A Framework for Comparing Alternative Formalisms for Plausible Reasoning

Eric J. Horvitz, David E. Heckerman, Curtis P. Langlotz

Medical Computer Science Group
Knowledge Systems Laboratory
Departments of Medicine and Computer Science
Stanford University
Stanford, California 94305

ABSTRACT

We present a logical relationship between a small number of intuitive properties for measures of belief and the axioms of probability theory. The relationship was first demonstrated several decades ago but has remained obscure. We introduce the proof and discuss its relevance to research on reasoning under uncertainty in artificial intelligence. In particular, we demonstrate that the logical relationship can facilitate the identification of differences among alternative plausible reasoning methodologies. Finally, we make use of the relationship to examine popular non-probabilistic strategies.

I INTRODUCTION

As artificial intelligence research has extended beyond deterministic problems, methodologies for reasoning under uncertainty or *plausible reasoning* have become increasingly central. Several competing approaches to reasoning in complex and uncertain settings have been formulated. These include probability [1], fuzzy logic [2], Dempster-Shafer theory [3], certainty factors [4], and multi-valued logics [5]. There has been debate on the theoretical and pragmatic benefits and disadvantages of these alternative strategies.

A particular focus of discussion has centered around the adequacy of probability theory for handling reasoning under uncertainty [6]. While probabilists have defended the use of probability, others have cited benefits achieved through the use of non-probabilistic formalisms [2, 3, 7, 4]. Such discussion has been heightened in recent years as the demand has grown for applicable methodologies for reasoning under uncertainty.

In this paper, we discuss the ramifications of a proof showing that the axioms of probability theory follow logically from a set of simple properties. We shall reformulate the work of R.T. Cox, a physicist interested in reasoning under uncertainty. Cox demonstrated, over forty years ago, that the axioms of probability theory are a necessary consequence of intuitive properties of measures of belief [8]. That is, if a set of simple properties are assumed, the axioms of probability theory must be accepted. Even though others, including Jaynes [9] and Tribus [10] have since demonstrated similar proofs, the work has remained obscure.

We think it is important that the artificial intelligence community become familiar with Cox's result. After clarifying the focus of this paper, we will present fundamental properties that Cox and Jaynes have asserted as necessary for any measure of belief. We will then discuss the relevance of the proof to current discussions within the artificial intelligence community on the use of alternative formalisms for plausible reasoning. Finally, we will describe how the proof can serve as a framework for analyzing and communicating differences about alternative methodologies for plausible reasoning. We will use the framework to critique the non-probabilistic methods of fuzzy logic, the Dempster-Shafer theory of belief functions, and the MYCIN certainty factor model.

*This work was supported in part by the Josiah Macy, Jr. Foundation, the Henry J. Kaiser Family Foundation, and the Ford Aerospace Corporation. Computing facilities were provided by the SUMEX-AIM resource under NIH grant RR-00785.

II THE LIMITS OF BELIEF ENTAILMENT

We intend to present a useful perspective on methodologies for the entailment of belief. We use the phrase *belief entailment* to refer to the consistent assignment of measures of belief to propositions, *in the context* of established belief. Belief entailment schemes, such as the MYCIN certainty factor model, fuzzy logic, and probability theory dictate the belief in Boolean combinations of propositions given measures of belief in component propositions. Entailment schemes also provide a mechanism whereby beliefs can be updated as new information becomes available.

Some have rightly pointed out that theories of belief entailment do not capture the rich semantics of plausible reasoning [7]. We stress that such methods are, indeed, only intended for the relatively simple task of the consistent assignment of measures of belief. We believe that belief entailment should be distinguished from the more encompassing task of plausible reasoning.

It is useful to decompose the problem of reasoning under uncertainty into three distinct components: problem formulation, initial belief assignment, and belief entailment. We use the term *problem formulation* to refer to the task of constructing the plausible reasoning problem. This consists of the process of enumerating important propositions as well as relations among propositions. The initial *assignment of belief* requires the direct assessment of belief or some procedure for constructing belief. Belief entailment occurs after a problem is formulated and an initial assignment of belief is completed.

