A Symbolic Generalization of Probability Theory *

Adnan Y. Darwiche and **Matthew L. Ginsberg**

Computer Science Department
Stanford University
Stanford, CA 94305

Abstract

This paper demonstrates that it is possible to relax the commitment to numeric degrees of belief while retaining the desirable features of the Bayesian approach for representing and changing states of belief. We first present an abstract representation of states of belief and an associated notion of conditionalization that subsume their Bayesian counterparts. Next, we provide some symbolic and numeric instances of states of belief and their conditionalizations. Finally, we show that patterns of belief change that make Bayesianism so appealing do hold in our framework.

Introduction

Representing states of belief and modeling their dynamics is an important area of research in AI that has many interesting applications. A number of formalisms for this purpose have been suggested in the literature [Aleliunas, 1988; Bonissone, 1987; Dubois and Prade, 1988; Ginsberg, 1988; Pearl, 1988; Shenoy, 1989; Spohn, 1990] but Bayesian formalisms [Pearl, 1988] seem to be among the best we know. Here, a state of belief is represented by a probability function over some set of propositions, and Bayes conditionalization is used to change a state of belief upon acquiring certain evidence.

The success and increasing popularity of Bayesian formalisms result largely from two factors. First, their admission of non–binary degrees of belief makes them more convenient than classical logic formalisms, for example, which support true and false propositions only. Second, the associated notion of Bayes conditionalization gives rise to many desirable patterns of belief change [Polya, 1954; Pearl, 1988]. Among these patterns are Polya's five patterns of plausible inference: examining a consequence, a possible ground, a conflicting conjecture, several consequences in succession, and circumstantial evidence [Polya, 1954]. For example, the first of these patterns says "The verification of a consequence renders a conjecture more credible" while the second says "Our confidence in a conjecture can only diminish when a possible ground for the conjecture has been exploded."

A significant problem with probability calculus[1], however, is that it commits one to numeric degrees of belief. The reason why this commitment is problematic was clearly expressed by Jon Doyle [Doyle, 1990]:

> One difficulty is that while it is relatively easy to elicit tentative propositional rules from experts and from people in general, it is considerably harder to get the commitment to particular grades of certainty ... Worse still, individual informants frequently vary in their answers to a repeated question depending on the day of the week, their emotional state, the preceding questions, and other extraneous factors ... Reported experiments show the numbers do not actually mean exactly what they mean, for the performance of most systems remains constant under all sort of small ($< 30\%$) perturbations in the precise values used.

Nevertheless, AI practitioners continue to have mixed feelings about probability calculus and other numerical approaches to uncertainty:

> Understandably, expert system designers have difficulty justifying their use of the numerical judgements in face of these indications of psychological and pragmatic unreality. Unfortunately, they have had to stick to their guns, since no satisfactory alternative has been apparent. [Doyle, 1990]

It is therefore of significant interest to the AI community to have a calculus that (1) does not commit to numbers, (2) admits non–binary degrees of belief, and (3) supports patterns of belief change that make probability calculus so successful. But is this possible? This paper answers "Yes." In the following sec-

*An earlier version of this paper will appear in the Proceedings of the Tenth National Conference on Artificial Intelligence (AAAI–92).

[1]We use the terms "probability calculus," "probability theory," and "Bayesianism" interchangeably in this paper.

tions, we present a belief calculus that enjoys the above properties.[2]

Representing states of belief

A state of belief can be viewed as an attribution of degrees of support to propositions.[3] To formalize this intuition, however, we need to choose particular representations of propositions and degrees of support, and to constrain the mappings from propositions to degrees of support so that they correspond to coherent states of belief.

Propositions Our account of propositions is to identify them with sentences of a propositional language $\mathcal{L}$ with the usual connectives $\neg, \wedge, \vee, \equiv$, and $\supset$. We use **false** to denote any contradictory sentence, and **true** to denote any tautologous sentence, in $\mathcal{L}$. The symbols A, B, and C denote sentences in $\mathcal{L}$, and $\models$ denotes logical entailment.

Degrees of support A degree of support is an abstract quantity. It is neither strictly numeric nor strictly symbolic. Degrees of support can be integers, rationals, and even logical sentences. We view a degree of support as a primitive concept that derives its meaning from the operations and relations that are defined on degrees of support $\mathcal{S}$. The symbols a, b, and c denote degrees of support in $\mathcal{S}$.

