (1) for some multivalued mapping Γ and some probability measure P.

Now consider two belief functions Bel_1 and Bel_2 on T, which we judge to be based on independent items of evidence. Each belief function will be based on its own probability space and its own multivalued mapping from that probability space to T. We may write S_1 and S_2 for the two probability spaces, P_1 and P_2 for the two probability measures, and Γ_1 and Γ_2 for the two multivalued mappings. Dempster's rule is a rule for combining Bel_1 and Bel_2 in order to obtain a belief function Bel on T that represents the pooling of two items of evidence. How do we describe Dempster's rule in terms of S_1, S_2, P_1, P_2, Γ_1, and Γ_2?

We can answer this question by using S_1, S_2, P_1, P_2, Γ_1, and Γ_2 to construct a probability space S, a probability measure P, and a multivalued mapping Γ from S to T. The belief function Bel given by combining Bel_1 and Bel_2 by Dempster's rule will be the belief function given by (1) using S, P, and Γ.

Independence of the two items of evidence means that we can make initial joint probability judgments about the two questions answered by S_1 and S_2 by forming the product measure $P_1 \times P_2$ on $S_1 \times S_2$. It also means that what an element s_1 of S_1 tells us about Q_2 does not affect what an element s_2 of S_2 tells us about Q_2; s_1 and s_2 together tell us only that the answer to Q_2 is in the intersection $\Gamma_1(s_1) \cap \Gamma_2(s_2)$. If this intersection is empty for some (s_1,s_2), then s_1 and s_2 are telling us contradictory things about Q_2, and one of them must be wrong. So we must condition the product measure $P_1 \times P_2$ on the set of (s_1,s_2) for which $\Gamma_1(s_1) \cap \Gamma_2(s_2)$ is not empty. We let S be the subset of $S_1 x S_2$ consisting of (s_1,s_2) for which $\Gamma_1(s_1) \cap \Gamma_2(s_2)$ is not empty, and we let P be the probability measure on S obtained by conditioning $P_1 x P_2$ on S. Finally, we let Γ be the multivalued mapping from S to T given by

$$\Gamma(s_1,s_2) = \Gamma_1(s_1) \cap \Gamma_2(s_2).$$

This completes the construction of S, P, and G, and hence the statement of Dempster's rule.

To summarize verbally: Dempster's rule says to form the product probability space, condition it by eliminating pairs that map to disjoint subsets of T, and then obtain a belief

function by mapping each remaining pair to the intersection of the subsets to which the two elements of the pair are mapped.

Compatibility Relations. The multivalued mapping Γ from S to T tells us, for each element s of S, which elements of T are possible answers to Q_2 if s is the correct answer to Q_1. It tells us, in other words, which s are compatible with which t. This information can also be represented by specifying the set C of all ordered pairs (s,t) such that s is compatible with t. This is a "relation"—a subset of the Cartesian product S×T. The relation C is related to the multivalued mapping Γ by

$$C = \{(s,t) | \, t \in \Gamma(s)\}$$

and

$$\Gamma(s) = \{t | \, (s,t) \in C\}.$$

In terms of C, (1) becomes

$$Bel(A) = P\{s | \, \{t | \, (s,t) \in C\} \subseteq A\},$$

or

$$Bel(A) = P\{s | \, \text{if } (s,t) \in C, \text{ then } t \epsilon A\}. \tag{2}$$

Shafer (1987b) calls C a *compatibility relation*, and he develops the mathematics of belief functions from (2). This is also the approach taken by Lowrance (1988) and by Shafer and Srivastava (1990).

In truth, the choice between multivalued mappings and compatibility relations is scarcely more than a choice of terms. Many treatments of mathematical set theory (e.g., Kelley 1955) define mappings as relations. Moreover, Dempster used the word "compatible" repeatedly in explaining the meaning of his multivalued mappings.

Random subsets. We have assigned a subset of T, Γ(s), to each s, and we have assigned a probability P(s) to each s. If we think of the probability as being attached to the subset instead of to s, then we have, in effect, defined a random subset of T.

The subset Γ(s) might be the same for different s. So in order to find the total probability that the random subset will be equal to B, we must add the probabilities of all the s for which Γ(s) is equal to B. The degree of belief Bel(A) is the total probability that the random subset is contained in A.

In this setting, Dempster's rule is a rule for combining two random subsets to obtain a third—or, more precisely, a rule for combining two probability distributions for random subsets to obtain a third. We assume that the random subsets are probabilistically independent, we intersect them, and then we condition the probability distribution for the intersection on its being non-empty.

This approach to the mathematics of belief functions was emphasized by Nguyen (1978), Goodman and Nguyen (1985), and Shafer, Shenoy, and Mellouli (1987). It is convenient for advanced mathematical exposition, because the idea of a random subset is well established among mathematical probabilists (Matheron 1975).

The Axiomatic Approach. Another approach is to characterize belief functions directly in terms of their mathematical properties. We simply list a set of axioms that a belief function Bel must satisfy. And we use a formula to define Dempster's rule for combining two belief functions Bel_1 and Bel_2. This was the approach of Shafer (1973, 1976a,b, 1979). It is related to earlier mathematical work by Choquet (1953).

In my 1976 monograph, I also gave a more transparent characterization for the case where T is finite. We assign a non-negative number m(B) to each subset B of T. Intuitively, m(B) is the probability that the random subset is equal to B. We require that $m(\varnothing)=0$, where $\varnothing$ is the empty set, and that the m(B) add to one. The function m is called the basic probability assignment. We define the function Bel by

$$Bel(A) = \Sigma \, \{m(B) \, | \, B \subseteq A\}. \tag{3}$$

As it turns out, the m(B) can then be recovered from the Bel(A) by the formula

$$m(B) = \Sigma \ \{(-1)^{|A|} Bel(A) \mid A \subseteq B\}. \tag{4}$$

The functions m and Bel are said to be Möbius transforms of each other (Rota 1964). Dempster's rule for combining two belief functions Bel_1 and Bel_2 can be defined by a relatively simple rule in terms of the corresponding basic probability assignments m_1 and m_2; we define m(B) by

$$m(B) = \frac{\Sigma\{m_1(B_1)m_2(B_2) \mid B_1 \cap B_2 = B\}^{\cdot}}{\Sigma\{m_1(B_1)m_2(B_2) \mid B_1 \cap B_2 \neq \varnothing\}}. \tag{5}$$

The belief function Bel resulting from the combination can then be obtained from m using (3).

Many other authors, including Gordon and Shortliffe (1984, 1985), have used basic probability assignments to explain the theory of belief functions.

Inner Probability. The idea of deriving minimal degrees of belief for some sets from probabilities for other sets has long been familiar in abstract probability theory in the context of "inner measures" (Halmos 1950) or "inner probabilities" (Neveu 1965). With attention to a few technicalities, we can relate belief functions to the idea of inner measure or inner probability.

It is easiest to explain this using the compatibility relation C that relates the frames S and T. The set C is itself a frame—it consists of the possible answers to the joint question formed by compounding the question answered by S and the question answered by T. A Bayesian approach to our problem would construct a probability measure on C. The belief-function approach stops short of completing this construction. We stop when we have constructed a probability measure P on S. We then extend P to an inner probability on C.

Let us review the definition of inner probability in the finite case. Recall that an algebra of subsets of a set X is a collection of subsets that includes X and the empty set $\varnothing$ and also includes $A \cap B$, $A \cup B$, $X - A$, and $X - B$ whenever it includes A and B. Given a probability measure Q defined only on an algebra $\Im$ of subsets of a finite set X, the inner probability of Q is the function Q_* defined by

$$Q_*(A) = \max \ \{Q(B) \mid B \in \Im \text{ and } B \subseteq A\} \tag{6}$$

for every subset A of X. Intuitively, $Q_*(A)$ is the degree to which the probabilities for the elements of $\Im$ force us to believe A.

Let $\Im$ denote the collection of all subsets of C of the form $C \cap (R \times T)$, where R is a subset of S. This is an algebra of subsets of C. Since the subset $C \cap (R \times T)$ of C has the same meaning (qua assertion about the answer to Q_1) as the subset R of S, it is natural to define a probability measure Q on $\Im$ by setting $Q(C \cap (R \times T)) = P(R)$. With this definition, (6) becomes

$$Q_*(A) = \max \ \{P(R) \mid R \subseteq S \text{ and } C \cap (R \times T) \subseteq A\} \tag{7}$$

for every subset A of C.

What belief should we give to a subset U of T? It is natural to answer by looking at the value of (7) for the subset of C that corresponds to U, namely $C \cap (S \times U)$. This is

$$Q_*(\ C \cap (S \times U)) \quad = \max \ \{P(R) \mid R \subseteq S \text{ and } C \cap (R \times T) \subseteq C \cap (S \times U)\}$$

$$= \max \ \{P(R) \mid R \subseteq S \text{ and if } s \in R \text{ and } (s,t) \in C, \text{ then } t \in U\}$$

$$= P\{s \mid \text{ if } (s,t) \in C, \text{ then } t \in U\},$$

which is the same as formula (2) for Bel(U).