Belief entailment methodologies are relatively well-developed. For example, there are a number of different axiomatic schemes to choose from. In contrast, aspects of problem formulation and belief construction currently pose significant challenges for artificial intelligence research. Problem formulation has proven to be particularly difficult; there has been continuing debate as to whether or not an axiomatic theory for problem formulation is possible at all [11, 12, 13].

From this point on, we shall explicitly distill away the problems and issues concerning problem formulation in our discussion of alternative methods for reasoning under uncertainty.

III FUNDAMENTAL PROPERTIES OF BELIEF

We now turn to the intuitive basis for probability theory. We shall assert a set of fundamental properties for measures of belief. The intuitive basis is a reformulation of the properties asserted by Cox, Jaynes, and Tribus as being essential for any measure of belief that could vary between truth and falsehood. We have attempted to make explicit all the properties used in the classic proof, including those that were not emphasized in the original work. We will enumerate seven properties for measures of belief and name each for reference.

Like many artificial intelligence researchers, Cox reasoned about the problem of uncertainty from a deductive perspective. He sought essential properties required of a measure that represented a *degree* of belief in the truth of a Boolean proposition, or propositions created through the combination rules of Boolean algebra.

The first property in our formulation asserts that propositions to which belief can be assigned must be well-defined. That is, propositions must be defined precisely enough so that it would be theoretically possible to determine whether a proposition is indeed true or false. The intention of this property is captured by the notion that a proposition should be defined clearly enough so that an omniscient being (often referred to in the decision analysis literature as a *clairvoyant*) could determine its truth or falsehood. We shall refer to this requirement as the *clarity* property.

A second assertion is that measures of degree of belief in the truth of a proposition should be able to vary continuously between values of certain truth and certain falsehood and that the continuum of belief can be represented by a single real number. We refer to the use of a single real number to represent continuous measures of belief as the *scalar continuity* property.

A third assertion in our formulation is that it is possible to assign a degree of belief to *any* proposition which is precisely defined. We refer to this property as the *completeness* property.

Upon what might a degree of belief depend? An individual's or computer program's degree of belief in a proposition should, of course, depend on the particular proposition under consideration. In addition, the degree of belief in a particular proposition can depend upon knowledge about other propositions. We refer to this as the *context dependency* property. We shall use the term $Q|e$ to represent the degree of belief in proposition Q by an individual with background knowledge e. The background knowledge e refers to information relevant to the belief in Q that is assumed or believed to be true.

In exploring the dependency of belief in one proposition on another, Cox specifically focused on the belief in the conjunction of two propositions given belief in each proposition. He asserted that the belief in the proposition QR should be related to the belief in Q alone as well as to the belief in R given that Q is true. That is, the belief in an event of interest should depend on one's belief in the event given the truth of some conditioning event as well as the degree of belief in the conditioning event itself. Formally, we assert that measures of belief should have the property that there exists some function F such that

$$QR|e = F(Q|e,\ R|Qe). \qquad (1)$$

The function is asserted to be continuous and monotonically increasing in both arguments when the other is held constant. We refer to the above property as the *hypothetical conditioning* property. This property is related to what is often referred to as hypothetical reasoning. Individuals commonly assign belief to events conditioned on the truth of other events. This property may be viewed as a specialization of *context dependency*.

Although the *hypothetical conditioning* property was stated by Cox, it can actually be proved from a weaker assumption about the relationship of belief in the conjunction of two propositions to belief in the individual component propositions. The proof considers functions of belief in two propositions that could generate a measure of belief in the conjunction of the propositions. Alternative arguments are eliminated based on contradiction and symmetry, leaving only the *hypothetical conditioning* form. This work is elegantly presented in Tribus [10].

Another assertion is that belief in the negation of a proposition Q, denoted ~Q, should be determined by the belief in the proposition itself. Formally, there should be a continuous monotonically decreasing function G such that

$$\sim Q|e = G(Q|e).$$

We refer to this as the *complementarity* property.