States of belief A state of belief is a mapping Φ from a language $\mathcal{L}$ into degrees of support $\mathcal{S}$.[4] This definition, however, admits some incoherent states of belief. For example, if $\mathcal{S} = \{\text{true, false}\}$, we may have a state of belief that assigns false to a proposition and to its negation. We would like to exclude such states. And we will do this by identifying and formalizing a set of intuitive properties about coherent states of belief. The following are the properties we have identified:

(A0) Equivalent sentences have the same support in the same state of belief.

(A1) Contradictory sentences have the same support across all states of belief.

(A2) Tautologous sentences have the same support across all states of belief, which is different from the support for contradictory sentences.

(A3) The support for $A \vee B$ is a function of the supports for A and for B when A and B are logically disjoint.

(A4) If $A \models B \models C$ and A has the same support as C, then B has also the same support as A and C.

Formalizing the above properties constrains the degrees of support $\mathcal{S}$, and the mappings form $\mathcal{L}$ to $\mathcal{S}$, as shown by the following theorem.

Theorem 1 *Let $\Phi : \mathcal{L} \longrightarrow \mathcal{S}$ be a state of belief, A and B be sentences in $\mathcal{L}$. Properties (A0)–(A4) hold iff:*

1. *$\Phi(A) = \Phi(B)$ if $\models A \equiv B$.*
2. *There exists a partial function $\oplus : \mathcal{S} \times \mathcal{S} \longrightarrow \mathcal{S}$ such that:*[5]
   - *$\Phi(A \vee B) = \Phi(A) \oplus \Phi(B)$ if $\models \neg(A \wedge B)$, and*
   - *$a \oplus b = b \oplus a$, $(a \oplus b) \oplus c = a \oplus (b \oplus c)$, and if $(a \oplus b) \oplus c = a$ then $a \oplus b = a$.*
3. *There exists a unique support $\mathbf{0}$ in $\mathcal{S}$ such that:*
   - *$\Phi(\mathbf{false}) = \mathbf{0}$, and*
   - *for all a, $a \oplus \mathbf{0} = a$.*
4. *There exists a unique support $\mathbf{1} \neq \mathbf{0}$ in $\mathcal{S}$ such that:*
   - *$\Phi(\mathbf{true}) = \mathbf{1}$, and*
   - *for all a, there exists b that satisfies $a \oplus b = \mathbf{1}$.*

The function $\oplus$ is called support summation and $\langle \mathcal{S}, \oplus \rangle$ is called a partial support structure.

	$\langle \mathcal{S}, \oplus \rangle$	$\mathbf{0}$	$\mathbf{1}$
Possibility	$\langle \{0,1\}, \max \rangle$	0	1
Probability	$\langle [0,1], + \rangle$	0	1
Disbelief	$\langle \{0,1,\ldots,\infty\}, \min \rangle$	∞	0
Objection	$\langle \mathcal{O}, \wedge \rangle$	**true**	**false**

Table 1: Examples of partial support structures.

The full paper shows that the first three partial support structures in Table 1 induce states of belief that correspond to the following, respectively: classical logic, probability calculus, and nonmonotonic logic based on preferential models [Kraus *et al.*, 1990].

In probability calculus, we assess our support for a sentence by providing a number in the interval $[0,1]$. If we have complete confidence in a sentence, we give it a support of one; otherwise, we give it a support of less than one. Another way to assess support for a sentence is to explicate the reason we have doubts about it. For example, given that Tweety is a bird, we may have doubts about its flying ability because "Tweety is wingless." This intuition motivates a class of states of belief where degrees of support are sentences, called *objections*, in a propositional language $\mathcal{O}$.

The support summation function in objection–based states of belief is logical conjunction. That is, the objection to $A \vee B$ is equivalent to A's objection conjoined with B's objection. For example, considering Table 2, the objection to $A \supset B$ = "Bird implies flies"

[2] Proofs, omitted due to space limitations, can be found in the full version of this paper.

[3] We use the term "degree of support" as a generalization of the term "degree of belief." Support could be *for* or *against* the proposition to which it is attributed.

[4] We assume that degrees of support are useful. That is, for all a in $\mathcal{S}$, there is a state of belief that attributes a to some sentence in its domain.