Thus a belief function is simply the inner measure of a probability measure—or, more precisely, the restriction to a subalgebra of an inner measure obtained from a probability measure on a possibly different subalgebra.

This connection between inner measures and belief functions must have been known for some time to many students of belief functions. To the best of my knowledge, however, it has appeared in the literature only in the past few years. The only references I know are Ruspini (1987) and Fagin and Halpern (1989a).

5. Belief-Function Degrees of Belief are not Lower Bounds

In this section, I will review the point, well established in the literature on belief functions, that belief-function degrees of belief should not be interpreted as bounds on unknown true probabilities. Such an interpretation seems plausible when we consider only a single belief function, but it breaks down when we consider belief functions that represent different and possibly conflicting items of evidence. Most importantly, a probability-bound interpretation is incompatible with Dempster's rule for combining belief functions. If we make up numbers by thinking of them as lower bounds on true probabilities, and we then combine these numbers by Dempster's rule, we are likely to obtain erroneous and misleading results.

In order to see how the degrees of belief given by a belief function might be thought of as lower bounds on probabilities, consider again my 0.9 belief that a limb fell on my car, and my zero belief that no limb fell on my car. These degrees of belief were derived from my 0.9 and 0.1 subjective probabilities for Betty being reliable or unreliable. Suppose these subjective probabilities were based on my knowledge of the frequency with which witnesses like Betty are reliable. Then I might think that the 10% of witnesses like Betty who are not reliable make true statements a definite (though unknown) proportion of the time and false statements the rest of the time. Were this the case, I could think in terms of a large population of statements made by witnesses like Betty. In this population, 90% of the statements would be true statements by reliable witnesses, x% would be true statements by unreliable witnesses, and (10-x)% would be false statements by unreliable witnesses, where x is an unknown number between 0 and 10. The total chance of getting a true statement from this population would be (90+x)%, and the total chance of getting a false statement would be (10-x)%. My degrees of belief of 0.9 and zero are lower bounds on these chances; since x is anything between 0 and 10, 0.9 is the lower bound for (90+x)%, and zero is the lower bound for (10-x)%.

As this example suggests, a *single* belief function is always a consistent system of probability bounds. For any belief function Bel over any finite frame T, there will exist a class of probability distributions $\wp$ such that

$$\text{Bel}(A) = \min_{P \in \wp} P(A) \tag{8}$$

for every subset A of T. (There are many ways of seeing that this is true. One way is to recall that Bel(A) is the sum of m(B) for all B contained in A. Consider the different probability distributions obtained by distributing the mass m(B), for each B, among the elements of B. If B is contained in A, then all the mass has to fall in A, but if B is not contained in A, then it is possible to distribute it all outside of A. Hence minimum probability that one of these distributions can give A is Bel(A).)

However, the degrees of belief given by belief functions should *not* be interpreted as lower bounds on some unknown true probability. Belief functions are not, in general, concerned with a well-defined reference population, and with learning about the frequencies in this population. And differences between belief functions do not, in general, reflect disagreements about unknown true probabilities. When Betty says a limb fell on my car, and Sally says nothing fell on my car, they are disagreeing about whether something fell on my car, not about the true probability of something having fallen on my car.

Were we to insist on a probability-bound interpretation of belief functions, then we would only be interested in groups of belief functions whose degrees of belief, when interpreted as probability bounds, can be satisfied simultaneously. But when belief functions are given their proper interpretation it is of no particular significance whether there exist probabilities that simultaneously satisfy the bounds defined by a whole group of belief functions. Consider two cases that might arise when we use belief functions to represent contradictory evidence from Betty and Sally:

> Case 1. Before hearing their testimony, we think highly of the reliability of both Betty and Sally. We represent Betty's evidence by a belief function that

9

gives a 95% degree of belief to a limb having fallen on my car, and we represent Sally's evidence by a belief function that gives a 95% degree of belief to nothing having fallen on my car. In this case, the two belief functions are contradictory as probability bounds; if the true probability of a limb having fallen on my car is greater than 95%, then the true probability of nothing having fallen on my car cannot also be greater than 95%.

Case 2. Before hearing their testimony, we think that both Betty and Sally are fairly unreliable. So in both belief functions, we assign a 35% degree of belief rather than a 95% degree of belief. In this case, the two belief functions define consistent probability bounds; the true probability of a limb having fallen on my car and of nothing having fallen on my car can both be greater than 35%.

From the belief-function point of view, there is no conceptual difference between these two cases. In both cases, we can combine the two belief functions by Dempster's rule. In both cases, there is conflict in the evidence being combined, and normalization is required.

It can be shown that if no normalization is required in the combination of a group of belief functions by Dempster's rule, then there do exist consistent probabilities that simultaneously bound all the belief functions being combined as well as the belief function that results from the combination. This may also happen when normalization is required, as in case 1 above, but we cannot count on this. In general, a probability-bound interpretation of belief functions is inconsistent with normalization (Zadeh 1986).

Probability bounds do provide another way to use the ideal picture of probability in subjective judgment. I have called this the *lower-probability* approach in order to distinguish it from the belief-function approach (Shafer 1981). The lower-probability approach has been elaborated by Smith (1961), Good (1962), Suppes (1974), Levi (1980), Kofler et al. (1984), Nilsson (1986), Snow (1986) and others. It does not, in general, use Dempster's rule.

When we use the lower-probability approach in a practical problem, we are drawing an analogy between actual evidence and knowledge of bounds on unknown true probabilities for the question of interest. Like the belief-function approach, the lower-probability approach is sometimes but not always appropriate and useful. To claim it is always appropriate is to fall into the trap of assuming that unknown true probabilities always exist. In general, they exist only if a population and sampling scheme are well defined. An unknown true probability for the truth of Betty's statement, for example, exists only if a the population of true and false statements of witnesses like Betty is well-defined.

In some cases, a lower-probability analysis will come very close to a Bayesian analysis. Indeed, if the bounds we consider are fairly tight—fairly close to defining a single probability distribution—then they may correspond to bounds we would consider to see how slight changes in our subjective probability distribution would affect the outcome of a Bayesian analysis. Some authors call this "sensitivity analysis" (Fishburn 1965, Isaacs 1963); others call it "robust Bayesian analysis" (Kadane 1984).

One contrast between the belief-function and lower-probability approaches is in how they combine a given belief function Bel on a frame T with evidence that establishes that the answer is in a subset A of T. The belief-function approach is to combine Bel by Dempster's rule with a new belief function that gives degree of belief one to A and degree of belief zero to every proper subset of A. This generalizes Bayesian conditioning; I have called it *Dempster's rule of conditioning*. The lower-probability approach is to consider all the probability distributions bounded from below by the belief function, condition them all on A, and then take lower bounds over the resulting conditional probabilities. The lower-probability approach produces weaker degrees of belief than the belief-function approach (Shafer 1981, Kyburg 1987).

Though applying (8) to an arbitrary class of probability distributions p does not always produce a belief function, it does produce a belief function surprisingly often (Wasserman 1990). Moreover, the lower-probability approach to conditioning a belief function produces another

belief-function (Fagin and Halpern 1989b). These results show that belief functions are of greater interest to the lower-probability approach than one might have expected, but they do reconcile the two approaches from the point of view of application and interpretation.

The point that belief functions are not lower bounds on probability measures is well accepted in the literature on belief functions. It is true that Dempster used the idea of bounds on unknown probabilities as a didactic tool in several of his articles on belief functions (which he then called "upper and lower probabilities") in the 1960s. The idea of probability bounds was never basic to Dempster's work, however; his work differed in this respect from the earlier work of Smith (1961) and Good (1962). In my 1976 monograph, where I introduced both the name "belief function" and the notation now associated with belief functions, I explicitly disavowed the interpretation of belief-function degrees of belief as lower bounds over classes of probabilities (p. ix). In later articles, I have amplified, emphasized, and repeated this disavowal (Shafer 1981, 1987b, Shafer and Srivastava 1990). Dempster and other proponents of belief functions have seconded the disavowal (Dempster 1982, Ruspini 1987, Smets 1988).

6. The Semantics of Belief Functions

As I noted in Section 4, the use of belief functions in practical problems requires metaphors that can guide us in relating practical problems to the theory and in assessing the strength of evidence numerically. This is the practical meaning of calls for a "semantics" for belief functions. In this section, I review the metaphors that I have suggested in the past, as well as some of the suggestions others have made.

In my mind, the simplest and most effective metaphor for belief functions is the metaphor of the witness who may or may not be reliable. In many cases, the example of such a witness can serve as a standard of comparison for the strength of evidence. We can assess given evidence by saying that it is comparable in strength to the evidence of a witness who has a certain chance of being reliable.

A witness testifying to a specific proposition leads to a relatively simple belief function— one that gives a specified degree of belief to that proposition and its consequences, and zero degree of belief to all other propositions. Arbitrarily complex belief functions can be built up by combining such simple belief functions (Shafer 1976a, p. 200), but in some cases we may want to produce complex belief functions more directly, in order to represent evidence that conveys a complex or mixed message but cannot be broken down into independent components. This requires more complex metaphors.