A final property focuses on logically equivalent propositions. Consider two propositions, Q and R where it is possible to show that Q logically implies R and vice-versa. In this case, we assert that $Q|e = R|e$ for any prior information e. In other words, if two propositions have the same truth value, an individual should believe each of them with equal conviction. We term this the *consistency* property.

In summary, we have presented seven fundamental properties for continuous measures of belief in the truth of propositions. We have termed these properties:

1. <u>Clarity</u>: Propositions should be well-defined.
2. <u>Scalar continuity</u>: A single real number is both necessary and sufficient for representing a degree of belief.
3. <u>Completeness</u>: A degree of belief can be assigned to any well-defined proposition.
4. <u>Context dependency</u>: The belief assigned to a proposition can depend on the belief in other propositions.
5. <u>Hypothetical conditioning</u>: There exists some function that allows the belief in a conjunction of propositions to be calculated from the belief in one proposition and the belief in the other proposition given that the first proposition is true.
6. <u>Complementarity</u>: The belief in the negation of a proposition is a monotonically decreasing function of the belief in the proposition itself.
7. <u>Consistency</u>: There will be equal belief in propositions that have the same truth value.

IV A LOGICAL MAPPING

As mentioned above, Cox and others have demonstrated that the above properties logically necessitate the axioms of probability theory. According to the proof, if one accepts the above intuitive properties, one must accept the axioms of probability. More precisely, it can be shown that if the intuitive properties of belief are assumed, there exists a continuous monotonic function ϕ such that

$$0 \leq \phi(Q|e) \leq 1$$

$$\phi(\text{TRUE}|e) = 1$$

$$\phi(Q|e) + \phi(\sim Q|e) = 1$$

$$\phi(QR|e) = \phi(Q|e) \cdot \phi(R|Qe)$$

These relations are the axioms of probability theory as they are commonly formulated. That is, $\phi(Q|e)$ satisfies the axioms of probability. Given the above fundamental properties, the *only* measure of belief in the truth of proposition Q in light of evidence E *must* be the probability of Q given E, written $p(Q|E)$ or some monotonic transformation of this quantity. Bayes' theorem follows directly from the last axiom above.

The proof of the necessary relationship between the intuitive properties and the axioms of probability theory is based on an analysis of solutions for the *functional forms* [14] implied by the intuitive properties. We recommend the referenced versions of the proof to the reader.

V RELEVANCE OF THE MAPPING TO AI RESEARCH

The logical mapping relating intuitive properties to the axioms of probability has important implications for artificial intelligence research. In the context of the mapping, if one subscribes to the simple intuitive properties of belief in systems that reason under uncertainty, one thereby agrees that the axioms of probability are theoretically sound for capturing all aspects of belief entailment; arguments for alternative entailment schemes based on theoretical or pragmatic considerations must involve the violation or modification of one or more of the enumerated properties.

In addition to serving as a proof of logical necessity between the intuitive properties and the axioms of probability theory, the Cox result can provide a useful perspective on the differences between alternative entailment methodologies. Investigation of differences between competing methods may be hampered when the formalisms are defined with axiomatic systems that are difficult to compare. As an example, the most common axiomatization of Dempster's theory of evidence [15] is in a form not particularly suited for comparison with the axioms of probability theory. As we will see in section VII, moving discussion into the realm of intuitive properties can highlight the fundamental differences between alternative belief entailment schemes.

The mapping can be especially helpful in identifying the basis of possible dissatisfaction with probability theory. An individual, harboring ill-defined dissatisfaction with probability theory, might be able to identify the sources of his uneasiness at the level of the intuitive properties.

VI A FRAMEWORK FOR COMPARING ALTERNATIVES

Cox's proof of a mapping between a set of intuitive properties and probability theory can serve as an integrative framework for identifying differences among alternative belief entailment schemes. We believe that the set of intuitive properties are so basic as to be relevant to any belief entailment scheme; the properties, or close analogs of them, were undoubtedly addressed in the creation of the methodologies. Ascertaining the status of each of the fundamental properties in a non-probabilistic methodology is usually straightforward.