[5] $a \oplus b$ is defined iff $a = \Phi(A)$ and $b = \Phi(B)$ for some Φ, A, B where $\models \neg(A \wedge B)$.

Sentence ($\mathcal{L}$)	*Objection* ($\mathcal{O}$)
Bird and flies.	Wingless.
Bird.	**false**.
Not bird.	Has feather.

Table 2: A partial objection–based state of belief.

can be computed by conjoining the objection to $A \wedge B$ = "Bird and flies" with the objection to $\neg A$ = "Not bird," which yields $\Phi(A \wedge B) \wedge \Phi(\neg A)$ = "Wingless and has feather."[6]

There is a close connection between objection–based states of belief and ATMS [Reiter and de Kleer, 1987], which rests on the following observation: the objection to a sentence can be viewed as an ATMS label for the negation of that sentence. For example, the objection to "Bird implies flies," "Wingless and has feather," can be viewed as a label for "Bird and does not fly."

Ordering degrees of support Degrees of support can be partially ordered using the support summation function. The intuition here is that the sum of two supports is at least as great as each of the summands.

Theorem 2 *Define support a to be "less than or equal to" support b (written $a \leq_\oplus b$) iff there is a support c satisfying $a \oplus c = b$. The relation $\leq_\oplus$ is a partial order on $\mathcal{S}$, and for all a in $\mathcal{S}$, $\mathbf{0} \leq_\oplus a \leq_\oplus \mathbf{1}$. $\leq_\oplus$ is called a* <u>*support order.*</u>

$\langle \mathcal{S}, \oplus \rangle$	$a \leq_\oplus b$
$\langle \{0,1\}, \max \rangle$	$a \leq b$
$\langle [0,1], + \rangle$	$a \leq b$
$\langle \{0,1,\ldots,\infty\}, \min \rangle$	$a \geq b$
$\langle \mathcal{O}, \wedge \rangle$	$b \models a$

Table 3: Examples of support orders.

A sentence that has support $\mathbf{0}$ will be called *rejected* and its negation will be called *accepted.* Note that if a sentence is accepted, it must have support $\mathbf{1}$, but the converse is not true. To consider an example, let degrees of support $\mathcal{S}$ be {possible, impossible}, and let support summation be defined as follows: $a \oplus b$ = possible unless $a = b$ = impossible. This makes $\mathbf{1}$ = possible. Moreover, a state of belief Φ may be such that Φ("Bird") = Φ("Not bird") = possible. Hence, a sentence and its negation may have support $\mathbf{1}$ and yet none of them may be accepted.

Changing states of belief

This section is mainly concerned with the following question: how should a state of belief change as a result of accepting a non–rejected sentence?

When we accept a sentence A in a state of belief Φ, we say that we have *conditionalized* Φ on A. Our goal in this section is to formalize this process.

Definition 1 *Let Φ be a state of belief with respect to $\langle \mathcal{S}, \mathcal{L}, \oplus \rangle$. If $A \in \mathcal{L}$ is not rejected by Φ, then a* <u>*conditionalization*</u> *of Φ on A (written Φ_A) is a state of belief, with respect to $\langle \mathcal{S}, \mathcal{L}, \oplus \rangle$, in which A is accepted.*

Given a state of belief Φ, there are many conditionalized states of belief Φ_A that satisfy the above definition. Some of these states correspond to plausible changes in a state of belief, but others do not. We would like to constrain conditionalization so that implausible belief changes are excluded. And we will do this by identifying and formalizing some intuitive properties of belief change. The following are the properties we have identified:

(A5) Accepting a non–rejected sentence retains all accepted sentences.

(A6) Accepting an accepted sentence leads to no change in a state of belief.

(A7) The support for A either increases or does not change after accepting $A \vee B$.

(A8) If A's support after accepting C equals its support after accepting $B \wedge C$, then B's support after accepting C equals its support after accepting $A \wedge C$.

(A9) If $A \vee B$ is equally supported by two states of belief, then accepting $A \vee B$ in each state does not introduce equality or order between the respective supports for A.

(A10) The support for A after accepting $A \vee B$ is a function of the initial supports for A and $A \vee B$.

Formalizing the above properties leads to a constructive definition of conditionalization as shown by the following theorem.