In some cases, we can obtain the more complex metaphor we require simply by extending the metaphor of the witness who may not be reliable. Shafer and Tversky (1985), for example, give a metaphor for consonant belief functions (belief functions such that $m(A)=0$ except for a nested family of subsets A) by imagining a witness whose statements may be meaningful to different degrees of exactitude, with different probabilities.

In Shafer (1981), I suggested a more abstract but very general metaphor—the metaphor of a randomly coded message. In this metaphor, we have probabilities for several codes that might be used to encode a message. We do not yet know what the message says, but we know it is true. We have this message in hand in its coded form, and we will try to decode it using each code, but the probabilities are judgments we make before this decoding. When we do decode using the different codes, we sometimes get nonsense, and we sometimes get a comprehensible statement. It seems sensible, in this situation, to condition our probabilities for the codes by eliminating the ones with which we get nonsense. The conditioned probability for each remaining code can then be associated with the statement we get by decoding using that code. These statements may be related in various ways; some may be inconsistent with each other, and some may be stronger than others. Thus we obtain the complexity of an arbitrary belief function.

11

In this metaphor, the independence of two belief functions means that two different people independently choose codes with which to send two possibly different (though both true) messages. Our uncertainties about the codes in the two cases remain independent unless possible codes imply contradictory messages. If s_1 is a possible code for the first person, and s_2 is a possible code for the second person, and the first message as decoded by s_1 contradicts the second message as decoded by s_2, then it cannot be true that these were the two codes used. We eliminate such pairs of codes and normalize the probabilities of the remaining possible pairs. The probability of each pair is then associated with the conjunction of the two implied messages. This is Dempster's rule.

Both the metaphor of the witness and the more general metaphor of the randomly-coded message can be presented in a way that forestalls the interpretation of belief-function degrees of belief in terms of bounds on probabilities. There is no probability model for the choice of the true message sent. The probabilities are only for the choice of codes. We might visualize these probabilities in terms of a repetition of the choice of codes, but since the true message can vary arbitrarily over this population of repetitions, the idea of this population does not lead to the idea of a true unknown probability for the true message or for the true answer to the question of interest.

An insistence on imposing an overall Bayesian model on both the truth of the message and the choice of a code will lead, of course, to the conclusion that Dempster's rule is wrong. This has been shown in detail by Williams (1982) and Good (1982); see also Shafer (1986a). Related attempts to relate Dempster's rule to overall Bayesian models include Freeling and Sahlin (1982) and Baron (1987).

Laskey (1989) has given some examples that further illustrate the difference between belief-function thinking and a Bayesian approach to the randomly-coded message metaphor. In Laskey's examples, strikingly different randomly coded messages produce the same belief function. These randomly coded messages would produce quite different Bayesian probabilities if combined with Bayesian priors about the true message.

What other possibilities are there for a semantics for belief functions? Krantz (1982) has explored the possibility of justifying the rules for belief functions, including Dempster's rule, using measurement-type axioms. This offers at least a partial non-probabilistic semantics for belief functions, since the axioms are concerned with comparisons of evidence and hence provide guidance in the assessment of evidential strength. I have been unable to convince myself, however, that this approach provides an understanding of the combination of evidence. In particular, it is difficult to justify Dempster's rule fully without a probabilistic basis for the concept of independence.

Pearl (1988) has developed a metaphor involving random switches to provide a semantics for belief functions. Again, however, I have not been able to convince myself that this metaphor provides a full basis for Dempster's rule. It does not seem to provide a basis for normalization.

In this section, I have used the word "semantics" in a process-oriented way. Semantics for a mathematical theory of evidence is guidance in using the theory to make quantitative judgments of the strength of evidence. The meaning of the resulting judgments is not independent of this process of judgment.

Classical logic and the classical frequentist and Bayesian theories of probability all have a stronger conception of semantics. In all three cases, we can say what statements of the theory mean without reference to how they are obtained. In logic, the meaning of statements can be explained in terms of possible models or worlds. In frequentist probability, the meaning of statements can be explained in terms of actual frequencies. In Bayesian probability, the meaning of statements can be explained in terms of a person's willingness to bet. I do not believe, however, that such process-independent semantics is a reasonable goal for AI. I agree with Winograd (1980) that AI must use statements that have no meaning "in any semantic system that fails to deal explicitly with the reasoning process." One fundamental aspect of the subjectivity of

judgments under uncertainty is the fact that these judgments depend on the process by which they are made as well as on the objective nature of the evidence.

7. *Sorting Evidence into Independent Items*

Dempster's rule should be used to combine belief functions that represent independent items of evidence. But when are items of evidence independent? How can we tell? These are probably the questions asked most frequently about belief functions.

The independence required by Dempster's rule is simply probabilistic independence, applied to the questions for which we have probabilities, rather than directly to the question of interest. In the metaphor of the randomly coded messages, this means that the codes are selected independently. In the more specialized metaphor of independent witnesses, it means that the witnesses (or at least their current properties as witnesses) are selected independently from well-defined populations.

Whether two items of evidence are independent in a real problem is a subjective judgment, in the belief-function as in the Bayesian approach. There is no objective test.

In practice, our task is to sort out the uncertainties in our evidence. When items of evidence are not subjectively independent, we can generally identify what uncertainties they have in common, thus arriving at a larger collection of items of evidence that are subjectively independent. Typically, this maneuver has a cost—it forces us to refine, or make more detailed, the frame over which our belief functions are defined.

We can illustrate this by adapting an example from Pearl (1988). Suppose my neighbor Mr. Watson calls me at my office to say he has heard my burglar alarm. In order to assess this testimony in belief-function terms, I assess probabilities for the frame
$$S_1 = \{\text{Watson is reliable, Watson is not reliable}\}.$$
Here Watson being reliable means he is honest and he can tell whether he is hearing my burglar alarm. I can use these probabilities to get degrees of belief for the frame
$$T = \{\text{My alarm sounded, My alarm did not sound}\}.$$
Putting a probability of 90%, say, on Watson being reliable, I get a 90% degree of belief that my burglar alarm sounded, and a 0% degree of belief that my burglar alarm did not sound.

I now call another neighbor, Mrs. Gibbons, who verifies that my alarm sounded. I can assess her testimony in the same way, by assessing probabilities for the frame
$$S_2 = \{\text{Gibbons is reliable, Gibbons is not reliable}\}.$$
Suppose I also put a probability of 95% on Gibbons being reliable, so that I again obtain a 95% degree of belief that my burglar alarm sounded, and a 0% degree of belief that it did not sound.

Were I to combine these two belief functions by Dempster's rule, I would obtain an overall degree of belief of 99.5% that my burglar alarm sounded. This is inappropriate, however, for the two items of evidence involve a common uncertainty—whether there might have been some other noise similar to my burglar alarm.

In order to deal with this problem, I must pull my skepticism about the possibility of a similar noise out of my assessment of Watson's and Gibbons' reliability, and identify my grounds for this skepticism as a separate item of evidence. So I now have three items of evidence—my evidence for Watson's honesty (I say honesty now instead of reliability, since I am not including here the judgment that there are no other potential noises in the neighborhood that Watson might confuse with my burglar alarm), my evidence for Gibbons' honesty, and my evidence that there are no potential noises in the neighborhood that sound like my burglar alarm.

These three items of evidence are now independent, but their combination involves more than the frame T. In its place, we need the frame $U = \{u_1, u_2, u_3\}$, where
$$u_1 = \text{My alarm sounded,}$$
$$u_2 = \text{There was a similar noise,}$$
$$u_3 = \text{There was no noise.}$$

(Let us exclude, for simplicity of exposition, the possibility that there were two noises, my alarm and also a similar noise.) My first two items of evidence (my evidence for Watson's and Gibbons' honesty) both provide a high degree of belief in $\{u_1,u_2\}$, while the third item (my evidence against the existence of other noise sources) provides a high degree of belief in $\{u_1,u_3\}$. Combining the three by Dempster's rule produces a high degree of belief in $\{u_1\}$.

A Bayesian approach to this problem would be somewhat different, but it too would involve refining the frame T to U or something similar. In the Bayesian case, we would ask whether the events "Watson says he heard a burglar alarm" and "Gibbons says she heard a burglar alarm" are subjectively independent. They are not unconditionally independent, but they are independent conditional on a specification of what noise actually occurred. I can exploit this conditional independence in assessing my subjective probabilities, but in order to do so, I must bring the possibility of other noises into the frame.

In the belief-function approach, one talks not about conditional independence of propositions, but rather about the overlapping and interaction of evidence. For further explanation and more examples, see Shafer (1976a, Chapter 8), Shafer (1981), Shafer (1984), Shafer (1987b) and Srivastava, Shenoy, and Shafer (1989).