How can we critique alternatives of probability in terms of the intuitive properties? It is useful to carefully identify the status of the seven intuitive properties in each entailment methodology. In most cases, the spirit of a non-probabilistic methodology can be captured by identifying a fundamental dissimilarity with one or two of the intuitive properties of probability theory. Although such a difference will often have the side effect of invalidating other intuitive properties, it may still be useful to focus on the primary property violation that best *captures the rationale* behind the creation of the method.

The identification of a primary property violation can focus debate on well-defined fundamental principles. Such a focus can be especially useful in discussions of perceived theoretical advantages of alternative belief entailment schemes. When the selection of a scheme is based on the pragmatics of computation or belief assessment, identifying a central property violation can be useful in characterizing problems that may arise in practice.

It may also be useful to categorize the differences between probability and alternative belief entailment methods. Such a categorization scheme can summarize agreement of any methodology with the intuitive properties. If we examine the status of the seven intuitive properties of belief, non-probabilistic strategies can be viewed to fall into one of the following categories:

1. Generalization: The elimination or weakening of particular intuitive properties.

2. Specialization: The addition of new fundamental properties or the strengthening of existing properties.

3. Self-inconsistency: The addition or strengthening of properties such that a logical inconsistency arises in the set of fundamental properties; the set of properties become self-inconsistent.

4. Substitution: The substitution of one or more properties for another such that the set does not fall into one of the above categories.

Armed with this intuitive framework, we will now explore specific examples of popular belief entailment methodologies that are often viewed as competing with probability theory. In particular we will examine fuzzy logic, the Dempster-Shafer theory of belief functions, and the MYCIN certainty factor model.

VII EXAMINATION OF ALTERNATIVE METHODS

A. Fuzzy Logic

There are currently at least two distinct forms of fuzzy logic used to manage uncertainty. Each deviates from the intuitive properties in a different way.

One form of fuzzy reasoning applied to managing uncertainty was introduced by Zadeh [2]. Fuzzy logicians using this methodology do not object to the use of probability theory when events are precisely defined. However, they argue that it is desirable to reason with imprecision in the definition of events in addition to uncertainty about their occurrence. They allow beliefs to be assigned to imprecise events as well as precise ones. This version of fuzzy logic theory includes fuzzy versions of Bayes' theorem [16]. Zadeh attempts to demonstrate the need to assign belief to fuzzy propositions in the following challenge:

> An urn contains *approximately* n balls of various sizes, of which *several* are *large*. What is the probability that a ball drawn at random is *large* [16]?

Returning to our intuitive properties, it appears that the central dissimilarity of this kind of fuzzy logic with probability theory occurs with the *clarity* property. This methodology weakens the *clarity* property in that it is assumed that events to which belief may be assigned remain ill-defined. We would classify this school of fuzzy logic as being a *generalization* of probability theory.

The identification that a central difference between this form of fuzzy logic and probability theory occurs at the level of the *clarity* property defines a particular focus for discussion about the benefits or rationale of the methodology. Analysis of the advantages and disadvantages of fuzzy logic should center on the rationale and ramifications of weakening the *clarity* property. Many probabilists have argued against the weakening of the *clarity* property by pointing out that imprecision in the specification of a proposition could always be converted to uncertainty about the occurrence of a related precise event that had similar or identical semantic content. It has also been proposed that probability distributions over variables of interest can capture the essence of fuzziness within the framework of probability [17].

It has also been argued that the use of imprecise propositions is inappropriate in making important decisions. The penalty for reasoning with fuzzy events is often obscured by the examples used in presentations of fuzzy set theory. Typical examples tend to center on reasoning about events with small potential losses and gains. For example, it is generally not very important whether or not a person of height 4' 10" is called "short." However, problems with using fuzzy events may be more apparent when large potential utility changes are associated with events. The cost for relying on a fuzzy entailment calculus is highlighted by the following example of a high-stakes situation:

> Stan finally received news about the growth on his chin. His physician, who was quite fond of fuzzy logic, reported to his nervous patient, "The test results *usually* mean that it is *somewhat likely* that you have cancer. As the tumor is *quite large* and *probably dangerous*, I will operate. You shouldn't worry; my patients *usually* survive such operations.