Theorem 3 *Assume (A0)–(A4), let Φ be a state of belief with respect to $\langle \mathcal{S}, \mathcal{L}, \oplus \rangle$, and let A, B be two sentences in $\mathcal{L}$ where A is not rejected by Φ. Properties (A5)–(A10) hold iff there exists a partial function $\oslash : \mathcal{S} \times \mathcal{S} \rightarrow \mathcal{S}$ such that:*[7]

- *$\Phi_A(B) = \Phi(A \wedge B) \oslash \Phi(A)$, and*
- *$\mathbf{0} \oslash a = \mathbf{0}$, $a \oslash \mathbf{1} = a$, $(a \oplus b) \oslash c = (a \oslash c) \oplus (b \oslash c)$, $a \oslash a = \mathbf{1}$, $a \oslash b \geq_\oplus a$, if $a \oslash c \leq_\oplus (=)\ b \oslash c$ then $a \leq_\oplus (=)\ b$, and if $a \oslash b = c \oslash d$ then $a \oslash c = b \oslash d$.*

The function $\oslash$ is called <u>*support scaling*</u> *and $\langle \mathcal{S}, \oplus, \oslash \rangle$ is called a* <u>*support structure.*</u>

The full paper shows that the first three scaling functions in Table 4 give rise to conditionalization rules that correspond to the following, respectively: augmenting conclusions in classical logic, Bayes conditionalization in probability calculus, and augmenting/retracting conclusions in nonmonotonic logic based on preferential models [Kraus *et al.*, 1990].

[6] Note that $A \supset B$ is equivalent to $(A \wedge B) \vee \neg A$.

[7] $a \oslash b$ is defined iff $a \leq_\oplus b$ and $b \neq \mathbf{0}$.

$\langle \mathcal{S}, \oplus \rangle$	$a \oslash b$
$\langle \{0,1\}, \max \rangle$	$\min(a,b)$
$\langle [0,1], + \rangle$	a/b
$\langle \{0,1,\ldots,\infty\}, \min \rangle$	$a - b$
$\langle \mathcal{O}, \wedge \rangle$	$\begin{cases} \textbf{true}, & \text{if } a \equiv \textbf{true}; \\ a \wedge \neg b, & \text{otherwise.} \end{cases}$

Table 4: Examples of support scaling.

Conditionalization of objection–based states of belief roughly states the following: the objection to B after accepting A is the initial objection to $A \wedge B$ minus the initial objection to A [Darwiche, 1992b]. To give an example, let us compute the objection to B = "Flies" conditioned on accepting A = "Bird." According to Theorem 3, this can be computed from the objection to $A \wedge B$ = "Bird and flies," and the objection to A = "Bird," which are given in Table 2. The desired objection, $\Phi_A(B)$, is then computed by $\Phi(A \wedge B) \wedge \neg\Phi(A)$ = "Wingless."

We conclude this section by noting that objection–based conditionalization is closely related to updating ATMS labels. A complete treatment of this connection, however, is not within the scope of this paper — the interested reader is referred to [Darwiche, 1992a].

Conditional and unconditional supports

For our framework to be useful in building artificial agents, the specification of a state of belief must be made intuitive enough so that a domain expert can naturally map his state of belief onto an artificial agent. This section discusses a function on degrees of support that helps achieve this goal.

A basic observation about human reasoning, claimed by Bayesian philosophers, is that it is more intuitive for people to specify their support for a sentence B (e.g., "The grass is wet") conditioned on accepting a relevant sentence A (e.g., "It rained") than to specify their unconditional support for B. It is therefore natural for domain experts to specify their states of belief by providing conditional supports. This is indeed the approach taken by most probabilistic representations where a domain expert provides statements of the form "$P(B|A) = p$," which reads as "If I accept A, then my probabilistic support for B becomes p."

One should note, however, that conditional supports are most useful when they can tell us something about unconditional supports. For example, conditional probabilities can be easily mapped into unconditional probabilities: $P(A \wedge B) = P(B|A)P(A)$. It is then important to ask whether the previous equality is an instance of a more general one that holds in our framework. This question is answered positively by the following theorem, which states that for every support structure there is a function on degrees of support that plays the same role as that played by numeric multiplication in probability theory.