8. *Statistical Inference and Frequency Reasoning*

Statistical inference is inference about frequencies from sample data. There are belief-function methods for statistical inference, as well as Bayesian and classical frequentist methods. In this section, I discuss the relevance of statistical inference, and of frequency reasoning more generally, to the general problem of subjective judgment. We can sometimes draw analogies between non-sample evidence and imaginary data and hence use Bayesian or belief-function methods of statistical inference in non-statistical problems. In order to do this successfully, however, we must distinguish clearly between frequencies and degrees of belief.

The problem of statistical inference has two distinctive features—two features that distinguish it from the general problem of subjective judgment. First, the population that defines the frequencies is well defined. Second, we have substantial sample data from that population. There are many reasons for wanting to know about frequencies in a particular population. Sometimes we want to know about a population for its own sake. Sometimes knowledge about frequencies in a population may serve as evidence about causal mechanisms that manifest themselves in the population. Sometimes knowledge about frequencies in a population can help us make subjective judgments about a question that can be regarded as having its answer drawn randomly from the population. In the last two cases, statistical inference is only part of the undertaking, but it is an important part.

The enduring philosophical conundrum of statistical inference is the fact that the sample data does not seem to provide a complete basis for making inferences about frequencies or about parameters in statistical models (models for frequencies). We must usually supplement this data with subjective judgment or with arbitrary choices. In the Bayesian case, we must supplement it with subjective probabilities based on evidence other than the sample data. In both the classical frequentist and belief-function cases, there are usually several methods of statistical inference for a given problem, and we must make an arbitrary choice among them.

Classical frequentism is largely concerned with methods of estimating frequencies (or parameters that determine frequencies) from sample data and with estimating the average error in the estimates. There are almost always competing estimators in a given problem, and the choice among these methods is not always clear-cut (Efron 1978). Bayesian methods give more complete and definite answers; they produce probability distributions for the frequencies or parameters. These probability distributions depend, however, on prior subjective opinions as well as on the sample data (Savage 1961, Lindley 1972).

14

Dempster, in his original work on belief functions (Dempster 1966, 1967a,b, 1968a,b, 1969, 1972), was motivated by the desire to obtain probability judgments based only on sample data, without dependence on prior subjective opinion. His work, together with later work on belief-function statistical inference (Krantz and Miyamoto 1983, Shafer 1976a,b, 1982a,b, Tritchler and Lockwood 1990, Walley 1987, Wasserman 1989, Weisberg 1969, and West 1971) has shown that this is possible. But this work has produced a variety of belief-function methods, not a single prescribed method. In Shafer (1982b), I argued that we should choose among these methods by considering the nature of the evidence for the statistical model. But the evidence for the statistical model is often too nebulous for this approach to be helpful.

Most problems of subjective judgment are not problems of statistical inference, because there is neither sample data nor a well-defined population. In Section 1, I pointed out that the Bayesian approach, in general, draws an analogy between our actual evidence and knowledge of frequencies in a population. But usually this is only an analogy. The population in question is purely imaginary. Moreover, the analogy is with knowledge of frequencies—not with knowledge of sample data.

In some cases, however, we can approach a problem of subjective judgment by drawing an analogy between certain items of evidence and imaginary sample data from a real or imaginary population. This will allow us to use Bayesian or belief-function statistical methods even though we do not have real sample data. Suppose, for example, that we are interested in a bird named Tweety. We want to know whether Tweety flies and whether Tweety is a penguin. We decide to make judgments about this by defining a population of birds in some way and thinking of Tweety as randomly selected from this population. We have some opinions, based on fragmentary information, hunches, partial memories, etc., about the proportion of birds in this population that fly and the proportion that are penguins. If we can assess the strength of this evidence by saying that it is equivalent to certain sample data, then we can express this strength in terms of likelihoods, and then we can combine these likelihoods with other evidence in either a Bayesian or a belief-function framework (Krantz 1990). The result would be probabilities about the frequencies in the population, and derivatively, probabilities about Tweety.

It is important, if we follow this strategy of using statistical methods for non-statistical evidence, to distinguish clearly between frequencies and degrees of belief. The frequencies, even if they are imaginary because the population is imaginary, should not be thought of as degrees of belief, because they are not taken as known. We are evaluating evidence for what these frequencies are.

I emphasize the distinction between frequencies and degrees of belief because it tends to disappear when we use the basic Bayesian analogy, without statistical methods. The basic Bayesian analogy is between our actual evidence and an ideal picture in which frequencies are known and hence equal to our degrees of belief. It is natural, when making this analogy, to say that the numbers we produce are both degrees of belief and guesses at frequencies in an imaginary population. This basic Bayesian analogy will be needed even if we are using Bayesian statistical methods, for it is needed in order to assess the prior subjective probabilities. Thus in a Bayesian statistical analysis for a non-statistical problem, we will have some "frequencies" that are degrees of belief, and other "frequencies" that are unknown.

In the case of belief functions, a careless equating of frequency with degree of belief is especially dangerous. If our strategy of subjective judgment involves applying belief-function methods of statistical inference to an imaginary population, we must be clear whether a given number is a guess at a frequency or a degree of belief. Combining degrees of belief by Dempster's rule may be appropriate. Combining different guesses about frequencies by Dempster's rule certainly is not.

Consider, for example, our guesses about the proportion of birds in Tweety's population that fly and the proportion that are penguins. These guesses should not be represented as belief functions over a set of statements about Tweety and then combined by Dempster's rule.

15

We can explain this in terms of our discussion of dependence in the preceding section. Both guesses bear on Tweety only through their accuracy as guesses about the population. This means that they have in common the uncertainty involved in choosing Tweety at random from the population. Depending on how we obtained the guesses, they may also have other uncertainties in common. Like every problem of dependence, we can deal with this problem within the belief-function approach by sorting out the uncertainties and properly refining our frame. In this case, we must bring the possible values for the population frequencies into the frame. We can then formalize the connection between these frequencies and Tweety as one of our items of evidence. We must also identify our sources of evidence about the frequencies, sort out their uncertainties, and use them to assess belief functions about what these frequencies are.

Judea Pearl, in his article in this volume, gives some examples of misleading results that can arise from representing conditional probabilities as belief functions and then combining these conditional probabilities by Dempster's rule. The error involved in these examples is the same as the error in the Tweety example; we are using Dempster's rule to combine fragments of information about frequencies in a given population. A legitimate belief-function treatment must deal explicitly with the unknown overall frequency distribution and with the evidence about it.

I do not believe, however, that applying belief-function statistical methods to non-statistical problems is the most promising direction for the use of belief functions. It is easy to talk about "the population of birds like Tweety." It is easy to talk about a population of repetitions for any particular observations whose evidential strength we want to assess. It is also easy to talk about causal models and conditional probabilities and populations of repetitions associated with them. But it is very hard to go beyond this talk and define these populations, even conceptually (Shafer 1982a). They usually do not provide a good starting point either for Bayesian or belief-function analyses.

I also do not think that purely statistical problems are the most important domain of application for belief functions. When adequate frequency information is available, there are usually many statistical methods from which to choose, and belief-function methods, though reasonable, may be more complicated and unfamiliar than standard methods. Belief-function representations of statistical evidence can be useful, however, when it is necessary to combine statistical and non-statistical evidence (Dempster 1990).

Belief functions are most useful precisely when it is not sensible to try to embed the joint occurrence of all our evidence in an imaginary population of repetitions. Belief functions are appropriate when different populations of repetitions, real or imagined, justify probability judgments for different items of evidence, and these populations bear on the question of interest and interact with each other in ways unique to each application. We saw this in the example of Section 3. The reliability of Sally and the reliability of Betty were both described in terms of imagined repetitions, repetitions in which Betty and Sally say many things about many things. But these reliabilities were linked to my car through what Betty and Sally said this particular time. Dempster's rule for combining their testimony was based on this unique aspect of the testimony, and it cannot be related to a population of repetitions.

9. Computation

The use of belief functions can involve challenging computational problems. In this section, I explain why belief functions are computationally complex, and I briefly review some ways of dealing with this complexity. These include Barnett's algorithm for the special case of belief functions focused on singletons, Thoma's fast Möbius transform, and tree propagation.

Why are belief functions computationally complex? A belief function Bel on a finite frame T with n elements is potentially far more complex than a probability measure P on the same frame. To specify a probability measure P, we need only n numbers—the probabilities $P(t)$ for each t in

T. To specify a belief function Bel, we may need as many as 2^n-2 numbers—the degrees of belief Bel(A) for every proper non-empty subset of T.

In practice, belief functions based on individual items of evidence do not approach this potential degree of complexity. Even if T is very large, we will try to sort our evidence into relatively simple items, each of which says something simple about T. We may have a problem, however, when we try to combine these belief functions by Dempster's rule.

The potential problems in implementing Dempster's rule can be seen using any of the mathematical formulations of Section 4. If we work with the basic probability assignment m, for example, then the straightforward approach is to implement formulas (3), (4) and (5). Each of these formulas is exponentially complex in itself, and each of them must be applied exponentially often—once for each subset of T.