A decision theorist might argue that, in general, lack of clarity as in the above problem will lead to lower expected utility of outcome. That is, a cost is incurred by foregoing the use of the *clarity* property. The comparison of fuzzy and precise versions of a problem would allow an actual penalty associated with loss of information to be calculated. Decision theorists might argue that imprecision may not be tolerable in certain high stakes situations.

We move next to an alternative fuzzy logic methodology. In this methodology [5], the degree of membership of a proposition P in the set of true propositions, denoted $\mu_T(P)$, is interpreted as the degree of belief in the hypothesis. That is,

$$\mu_T(P) \equiv P|e. \qquad (2)$$

We should note that a logical equivalency between this brand of fuzzy reasoning and forms of multi-valued logic has been demonstrated [5]. In this approach, it is assumed that

$$\mu_T(AB) = MIN(\mu_T(A), \mu_T(B)).$$

Given the correspondence (2), we see that this brand of fuzzy methodology is not consistent with the *hypothetical conditioning* property. Therefore, this form of fuzzy reasoning falls into the *substitution* category above.

Probabilists who accept $\mu_T(A)$, $\mu(B)_T$, and $\mu_T(AB)$ as measures of belief would object to the violation of the *hypothetical conditioning* relation. They might argue that the final belief in the conjunction $\mu_T(AB)$ is not necessarily dependent solely on $\mu_T(A)$ and $\mu_T(B)$. The violation of *hypothetical conditioning* in this case is tantamount to imposing independence or uniform conditional dependence (equivalent dependency among all propositions) where such a relationship may not exist.

B. Dempster-Shafer

In the Dempster-Shafer (DS) theory [3], two separate measures of belief can be assigned to each proposition P. These measures are referred to as the "belief" and "plausibility" in P, denoted Bel(P) and Plaus(P) respectively. Also, Bel(P) is not directly related to Bel(~P); instead, Bel(P) = 1 - Plaus(~P). Similarly, Plaus(P) = 1 - Bel(~P). Thus, the theory appears to differ from probability theory with respect to the *scalar continuity* property as well as the *complementarity* property. However, an examination of the original motivation for the theory reveals a more fundamental difference; the DS theory allows for the existence of well-defined hypotheses to which degrees of belief cannot be assigned. Thus, it seems that the central issue behind the development of the DS theory is the weakening of the *completeness* property. The fact that two numbers can be attached to the belief in any hypothesis is a consequence of this more fundamental difference between the two theories. To illustrate this, consider the following problem taken from Shafer [18]:

> Is Fred, who is about to speak to me, going to speak truthfully, or is he, as he sometimes does, going to speak carelessly, saying something that comes to his mind, paying no attention to whether it is true or not? Let S denote the possible answers to this question; S = {truthful, careless}. Suppose I know from experience that Fred's announcements are truthful reports on what he knows about 80% of the time and are careless statements the other 20% of the time. Then I have a probability measure p over S: p{truthful} = .8, p{careless} = .2.
>
> Are the streets outside slippery? Let T denote the possible answers to this question; T = {yes, no}. And suppose Fred's answer to this question turns out to be, "The streets outside are slippery." Taking account of this, I have a compatibility relation between S and T; "truthful" is compatible with "yes" but not with "no," while "careless" is compatible with both "yes" and "no."