Theorem 4 *Let $\langle \mathcal{S}, \oplus, \oslash \rangle$ be a support structure, Φ be a state of belief with respect to $\langle \mathcal{S}, \mathcal{L}, \oplus \rangle$, and A, B be two sentences in $\mathcal{L}$ where A is not rejected by Φ. Properties (A0)–(A10) imply the existence of a partial function $\otimes : \mathcal{S} \times \mathcal{S} \rightarrow \mathcal{S}$ such that:*[8]

- *$\Phi(A \wedge B) = \Phi_A(B) \otimes \Phi(A)$, and*
- *$(a \oslash b) \otimes b = (a \otimes b) \oslash b = a$, $\mathbf{0} \otimes a = \mathbf{0}$, $a \otimes \mathbf{1} = a$, $a \otimes b \leq_\oplus b$, if $a \leq_\oplus b$ then $a \otimes c \leq_\oplus b \otimes c$, $a \otimes b = b \otimes a$, and $(a \otimes b) \otimes c = a \otimes (b \otimes c)$.*

The function $\otimes$ is called <u>support unscaling.</u>

$\langle \mathcal{S}, \oplus \rangle$	$a \otimes b$
$\langle \{0,1\}, \max \rangle$	$\min(a,b)$
$\langle [0,1], + \rangle$	ab
$\langle \{0,1,\ldots,\infty\}, \min \rangle$	$a + b$
$\langle \mathcal{O}, \wedge \rangle$	$a \vee b$

Table 5: Examples of support unscaling.

The support unscaling function in objection–based states of belief is logical disjunction. We have previously computed the objection to B = "Flies" conditioned on accepting A = "Bird" to be $\Phi_A(B)$ = "Wingless." So let us now compute the objection to $A \wedge B$ = "Bird and flies." Theorem 4 tells us that we need the objection to A for this computation, which is given in Table 2. The desired objection, $\Phi(A \wedge B)$, is then computed by $\Phi_A(B) \vee \Phi(A)$ = "Wingless."

Patterns of plausible reasoning

The ultimate objective of many works in AI — most notably nonmonotonic logics — is to capture patterns of plausible reasoning in nonnumerical terms. George Polya (1887–1985) was one of the first mathematicians to attempt a formal characterization of qualitative human reasoning. Polya identified five main patterns of plausible reasoning in [Polya, 1954, Chapter XV] and demonstrated that they can be formalized using probability theory. Pearl highlighted these patterns in his recent book [Pearl, 1988] and took them — along with other patterns such as nonmonotonicity, abduction, explaining–away and the hypothetical middle [Pearl, 1988, Page 19] — as evidence for the indispensability of probability theory in formalizing plausible reasoning. In his own words:

> We take for granted that probability calculus is unique in the way it handles context–dependent information and that no competing calculus exists that closely covers so many aspects of plausible reasoning [Pearl, 1988, Page 20].

[8] $a \otimes b$ is defined iff there is c satisfying $c \leq_\oplus b \neq \mathbf{0}$ and $c \oslash b = a$.

This section shows that four of Polya's patterns of plausible reasoning hold in our framework. First, however, we need to formally define certain terms that Polya used in stating his patterns. To "verify" or "prove" a proposition is to accept it. To "explode" a proposition is to reject it. And, the "credibility of," "confidence in," and "belief in" a proposition are all equivalent terms.

Definition 2
A state of belief Φ supports A less than or equally to B iff $\Phi(A) \leq_\oplus \Phi(B)$.

Given Theorem 2, it should be clear that the relation less–than–or–equally–supported–to is a partial ordering.

Definition 3 *A state Φ believes A less than or equally to B iff Φ supports A less than or equally to B, and supports $\neg B$ less than or equally to $\neg A$.*

The second part of the above definition may seem redundant, but it generally is not. Table 6 provides a counterexample.

Given		*Computed*	
$\Phi(\texttt{bird} \wedge \texttt{fly})$	1	$\Phi(\texttt{bird})$	1
$\Phi(\texttt{bird} \wedge \neg\texttt{fly})$	1	$\Phi(\neg\texttt{bird})$	0
$\Phi(\neg\texttt{bird} \wedge \texttt{fly})$	0	$\Phi(\texttt{fly})$	1
$\Phi(\neg\texttt{bird} \wedge \neg\texttt{fly})$	0	$\Phi(\neg\texttt{fly})$	1

Table 6: A state of belief with respect to the structure $\langle\{0,1\}, \max\rangle$, where $0 \leq_{\max} 1$. `bird` is equally supported to `fly` although it is more believed.