One simple way to see the complexity that results from Dempster's rule is to count the potential focal elements of the belief functions involved. A *focal element* for a belief function with basic probability assignment m is a subset A such that m(A) 0. We might begin with belief functions with only a few focal elements. But we can see from (5) that the result of combining a belief function with r focal elements and a belief function with s focal elements will be a belief function with as many as r·s focal elements. Thus we may obtain a very complex belief function by combining a moderate number of belief functions that individually are not complex at all.

The belief functions we want to combine are usually initially defined on many different frames. This is because the different items of evidence we consider all bear on slightly different questions. In order to combine the belief functions, however, we must consider all these questions together; we must work in a frame whose elements answer all the questions at once. This frame may be enormous, even though the initial frames are all small. It is in this enormous frame, in theory, that Dempster's rule must be carried out. This is why the computational complexity of belief functions is a real problem.

Kong (1986a, 1988) introduced the idea of discussing belief-function combination in terms of variables. Conceptually, we can think of the elements of a frame T as the possible values of a variable X. In a problem of subjective judgment, we may be primarily interested in a variable X_1, with a relatively small frame T_1. But as we saw in Section 7, the different items of evidence that we are interested in, and their interaction, may force us to consider further variables. In general, we may be forced to consider a whole list of variables, say $X_2,X_3,...X_n$, with frames $T_2,T_3,...T_n$, respectively. Perhaps each item of evidence we consider is relevant to a few of these variables and hence each belief function we assess involves only the frame for a few variables. The first item of evidence might be relevant to the variables X_1, X_4 and X_5, for example, and hence be representable by a belief function on the Cartesian product $T_1 x T_4 x T_5$. But if we want to combine all the evidence, we find ourselves working in the frame $T_1 x T_2 x T_3 x...x T_n$, which may be enormous if n is large.

Barnett's Algorithm. Barnett (1981) gave algorithms for Dempster's rule for the special case in which each belief function supports either a single element of T or the complement of a single element. (More precisely, the focal elements for each belief function are either singletons or complements of singletons.) These algorithms are much better than the general algorithms; they are linear rather than exponential in the frame size.

In order to understand the potential and the limits of Barnett's algorithm, we must take into account the multivariate nature of belief-function computation. Individual items of evidence may in some cases produce belief functions that satisfy Barnett's conditions on a given frame. But as we study our problem, we will bring in new variables with new frames. A belief function that has singletons and their complements as focal elements in the frame $T_1 x T_4 x T_5$ will not have singletons and their complements as focal elements in the more refined frame $T_1 x T_2 x T_3 x...x T_n$. In this more refined frame, the focal elements will be elements of a given partition and their complements. Thinking of the belief function in this larger frame does not invalidate Barnett's algorithm; it can be thought of as working on the partition. But a different item of evidence,

17

assessed as a belief function satisfying Barnett's conditions on $T_3xT_4xT_6$, say, will involve a different partition in $T_1xT_2xT_3x...xT_n$, and Barnett's algorithm will not enable us to combine the two belief functions.

In practice, therefore, Barnett's algorithm must be supplemented by techniques for combining belief functions on related frames. Shafer and Logan (1987) successfully used Barnett's algorithm in this way in the case of hierarchical evidence.

Fast Möbius Transforms. With care, we can reduce the complexity of Dempster's rule even in the general case, because much of the computation in applying (3) (4) and (5) for each subset is repetitive. Thoma (1989) has shown how this repetition can be eliminated; the result is a fast Möbius transform, analogous to the fast Fourier transform.

It remains to be seen how useful the fast Möbius transform will be in practice. It is clear, however, that it is not enough to make arbitrary belief-function computations feasible. Even reducing the computational complexity from exponential to linear is not enough if the frame is enormous.

Propagation in Trees. More recent work on computation has focused on exploiting the pattern of evidence to reduce computation on very large frames to computation on many smaller frames. This can be done whenever individual items of evidence bear directly only on clusters of variables, and these clusters can be arranged in a join tree (Shafer, Shenoy and Mellouli 1987, Dempster and Kong 1988, Shafer and Shenoy 1988).

A tree in which the nodes are clusters of variables is called a *join tree* if whenever a variable is contained in two nodes of the tree, it is also contained in all the nodes on the path between these two nodes. Figure 1 shows an example of such a tree.


Figure 1. A join tree with ten variables.

Suppose we have independent items of evidence bearing on each cluster of variables in this tree. Each of these we represent by a belief function on the corresponding frame. The evidence bearing on $X_2, X_3,$ and X_4, for example is represented by a belief function on $T_2xT_3xT_4$. We want to combine all these belief functions on $T_1xT_2xT_3x...xT_{10}$. It turns out that we can do this (at least

we can get the values of the resultant belief function for statements about individual variables and the given clusters of variables) by applying Dempster's rule repeatedly within the smaller frames corresponding to the clusters in the tree. The information from one cluster relevant to another cluster can be passed along the path between the two, by means of messages between neighboring nodes that take the form of belief functions for the variables that the neighboring nodes have in common.

Work on belief functions in trees was initiated by the work of Gordon and Shortliffe (1984, 1985). They were concerned with the problem of combining evidence by Dempster's rule when different items of evidence are relevant to different levels of specificity in a hierarchy of diseases. Gordon and Shortliffe's suggestion of approximating Dempster's rule in this case was strengthened by Shafer and Logan (1987) to an efficient method of computing the exact results of the rule. Shenoy and Shafer (1986), Kong (1986a), and Shafer, Shenoy, and Mellouli (1987) explained how this treatment of hierarchical evidence can be understood as a special case of propagation in trees.

Bayesian and Belief-Function Propagation. Interest in belief-function propagation was inspired not only by Gordon and Shortliffe's work but also by the work on Bayesian propagation by Pearl (1986), which was based on models of conditional independence. How is the algorithm for belief-function propagation in trees related to algorithms for Bayesian propagation in trees developed by Pearl and by Lauritzen and Spiegelhalter (1988)?

It is possible to regard the Bayesian algorithms as special cases of the belief-function algorithm (Shenoy and Shafer 1986, Shafer, Shenoy, and Mellouli 1987, Zarley et al. 1988). In order to do this, however, each probability transition matrix, which provides the probabilities for one node as conditional probabilities given another node, must be represented by a belief function. This can be done, but the way in which it is done is neither efficient nor conceptually satisfying. The probability transition matrices represent different fragments of information about a single frequency distribution, and, as I explained in Section 8, such fragments should not, in general, be represented by belief functions and combined by Dempster's rule.

Rather than treating the Bayesian algorithm as a special case of the belief-function algorithm, therefore, we should think of them as two special cases of a more general algorithm. The scope of this more general algorithm has been demonstrated by Shenoy and Shafer (1990). It applies in any case where we have operations of marginalization and combination that satisfy a few simple axioms (Shenoy 1989a,b). As it turns out, this viewpoint clarifies Bayesian as well as belief-function propagation (Shafer and Shenoy 1988, 1990).

In both the belief-function and Bayesian cases, join trees are more than computational tools. They are also conceptual tools, tools that we use in sorting out our evidence. In the Bayesian case, they provide a graphical representation of the conditional independence structure that is needed to make probability judgments manageable. In the belief function case, they provide a representation of the sorting into independent uncertainties discussed in Section 7 (Dempster 1990).

Networks. When we sort our evidence into independent items, the clusters of variables that result may be such that they cannot be arranged in a join tree. This would be case in Figure 1, for example, if the variable X_3 were added to the cluster X_6,X_7. How to handle the computational problem in this case is a difficult problem in general.

The most satisfactory solution to the problem, if possible, is to find a join tree with slightly larger clusters such that each of the clusters with which we began can fit into a cluster in the tree. In the case of Figure 1 with X_3 added to the cluster X_6,X_7, for example, we can obtain such a tree by also adding X_3 to X_2,X_5,X_6.

Bayesian propagation also sometimes requires that we embed clusters of variables in a join tree. Lauritzen and Spiegelhalter (1988) give one way of doing this; Kong (1986a,b) gives another way, which usually results in a more manageable tree. How to embed collections of clusters in trees in the most efficient way is the subject of a growing literature (Rose 1970,

19

Bertele and Brioschi 1972, Tarjan and Yannakakis 1984, Arnborg et al. 1987, Mellouli 1987, Zhang 1990).

10. Implementing Belief Functions in Artificial Intelligence

Belief functions have been implemented in a wide variety of expert systems, and I am not prepared to evaluate or even list these implementations. The implementations I find most convincing, however, are those designed for interactive use. These include Gister (Lowrance, Garvey, and Strat 1986), Russell Almond's program (Almond 1988), DELIEF (Zarley et al. 1988), AUDITOR'S ASSISTANT (Shafer, Shenoy, and Srivastava 1988), and MacEvidence (Hsia and Shenoy 1989). These systems help human users build and evaluate belief networks. They require the user to make the judgments of independence that justify the network and to provide the numerical judgments of support based on each item of evidence.