If one wanted to use probability theory to determine the belief in the hypothesis that the streets outside are slippery given Fred's report, additional information would be needed. In particular, one's *prior* belief that the streets outside are slippery and the conditional belief that Fred will be correct given that he is careless will be required. If r is the needed prior belief that the streets outside are slippery and if s is the conditional belief that Fred will be correct given that he is careless, the belief of interest can be calculated using Bayes' theorem:

$$p(slippery|report) = \frac{.8r + .2rs}{.8r + .2rs + .2(1-r)(1-s)}$$

In DS theory, one is allowed to assert that r and s *cannot be assessed*. To make up for this lack of information, the theory uses the "compatibility relation" described above in order to define beliefs relevant to the problem. The DS "belief" and "plausibility" that the roads are slippery ("yes") are given by

$$Bel(\{"yes"\}) = \sum_{x|xCy \Rightarrow y \subseteq \{"yes"\}} p(x)$$

$$= p("truthful") = .8$$

$$Plaus(\{"yes"\}) = \sum_{x|xCy \Rightarrow y \cap \{"yes"\} \neq 0} p(x)$$

$$= p("truthful")+p("careless") = 1.$$

where xCy means that x in S and y in T are compatible. Thus, the violation of the *scalar continuity* and *complementarity* properties arises from a weakening of the *completeness* property. Based in the weakening of this property, DS can be considered a *generalization* of probability theory.

Many have objected to the weakening of the *completeness* property. For example, most decision analysts would insist that a personal measure of belief can be assigned to *any* well-defined proposition when placed in the context of a decision. There has been research in the decision analysis community focusing on the pragmatics of assessing belief in any well-defined proposition.

C. Certainty factors

We now turn to the MYCIN certainty factor model used for belief entailment in a number of rule-based systems. The MYCIN certainty factor model [4] can be shown to be self-inconsistent [19, 20]. Thus, the original certainty factor model falls into the third category above. There are several ways to demonstrate inconsistency in the model. We will outline one of these approaches here. The model is an augmentation to the rule-based representation paradigm. Knowledge is represented as rules of the form IF <evidence> THEN <hypothesis>. To each rule is attached a *certainty factor*, denoted CF(H,E), which is intended to represent the *change* in belief in hypothesis H given that evidence E becomes known. The definition of CF(H,E) is given in the original work:

$$CF(H,E) = \begin{cases} \dfrac{p(H|E) - p(H)}{1 - p(H)} & p(H|E) > p(H) \\ \\ \dfrac{p(H|E) - p(H)}{p(H)} & p(H) > p(H|E) \end{cases} \qquad (3)$$

where p(H) is the *prior* probability of H and p(H|E) is the *posterior* probability of H given E.

One component of the model involves a prescription for combining certainty factors. For example, suppose two pieces of evidence E_1 and E_2 bear on hypothesis H. In the model, the two certainty factors $CF(H,E_1)$ and $CF(H,E_2)$ are combined to give an effective certainty factor, $CF(H,E_1 \wedge E_2)$, for the rule IF $E_1 \wedge E_2$ THEN H with the following function:

$$z = \begin{cases} x + y - xy & x,y \geq 0 \\ \\ \dfrac{x + y}{1 - \min(|x|,|y|)} & x \cdot y < 0 \\ \\ x + y + xy & x,y < 0 \end{cases} \qquad (4)$$

where $x = CF(H,E_1)$, $y = CF(H,E_2)$, and $z = CF(H,E_1 \wedge E_2)$.

An inconsistency follows from the two relations above. From (4), it follows that $CF(H,E_1 \wedge E_2) = CF(H,E_2 \wedge E_1)$. That is, the combination of evidence is commutative. However, it can be shown that the definition of certainty factors, (3), prescribes non-commutative combination of evidence.

Recent work has focused on removing inconsistencies in the certainty factor model [19]. The consistent reformulation of the MYCIN certainty factor model falls into category 2 above; it can be shown that the certainty factor model is a *specialization* of probability in that assumptions of conditional independence are imposed by the methodology. For example, it can be shown that (4) is consistent with Bayes' theorem only if E_1 and E_2 are conditionally independent given H and its negation.

Although the certainty factor model is computationally efficient, many probabilists would feel the methodology was still unjustified because of its imposition of potentially invalid independence assumptions. They might seek a method whereby the tradeoff between computational efficiency and correctness can be controlled. Indeed, methods in which it is possible to selectively ignore dependencies that are not worth the computational effort to consider are currently being investigated [21].