The reason for this asymmetry with probability calculus is that, in general, the support for a proposition does not determine the support for its negation, as is the case in probability calculus.

Theorem 5 *Less–than–or–equally–believed–to is a partial ordering.*

We are now ready to state and prove Polya's patterns.

- *Examining a consequence:*

  The verification of a consequence renders a conjecture more credible.[Polya, 1954, Page 120]

Theorem 6 *If $A \supset B$ is accepted, and B is not rejected, by a state of belief Φ, then $\Phi_B(A) >_\oplus \Phi(A)$ unless $\Phi(A) = \mathbf{0}$ or $\Phi(B) = \mathbf{1}$.*

- *Examining several consequences in succession:*

  The verification of a new consequence enhances our confidence in the conjecture, unless the new consequence is implied by formerly verified consequences.[Polya, 1954, Page 125]

Theorem 7 *If $A \supset C_1, \ldots, A \supset C_n$ are accepted, and $C_1 \wedge \ldots \wedge C_n$ is not rejected, by a state of belief Φ, then $\Phi_{C_1 \wedge \ldots \wedge C_n}(A) >_\oplus \Phi_{C_1 \wedge \ldots \wedge C_{n-1}}(A)$ unless $\Phi_{C_1 \wedge \ldots \wedge C_{n-1}}(C_n) = \mathbf{1}$.*

The patterns of examining a possible ground and examining a conflicting conjecture are omitted for space limitations and can be found in the full version of this paper.

Discussion and related work

- An important attraction of probabilistic states of belief is that they can be specified using Probabilistic Causal Networks (PCNs) [Pearl, 1988]. PCNs are easy to construct and serve as models for computing unconditional and conditional probabilities. There are parallel constructs for specifying abstract states of belief, called Abstract Causal Networks (ACNs) [Darwiche, 1992a; Darwiche, 1992b]. ACNs are also easy to construct and serve as models for computing unconditional and conditional degrees of support.
- The four belief calculi that we presented so far are not the only instances of our framework. Table 7, for example, depicts two more calculi.

	Improbability	Consequence
$\mathcal{S}$	$[0,1]$	A propositional language
$a \oplus b$	$a + b - 1$	$a \vee b$
$a \oslash b$	$(a-b)/(1-b)$	$\begin{cases} \textbf{false}, & \text{if } a \equiv \textbf{false}; \\ a \vee \neg b, & \text{otherwise.} \end{cases}$
$\mathbf{0}$	1	**false**
$\mathbf{1}$	0	**true**
$a \leq_\oplus b$	$a \geq b$	$a \models b$
$a \otimes b$	$a + b - ab$	$a \wedge b$
$\Phi(A)$	Improb. of A	Consequence of A

Table 7: Improbability and Consequence calculi.

If a domain expert is not satisfied with the calculi proposed in this paper, all he has to do is the following: choose a set of supports $\mathcal{S}$ that he feels more comfortable with and accept properties (A0)–(A10) with respect to this choice. The results of this paper show that there must exist a support structure, $\langle\mathcal{S}, \oplus, \oslash\rangle$, which gives rise to a new belief calculus that shares with probability calculus its desirable properties.
- It is typical of multivalued logics [Bonissone, 1987; Ginsberg, 1988] and generalizations of probability calculus [Aleliunas, 1988] to assume at least one of the following axioms:

$\Phi(A \wedge B)$ is a function of $\Phi(A)$ and $\Phi(B)$, or

$\Phi(\neg A)$ is a function of $\Phi(A)$.

None of the six calculi presented in this paper satisfy the first axiom, and only probability and improbability calculi satisfy the second axiom.