Interactive systems seem appropriate to belief functions, since the theory practically requires that the relation between evidence and questions of interest should be unique to each application. Many probabilistic systems—such as the HUGIN system for medical diagnosis (Anderson et al. 1989)—apply the same conditional independence structure and, for the most part, the same numerical judgments to each new case. This means relating the entire structure of the evidence in each case to the same population of repetitions. As I have already argued, belief functions are appropriate to situations where this direct frequentist application of probability is not possible— situations where different populations of repetitions, real or imagined, justify probability judgments for different items of evidence, and these populations bear on the question of interest and interact with each other in ways unique to each application. Such uniqueness means that the belief network and the numerical judgments must be constructed anew for each case.

The next step, of course, is to automate this constructive process. This is difficult, but some progress has been made (Lowrance 1988, Andress and Kak 1988, Biswas and Anand 1989, Laskey et al. 1989).

Ultimately, it would be desirable to automate not only the construction of belief networks but also numerical assessment. I am not aware of current work in this direction, but current ideas in distributed memory (Kanerva 1989) encourage the idea that there is enough independence in such memories to permit the use of Dempster's rule.

11. Other Topics

In this final section, I briefly discuss several topics: generalizations of belief-function theory, decision methods based on belief functions, methods for reaching consensus using belief functions, work on belief-function weights of evidence, and other mathematical work, especially on infinite frames.

Generalizations. One way of generalizing belief-function theory is to retain the class of belief functions and to generalize Dempster's rule for combining belief functions. I have discussed some generalizations that allow for dependent evidence (1987a,b), but I am not aware of practical applications of these generalizations. In my view, the theory of belief functions should be used as a way of examining evidence (Dempster 1988), and in the examples I have thought about, this seems to lead to sorting out the independencies in the evidence, as in the example of Section 7.

Another approach is to retain the class of belief functions but to generalize the idea of a compatibility relation to allow for partial or probabilistic compatibility (Kohlas 1987, Yen 1989). In some cases, this idea can be seen not as a generalization at all, but rather as an introduction of frames intermediate between the S and T of Section 4. In other cases, the probabilistic basis of belief functions is lost, and the rationale for the generalization is not clear (Dubois and Prade 1986).

Finally, we can generalize the class of belief functions to more general set functions. Going in one direction, this takes us to the lower-probability theory discussed in Section 5 and even to versions of lower-probability theory that generalize probability bounds (Walley and Fine 1982, Wasserman and Kadane 1990). Going in another direction, away from probability but towards a variety of rules of combination, we enter the vast literature on fuzzy sets (Dubois and Prade 1986, 1988).

Decision Methods. Tom Strat, in his article in this volume, sets out a natural approach to decision theory within the belief function framework. The basic idea, that of upper and lower expectations, can also be found in Choquet (1953) and Dempster (1966). Examples and practical elaborations are given by Wesley (1984, 1986) and Wesley, Lowrance, and Garvey (1983).

Variations on this basic idea are possible, but I am not aware of extensive explorations of them. Two brief explorations can be found in Shafer (1986b; response to discussion) and Dempster and Kong (1987).

Another approach to relating belief functions to decision theory is to apply the usual von Neumann-Morgenstern theory to the linear space of belief functions. Jaffray (1989) has developed this approach.

Consensus for Belief Functions. Rationales for average or consensus belief functions for group decision have been developed by Wagner (1990) and Williams (1990). In Shafer (1986c), however, I argued against global consensus methods and in favor of direct reassessment by the group of the different items of evidence that contribute to a belief function.

Infinite Frames. Because my 1976 book dealt only with the case where the frame T is finite, some readers were left with the impression that the theory of belief functions had been developed only for this case. In fact, however, Dempster's articles had dealt from the outset with the continuous case. This is natural in the context of parametric statistical inference, since even the set of possible values for a single unknown probability is a continuous interval.

Dempster's early work included a treatment of the mathematics of belief functions generated by random closed intervals (Dempster 1968a). This topic was also treated by Strat (1984).

In general, the mathematical study of belief functions on continuous spaces will employ various regularity conditions, analogous to countable additivity or to familiar topological conditions in probability theory. These include topological conditions (Zhang and Wang 1987) or mere sequential continuity, analogous to countable additivity (Ross 1986). My own contributions to belief functions (Shafer 1976b, 1979) on continuous frames emphasized the condition of condensability, which was satisfied by the belief functions studied by Dempster. A condensable belief function is one for which the plausibility of every set A can be approximated by *finite* subsets of A; this expresses very strongly the intuition that in the case of subjective judgment, continuous mathematics should not be considered more than a convenient approximation to something fundamentally finite.

Weights of Evidence. My original purpose in writing my 1976 monograph was to explain an unsolved mathematical problem, which I called the *weight of conflict conjecture*. This conjecture derived from the idea of associating a weight of evidence to a belief function that supports a single subset of a frame to a certain degree. Such a belief function is called a *simple support function*. If A is the subset it supports, then the weight of evidence in favor of A is -log(1-Bel(A)). Weights of evidence add when simple support functions are combined by Dempster's rule. As I showed, all belief functions on finite frames can be obtained by combining simple support functions, reducing to coarser frames, and taking limits. Thus a belief function can always be understood in terms of a collection of weights of evidence. There is a possibility of competing representations by weights of evidence, however, because a given belief function can be obtained by coarsening in more than one way. The weight-of-conflict conjecture, if true, pointed to a way in which the competing representations would be in agreement. This conjecture remains unsettled, but the conjectured agreement was established by a different argument by

Zhang (1986). The possibility of using weights of evidence in belief-function applications remains largely unexploited.

    Other Mathematical Advances. Berres (1987) showed that products of belief functions are always belief functions. This has possible applications in the context of discounting belief functions (Shafer 1982a). Berres (1988) gave a comprehensive review of the relation between belief functions and Sugeno's λ-additive measures (Sugeno 1974). Another interesting advance is the extension to belief functions of ideas of information and entropy that have been associated with probability distributions (Smets 1983, Yager 1983).

References

Almond, Russell (1988). Fusion and propagation in graphical belief models. Research Report S-121, Department of Statistics, Harvard University.

Anderson, Stig K., Kristian G. Olesen, Finn V. Jensen, and Frank Jensen (1989). HUGIN—A shell for building Bayesian belief universes for expert systems. *IJCAI-89* 1080-1085.

Andress, K. and A. Kak (1988). Evidence accumulation and flow of control. *AI Magazine* **9:2** 75-94.

Arnborg, S., D.G. Corneil, and A. Proskurowski (1987). Complexity of finding embeddings in a k-tree. *SIAM Journal of Algebraic and Discrete Methods* **8** 277-284.

Baron, Johnathan (1987). Second-order probabilities and belief functions. *Theory and Decision* **23** 25-36.

Barnett, J.A. (1981). Computational methods for a mathematical theory of evidence. *IJCAI-81* 868-875.

Berres, Manfred (1987). On a multiplication and a theory of integration for belief and plausibility functions. *Journal of Mathematical Analysis and Applications* **121** 487-505.

Berres, Manfred (1988). λ-additive measures on measure spaces. *Fuzzy Sets and Systems* **27** 159-169.

Bertele, U. and F. Brioschi (1972). *Nonserial Dynamic Programming*. Academic Press.

Biswas, Gutam, and Tejwansh S. Anand (1989). Using the Dempster-Shafer scheme in a mixed-initiative expert system shell. Pp. 223-239 of *Uncertainty in Artificial Intelligence* **3**, L.N. Kanal, T.S. Levitt, and J.F. Lemmer, eds., North-Holland.

Brooks, Rodney A. (1986). A robust layered control system for a mobile robot. *IEEE Journal of Robotics and Automation* **RA-2:1** 14-23.

Cheeseman, Peter (1988). An inquiry into computer understanding (with discussion). *Computational Intelligence* **4** 58-66.

Choquet, Gustave (1953). Theory of capacities. *Annales de l'Institut Fourier, Grenoble* **5** 131-295.

Cohen, Paul R. (1987). The control of reasoning under uncertainty: A discussion of some programs. *The Knowledge Engineering Review* **2** 5-25.

Cox, Richard T. (1961). *The Algebra of Probable Inference*. The Johns Hopkins Press.

Dawes, Robyn M., David Faust, and Paul E. Meehl (1989). Clinical versus actuarial judgment. *Science* **243** 1668-1674.

Dempster, A.P. (1966). New methods for reasoning towards posterior distributions based on sample data. *Annals of Mathematical Statistics* **37** 355-374.

Dempster, A.P. (1967a). Upper and lower probabilities induced by a multivalued mapping. *Annals of Mathematical Statistics* **38** 325-339.

Dempster, A.P. (1967b). Upper and lower probability inferences based on a sample from a finite univariate population. *Biometrika* **54** 515-528.

Dempster, A.P. (1968a). Upper and lower probabilities generated by a random closed interval. *Annals of Mathematical Statistics* **39** 957-966.

Dempster, A.P. (1968b). A generalization of Bayesian inference (with discussion). *Journal of the Royal Statistical Society, Series B* **30** 205-247.