VIII CONCLUSION

We have presented a logical mapping between several intuitive properties and the axioms of probability theory, and have given examples of how the mapping can be useful in identifying the fundamental differences between probability theory and non-probabilistic methodologies. We believe that the framework can help clarify discussions about alternative belief entailment schemes, whether they currently exist or result from future research. We recommend that investigators who seek a method for reasoning under uncertainty review the fundamental properties of measures of belief to gain an intuitive perspective on the nature of probability theory and on the relationship of non-probabilistic alternatives to probability.

ACKNOWLEDGEMENTS

We thank Greg Cooper, Arthur Dempster, Larry Fagan, Benjamin Grosof, Judea Pearl, Glenn Shafer, Peter Szolovits, Ted Shortliffe, Michael Wellman, and Lotfi Zadeh for useful discussions.

REFERENCES

1. Pearl, J. Fusion, propagation, and structuring in Bayesian networks. Presented at the Symposium on Complexity of Approximately Solved Problems, Columbia University, 1985

2. Zadeh, L.A. The role of fuzzy logic in the management of uncertainty in expert systems. *Fuzzy Sets and Systems* 11 (1983), 199-227.

3. Shafer, G.. *A Mathematical Theory of Evidence.* Princeton University Press, 1976.

4. Shortliffe, E. H. and Buchanan, B. G. A model of inexact reasoning in medicine. *Mathematical Biosciences 23* (1975), 351-379.

5. Gaines, B.R. Fuzzy and probability uncertainty logics. *Information and Control 38* (1978), 154-169.

6. Cheeseman, P. In defense of probability. Proceedings of the Ninth International Joint Conference on Artificial Intelligence, IJCAI-85, 1985.

7. Cohen, P. R. *Heuristic Reasoning About Uncertainty: An Artificial Intelligence Approach.* Ph.D. Th., Computer Science Department, Stanford University, Aug. 1983.

8. Cox, R. Probability, frequency and reasonable expectation. *American Journal of Physics 14,* 1 (January-February 1946), 1-13.

9. Jaynes, E.T. Plausible reasoning. Chapter in unpublished course notes

10. Tribus, M. What do we mean by rational? In *Rational Descriptions, Decisions and Designs,* Pergamon Press, New York, 1969.

11. Buchanan, B.G. *Logics of Scientific Discovery.* Ph.D. Th., University of Michigan, Ann Arbor, 1966.

12. Hanson, N.R.. *Patterns of Discovery.* Cambridge University Press, Cambridge, 1958.

13. Simon, H.A. Does scientific discovery have a logic? *Philosophy of Science 40* (1973), 471-480.

14. Aczel, J.. *Lectures on Functional Equations and Their Applications.* Academic Press, New York, 1966.

15. Dempster, A.P. Upper and lower probabilities induced by a multivalued mapping. *Ann. Math. Statistics 38* (1967), 325-339.

16. Zadeh, L.A. Fuzzy probabilities and their role in decision analysis. In *Proceedings of the Fourth MIT/ONR Workshop on Distributed Information and Decision Systems,* MIT, 1982.

17. Cheeseman, P. Probabilistic versus fuzzy reasoning. In *Uncertainty and Probability in Artificial Intelligence,* North-Holland, New York, 1986.

18. Shafer, G. Probability judgment in artificial intelligence. In *Uncertainty in Artificial Intelligence,* North-Holland, 1986.

19. Heckerman, D.E. Probabilistic Interpretations for MYCIN's Certainty Factors. In *Uncertainty in Artificial Intelligence,* North Holland, New York, 1986.

20. Horvitz, E. J., and Heckerman, D. E. The Inconsistent Use of Measures of Certainty in Artificial Intelligence Research. In *Uncertainty in Artificial Intelligence,* North Holland, New York, 1986.

21. Howard, R. A., Matheson, J. E. Influence Diagrams. In *Readings on the Principles and Applications of Decision Analysis,* Howard, R. A., Matheson, J. E., Eds., Strategic Decisions Group, Menlo Park, CA, 1981, ch. 37, , pp. 721-762.