■ One can define a notion of qualitative conditional influence that generalizes the probabilistic notion defined for Qualitative Probabilistic Networks [Wellman, 1988]. In a state of belief Φ, we say that A positively influences B given C if and only if

$$\Phi_{\neg A \wedge C \wedge D}(B) \leq_{\oplus} \Phi_{A \wedge C \wedge D}(B),$$

for all D where $\neg A \wedge C \wedge D$ and $A \wedge C \wedge D$ are not rejected by Φ. Negative and zero influences are similar.
■ One criticism of the work presented here may be that its conception of a state of belief is restrictive because of Property (A3). As a result of this property, states of belief such as those represented by Dempster's basic probability assignments [Shafer, 1976] do not fit in our framework.
■ It is not clear whether Polya's pattern of circumstantial evidence holds in our framework. This pattern says that "If a certain circumstance is more credible with a certain conjecture than without it, the proof of that circumstance can only enhance the credibility of that conjecture."[Polya, 1954, Page 127].
■ Some important questions remain unanswered about our framework. For example, what additional properties of states of belief and belief change would commit us to Bayesianism? Moreover, what additional properties of belief change would force the uniqueness of support scaling, thus, reducing a support structure into a pair $\langle \mathcal{S}, \oplus \rangle$? Finally, is their an abstract decision theory that subsumes probabilistic decision theory, in the same way that our framework subsumes Bayesianism?

Conclusion

We have presented an abstract framework for representing and changing states of belief, which subsumes Bayesianism. At the heart of the framework is a mathematical structure, $\langle \mathcal{S}, \oplus, \oslash \rangle$, called a support structure, which contains all the information needed to represent and conditionalize states of belief. We have also presented symbolic and numeric instances of support structures, and have shown that our framework supports some patterns of plausible reasoning that have been considered unique to numeric formalisms.

Acknowledgements

The first author has benefited greatly from comments of Ahmad Darwiche, Don Geddis, Adam Grove, Jinan Hussain, Alon Levy, H. Scott Roy, Ross Shachter, Yoav Shoham, and the PRINCIPIA group at Stanford University. This work has been supported by the Air Force Office of Scientific Research under grant number 90-0363, by NSF under grant number IRI89-12188, and by DARPA/Rome Labs under grant number F30602-91-C-0036.

References

Aleliunas, Romas 1988. A new normative theory of probability logic. In *Proceedings of the Canadian Artificial Intelligence Conference.* Morgan Kaufmann Publishers, Inc., San Mateo, California. 67–74.

Bonissone, Piero 1987. Summarizing and propagating uncertain information with triangular norms. *International Journal of Approximate Reasoning* 1:71–101.

Darwiche, Adnan Y. 1992a. Forthcoming Ph.D. Thesis.

Darwiche, Adnan Y. 1992b. Objection–based causal networks. In *Proceedings of the Eighth Conference on Uncertainty in Artificial Intelligence (UAI).* 67–73.

Doyle, Jon 1990. Methodological simplicity in expert system construction: The case of judgements and reasoned assumptions. In Shafer, Glenn and Pearl, Judea, editors 1990, *Readings in Uncertain Reasoning.* Morgan Kaufmann Publishers, Inc., San Mateo, California. 689–693.

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

Ginsberg, Matthew L. 1988. Multivalued logics: a uniform approach for reasoning in artificial intelligence. *Computational Intelligence* 4:265–316.

Kraus, S.; Lehmann, D.; and Magidor, M. 1990. Nonmonotonic reasoning, preferential models and cumulative logics. *Artificial Intelligence* 24(1-2):167–207.

Pearl, Judea 1988. *Probabilistic Reasoning in Intelligent Systems: Networks of Plausible Inference.* Morgan Kaufmann Publishers, Inc., San Mateo, California.

Polya, George 1954. *Patterns of Plausible Inference.* Princeton University Press, Princeton, NJ.

Reiter, Ray and de Kleer, Johan 1987. Foundations of assumption-based truth maintenance systems: Preliminary report. In *Proceedings of AAAI.* AAAI. 183–188.

Shafer, Glenn 1976. *A Mathematical Theory of Evidence.* Princeton University Press, Princeton, NJ.

Shenoy, Parkash P. 1989. A valuation–based language for expert systems. *International Journal of Approximate Reasoning* 5(3):383–411.

Spohn, Wolfgang 1990. A general non–probabilistic theory of inductive reasoning. In Kanal, L.; Shachter, R.; Levitt, T.; and Lemmer, J., editors 1990, *Uncertainty in Artificial Intelligence 4.* Elsevier Science Publishers. 149–158.

Wellman, Michael P. 1988. Qualitative probabilistic networks for planning under uncertainty. In Lemmer, J. and Kanal, L., editors 1988, *Uncertainty in Artificial Intelligence*, volume 2. Morgan Kaufmann Publishers, Inc., San Mateo, California.