Dempster, A.P. (1969). Upper and lower probability inferences for families of hypotheses with monotone density ratios. *Annals of Mathematical Statistics* **40** 953-969.

Dempster, A.P. (1972). A class of random convex polytopes. *Annals of Mathematical Statistics* **43** 260-272.

Dempster, A.P. (1982). Comment. *Journal of the American Statistical Association* **77** 339-341.

Dempster, A.P. (1988). Probability, evidence, and judgment (with discussion). Pp. 284-298 of *Decision Making: Descriptive, Normative, and Prescriptive Interactions,* David E. Bell, Howard Raiffa, and Amos Tversky, eds., Cambridge.

Dempster, A.P. (1990). Construction and local computation aspects of network belief functions. Chapter 6 of *Influence Diagrams, Belief Nets, and Decision Analysis*, R.M. Oliver and J.Q. Smith, eds., Wiley.

Dempster, A.P., and Augustine Kong (1987). Discussion of papers by Lindley, Shafer, and Spiegelhalter. *Statistical Science* **1** 32-36.

Dempster, A.P., and Augustine Kong (1988). Uncertain evidence and artificial analysis. *Journal of Statistical Planning and Inference* **20** 355-368.

Dubois, Didier, and Henri Prade (1986). A set-theoretic view of belief functions. *International Journal of General Systems* **12** 193-226.

Dubois, Didier, and Henri Prade (1988). *Possibility Theory: An Approach to Computerized Processing of Uncertainty*. Plenum Press.

Duda, Richard O., and Peter E. Hart (1973). *Pattern Classification and Scene Analysis.* Wiley.

Efron, Bradley (1978). Controversies in the foundations of statistics. *American Mathematical Monthly* **85** 231-246.

Fagin, Ronald, and Joseph Y. Halpern (1989a). Uncertainty, belief, and probability. *IJCAI-89* 1161-1167.

Fagin, Ronald, and Joseph Y. Halpern (1989b). Updating beliefs vs. combining beliefs. Unpublished report.

Fishburn, Peter C. (1965). Analysis of decisions with incomplete knowledge of probabilities. *Operations Research* **13** 217-237.

Freeling, A.N.S., and Nils-Eric Sahlin (1983). Combining evidence. Pp. 58-74 of Gärdenfors et al. (1983).

Gärdenfors, Peter, Bengt Hansson, and Nils-Eric Sahlin, eds. ((1983). *Evidentiary Value: Philosophical, Judicial and Psychological Aspects of a Theory.* Lund.

Good, I.J. (1962). The measure of a non-measureable set. Pp. 319-326 of *Logic, Methodology, and Philosophy of Science*, Nagel, Suppes, and Tarski, eds., Stanford.

Good, I.J. (1982). Discussion of "Lindley's paradox," by Glenn Shafer. *Journal of the American Statistical Association* **77** 342-344.

Goodman, Irwin R., and Hung T. Nguyen (1985). *Uncertainty Models for Knowledge-Based Systems*. North-Holland.

Gordon, Jean, and Edward H. Shortliffe (1984). The Dempster-Shafer theory of evidence. In *Rule-Based Expert Systems: The MYCIN Experiments of the Stanford Heuristic Programming Project*, B.G. Buchanan and E.H. Shortliffe, eds., Addison-Wesley.

Gordon, Jean, and Edward H. Shortliffe (1985). A method for managing evidential reasoning in hierarchical hypothesis spaces. *Artificial Intelligence* **26** 323-358.

Halmos, Paul R. (1950). *Measure Theory*. Van Nostrand.

Hsia, Yen-Teh, and Prakash P. Shenoy (1989). An evidential language for expert systems. Pp. 9-16 of *Methodologies for Intelligent Systems*, Volume 4, Z. Ras, ed., North-Holland.

Isaacs, Herbert H. (1963). Sensitivity of decisions to probability estimation errors. *Operations Research* **11** 536-551.

Jaffray, Jean-Yves (1989). Linear utility theory for belief functions. *Operations Research Letters* **8** 107-112.

Kadane, J. (1984). *Robustness in Bayesian Statistics.* North-Holland.

Kanerva, Pentti (1988). *Sparse Distributed Memory*. MIT Press.

Kelley, John (1955). *General Topology*. Van Nostrand-Reinhold.

Kofler, E., Z.W. Kmietowicz, and A.D. Pearman (1984). Decision making with linear partial information. *Journal of the Operational Research Society* **35** 1079-1090.

Kohlas, Jürg (1987). Conditional belief structures. Technical report 131, University of Fribourg, Switzerland.

Kong, Augustine (1986a). Multivariate belief functions and graphical models. Doctoral dissertation, Department of Statistics, Harvard University.

Kong, Augustine (1986b). Construction of a tree of cliques from a triangulated graph. Technical Report S-118, Department of Statistics, Harvard University.

Kong, Augustine (1988). A belief function generalization of Gibbs ensembles. Technical Report 239, Department of Statistics, University of Chicago.

Krantz, David (1982). Foundations of the theory of evidence. Invited address. Society for Mathematical Psychology.

Krantz, David (1990). Judgments of likelihood evidence. Technical report, Department of Psychology, Columbia University.

Krantz, David, and J. Miyamoto (1983). Priors and likelihood ratios as evidence. *Journal of the American Statistical Association* **78** 418-423.

Kyburg, Henry E. (1987). Bayesian and non-Bayesian evidential updating. *Artificial Intelligence* **31** 297-294.

Laskey, Kathryn B. (1989). Belief in belief functions: An examination of Shafer's canonical examples. In *Uncertainty in Artificial Intelligence* **3**, L.N. Kanal, T.S. Levitt, and J.F. Lemmer, eds., North-Holland.

Laskey, Kathryn B., Marvin S. Cohen, and Anne W. Martin (1989). Representing and eliciting knowledge about uncertain evidence and its implications. *IEEE Transactions on Systems, Man, and Cybernetics* **19** 536-557.

Lauritzen, Steffen L., and David J. Spiegelhalter (1988). Local computations with probabilities on graphical structures and their application to expert systems (with discussion). *Journal of the Royal Statistical Society, Series B* **50** 157-224.

Levi, Isaac (1980). *The Enterprise of Knowledge*. MIT Press.

Lindley, Dennis V. (1972). *Bayesian Statistics: A Review*. Society for Industrial and Applied Mathematics.

Lowrance, John D. (1988). Automated argument construction. *Journal of Statistical Planning and Inference* **20** 369-387.

Lowrance, John D. Thomas D. Garvey, and Thomas M. Strat (1986). A framework for evidential-reasoning systems. *AAAI-86* 896-903.

Matheron, G. (1975). *Random Sets and Integral Geometry.* Wiley.

Mellouli, Khaled (1987). On the propagation of beliefs in networks using the Dempster-Shafer theory of evidence. Doctoral dissertation, School of Business, University of Kansas.

McDermott, Drew V. (1987). Critique of pure reason (with discussion). *Computational Intelligence* **3** 151-160.

Neveu, Jacques (1965). *Mathematical Foundations of the Calculus of Probability*. Holden-Day.

Nguyen, Hung T. (1978). On random sets and belief functions. *Journal of Mathematical Analysis and Applications* **65** 531-542.

Nilsson, Nils. J. (1965). *Learning Machines.* McGraw-Hill.

Nilsson, Nils J. (1986). Probabilistic logic. *Artificial Intelligence* **28** 71-87.

Pearl, Judea (1986). Fusion, propagation and structuring in belief networks. *Artificial Intelligence* **29** 241-288.

Pearl, Judea (1988). *Probabilistic Reasoning in Intelligent Systems*. Morgan Kaufmann.

Rose, D. J. (1970). Triangulated graphs and the elimination process. *Journal of Mathematical Analysis and Applications* **32** 597-609.

Ross, David (1986). Random sets without separability. *Annals of Probability* **14** 1064-1069.

Rota, Gian-Carlo (1964). Theory of Möbius functions. *Zeitschrift für Wahrscheinlichkeitstheorie and Verwandte Gebiete* **2** 340-368.

Ruspini, Enrique H. (1987). Epistemic logics, probability, and the calculus of evidence. *IJCAI-87* 924-931.

Savage, Leonard J. (1954). *The Foundations of Statistics.* New York: Wiley.

Savage, Leonard J. (1961). The foundations of statistics reconsidered. Pp. 575-586 of *Proceedings of the Fourth Berkeley Symposium on Mathematical Statistics and Probability*, Volume 1, Jerzy Neyman, ed., University of California Press.

Shafer, Glenn (1973). Allocations of probability. Doctoral dissertation, Department of Statistics, Princeton University.

Shafer, Glenn (1976a). *A Mathematical Theory of Evidence*. Princeton University Press.

Shafer, Glenn (1976b). A theory of statistical evidence. Pp. 365-436 of *Foundations of Probability Theory, Statistical Inference, and Statistical Theories of Science*, Volume II, W.L. Harper and C.A. Hooker, eds., Reidel.

Shafer, Glenn (1978). Non-additive probabilities in the work of Bernoulli and Lambert. *Archive for History of Exact Sciences* **19** 309-370.

Shafer, Glenn (1979). Allocations of probability. *Annals of Probability* **7** 827-839.

Shafer, Glenn (1981). Constructive probability. *Synthese* **48** 1-60.

Shafer, Glenn (1982a). Lindley's paradox (with discussion). *Journal of the American Statistical Association* **77** 325-351.

Shafer, Glenn (1982b). Belief functions and parametric models (with discussion). *Journal of the Royal Statistical Society, Series B* **44** 322-352.

Shafer, Glenn (1984). The problem of dependent evidence. Working Paper 164, School of Business, University of Kansas.

Shafer, Glenn (1986a). The combination of evidence. *International Journal of Intelligent Systems* **1** 155-179.

Shafer, Glenn (1986b). Savage revisited (with discussion). *Statistical Science* **1** 463-501.

Shafer, Glenn (1986c). Discussion of "Combining probability distributions," by C. Genest and J.V. Zidek. *Statistical Science* **1** 135-137.

Shafer, Glenn (1987a). Probability judgment in artificial intelligence and expert systems (with discussion). *Statistical Science* **2** 3-44.

Shafer, Glenn (1987b). Belief functions and possibility measures. Pp. 51-84 of *The Analysis of Fuzzy Information*, Volume 1, James C. Bezdek, ed., CRC Press.

Shafer, Glenn (1988). Discussion of "An inquiry into computer understanding," by Peter Cheeseman. *Computational Intelligence* **4** 121-124.

Shafer, Glenn (1990). The unity of probability. Pp. 95-126 of *Acting Under Uncertainty: Multidisciplinary Conceptions*. George von Furstenberg, ed., Kluwer.

Shafer, Glenn, and Roger Logan (1987). Implementing Dempster's rule for hierarchical evidence. *Artificial Intelligence* **33** 271-298.

Shafer, Glenn, and Prakash P. Shenoy (1988). Local computation in hypertrees. Working Paper 201, School of Business, University of Kansas.

Shafer, Glenn, and Prakash P. Shenoy (1990). Probability propagation. To appear in *Annals of Mathematics and Artificial Intelligence*.

Shafer, Glenn, Prakash P. Shenoy, and Khaled Mellouli (1987). Propagating belief functions in qualitative Markov trees. *International Journal of Approximate Reasoning* **1** 349-400.

Shafer, Glenn, Prakash P. Shenoy, and Rajendra P. Srivastava (1988). AUDITOR'S ASSISTANT: A Knowledge-Engineering Tool for Audit Decisions. Pp. 61-84 of *Proceedings of the 1988 Touche-Ross/ University of Kansas Symposium*.

Shafer, Glenn and Rajendra P. Srivastava (1990). The Bayesian and belief-function formalisms: A general perspective for auditing. To appear in *Auditing: A Journal of Practice and Theory*, Volume 9 Supplement.

Shafer, Glenn and Amos Tversky (1985). Languages and designs for probability judgment. *Cognitive Science* **9** 309-339.

Shenoy, Prakash P. (1989a). A valuation-based language for expert systems. *International Journal of Approximate Reasoning* **3** 383-411.

Shenoy, Prakash P. (1989b). On Spohn's rule for revision of beliefs. Working Paper 213, School of Business, University of Kansas.

Shenoy, Prakash P., and Glenn Shafer (1986). Propagating belief functions using local computations. *IEEE Expert* **1:3** 43-52.

Shenoy, Prakash P., and Glenn Shafer (1990). An axiomatic framework for Bayesian and belief-function propagation. To appear in *Uncertainty in Artificial Intelligence* **4**, R. Shachter, T.S. Levitt, J.F. Lemmer, and L.N. Kanal, eds., North-Holland.

Smets, Philippe (1983). Information content of an evidence. *International Journal of Man-Machine Studies* **19** 33-43.

Smets, Philippe (1988). Belief functions. Pp. 253-286 of *Non-Standard Logics for Automated Reasoning*, Philippe Smets, Abe Mamdani, Didier Dubois, and Henri Prade, eds., Academic Press.

Smith, C.A.B. (1961). Consistency in statistical inference and decision (with discussion). *Journal of the Royal Statistical Society, Series B* **23** 1-25.

Snow, Paul (1986). Bayesian inference without point estimates. *AAAI-86* 233-237.

Speigelhalter, D.J. (1986). A statistical view of uncertainty in expert systems. Pp. 17-56 of *Artificial Intelligence and Statistics*, W. Gale, ed. Addison-Wesley.

Srivastava, Rajendra P., Prakash P. Shenoy, and Glenn Shafer (1989). Audit risk: A belief-function approach. Working Paper 212, School of Business, University of Kansas.

Strat, Tom (1984). Continuous belief functions for evidential reasoning. *AAAI-84* 308-313.

Sugeno, M. (1974). Theory of fuzzy integrals and its applications. Doctoral dissertation, Tokyo Institute of Technology.

Suppes, Patrick (1974). The measurement of belief. *Journal of the Royal Statistical Society, Series B* **36** 160-191.

Tarjan, R. E. and M. Yannakakis (1984). Simple linear time algorithms to test chordality of graphs, test acyclicity of hypergraphs, and selectively reduce acyclic hypergraphs. *SIAM Journal of Computing* **13** 566-579.

Thoma, Hans Mathis (1989). Factorization of belief functions. Doctoral Dissertation, Department of Statistics, Harvard University.

Trichtler, David, and Gina Lockwood (1990). Modelling the reliability of paired comparisons. To appear in *Journal of Mathematical Psychology*.

Wagner, Carl (1990). Consensus for belief functions and related uncertainty measures. To appear in *Theory and Decision*.

Walley, Peter (1987). Belief-function representations of statistical evidence. *Annals of Statistics* **15** 1439-1465.

Walley, Peter, and Terrance L. Fine (1982). Towards a frequentist theory of upper and lower probability. *Annals of Statistics* **10** 741-761.

Wasserman, Larry Alan (1989). Belief functions and statistical inference. Technical report 458, Department of Statistics, Carnegie Mellon University.

Wasserman, Larry Alan (1990). Prior envelopes based on belief functions. To appear in *Annals of Statistics*.

Wasserman, Larry Alan, and Jay Kadane (1990). Bayes's theorem for Choquet capacities. To appear in *Annals of Statistics*.

Weisberg, H. (1969). Upper and lower probability inferences from ordered multinomial data. Doctoral dissertation, Department of Statistics, Harvard University.

Wesley, Leonard P. (1984). Reasoning about control: an evidential approach. Technical report 324, Artificial Intelligence Center, SRI International.

Wesley. Leonard P. (1986). Evidential knowledge-based computer vision. *Optical Engineering* **25:3** 363-379.

Wesley, Leonard P., John D. Lowrance, and Thomas D. Garvey (1983). Reasoning about control: The investigation of an evidential approach. *IJCAI-83* 203-206.

West, Sandra (1971). Upper and lower probability inference for the logistic function. Doctoral dissertation, Department of Statistics, Harvard University.

Williams, Peter M. (1982). Discussion of "Belief functions and parametric models," by Glenn Shafer. *Journal of the Royal Statistical Society, Series B* **44** 341-342.

Williams, Peter M. (1990). Averaging and eliciting expert opinion. To appear in *Artificial Intelligence*.

Winograd, Terry (1980). Extended inference modes in reasoning by computer systems. *Artificial Intelligence* **13** 5-26.

Yager, Ronald R. (1983). Entropy and specificity in a mathematical theory of evidence. *International Journal of General Systems* **9** 249-260.

Yen, John (1989). GERTIS: A Dempster-Shafer approach to diagnosing hierarchical hypotheses. *Communications of the ACM* **32:5** 573-585.

Zadeh, Lotfi A. (1978). Fuzzy sets as a basis for a theory of possibility. *Fuzzy Sets and Systems* **1** 3-28.

Zadeh, Lotfi A. (1986). A simple view of the Dempster-Shafer theory of evidence and its implications for the rule of combination. *AI Magazine* **7:2** 85-90.

Zarley Debra, Yen-Teh Hsia, and Glenn Shafer (1988). Evidential reasoning using DELIEF. *AAAI-88* 205-209.

Zhang, Lianwen (1986). Weights of evidence and internal conflict for support functions. *Information Sciences* **38** 205-212.

Zhang, Lianwen, and Wang Peizhuang (1987). A generalization of Choquet's theorem to measurable spaces. *Kexue Tongbao* **32** 1596-1601.

Zhang, Lianwen (1990). Studies on hypergraphs. Working Paper 215, School of Business, University of Kansas.

Improved references:

Almond, Russell George (1989). Fusion and Propagation of Graphical Belief Models: An Implementation and an Example. Doctoral dissertation (Research Report S-130). Department of Statistics, Harvard University.

Wasserman, Larry Alan (1990). Prior envelopes based on belief functions. *The Annals of Statistics* **18** 454-464.