

# Agent-oriented epistemic reasoning: subjective conditions of knowledge and belief<sup>1</sup>

Daniel G. Schwartz<sup>2</sup>

*Department of Computer Science, Florida State University, Tallahassee, Florida 32306-4530*

## Abstract

This paper introduces a formal system  $\Sigma$  of *subjective epistemic reasoning* that encodes a method of reasoning with conditions of *knowledge* and *belief*. The conditions are subjective in that they are taken from the perspective of an agent’s perception of his own state of knowledge or belief with respect to his observable world. Belief is measured along a series of linguistic degrees, e.g., *strongly believes*, *fairly confidently believes*, *somewhat disbelieves*, etc., and *knowledge* is taken as unequivocal belief. The system employs a novel, dual-leveled language that follows fuzzy logic by interpreting the logical *or* and *and* as the arithmetical *max* and *min*. Numerous properties of  $\Sigma$ , illustrating its intuitive appeal for the intended purpose, are derived.

*Keywords:* Agents, Agent-oriented reasoning, Approximate reasoning, Belief, Dual-leveled semantics, Epistemic logic, Fuzzy logic, Knowledge, Knowledge representation.

## 1 Introduction

A central aim of the field of Artificial Intelligence (AI) is to create computational agents that replicate the activities of the human mind. A major effort in this regard has sought to devise models of natural reasoning processes wherein a knowledge base, representing an agent’s view of its operational environment, changes and evolves over time. This emerged first as the study of nonmonotonic reasoning and subsequently as the study of belief revision. These investigations in turn evoked a special concern with “knowledge” and “belief” and the task of devising logical formalisms that capture everyday human reasoning with these ideas. In due course, the chapter from Philosophy that deals with these issues, known as *epistemic logic*, became folded into AI as part of the ongoing developments.

This paper briefly reviews some of the foundational works in this area, with the aim of distinguishing first between knowledge and belief, and then between objective and subjective knowledge. In this context it argues that, whereas most efforts to formalize our reasoning about knowledge have concerned the objective variety, from an agent-oriented perspective the subjective variety is actually more pertinent.

It next introduces the syntax and semantics for a logic  $\Sigma$  of *subjective epistemic reasoning* that is aimed at capturing the associated notions. This includes linguistic degrees of belief such as *strongly believes*, *fairly confidently believes*, *somewhat disbelieves*, etc., as well as *knowledge*, where this is taken as conviction or unequivocal belief. It is shown how  $\Sigma$  achieves the desired end by delineating some key formulas of its language of that are validated by the semantics. Taken together, these results show that  $\Sigma$  bears considerable intuitive appeal for the intended applications.

---

<sup>1</sup> This work was supported in part by the US Army Research Office, grant number DAAD19-01-1-0502.

<sup>2</sup> Email: schwartz@cs.fsu.edu

## 2 Knowledge and Belief

The classical view of the relationship between knowledge and belief, attributed to Plato’s “Theaetetus” [cf. 22], is that

“knowledge is justified true belief”

This view was challenged in 1963 by Gettier [11] with two examples showing how it is possible for someone (an agent) to be justified in believing a proposition that is true, but nonetheless cannot be said to *know* the proposition is true inasmuch as the justifications for his belief are predicated on assumptions that are false. In effect, Gettier’s argument is that one cannot be said to *know* that a proposition is true without also knowing *why* it is true.

This sparked a series of investigations by others, mostly focusing on the extent to which justified beliefs may be regarded as true. The next 15 years of such studies were summarized by Lenzen [22], who affirmed Gettier’s conclusions and took these to imply that the concept of knowledge cannot be reduced to other epistemic notions. This view, in turn, has been challenged by Voorbraak [37] in the context of a detailed logical formulation based on Kripke-style possible-worlds semantics. A chapter of [37] argues that, in order to be knowledge, a belief need only be justified by the *same reasons* for which it is true. Thus, in this sense, the classical view can be reaffirmed.

Whether one accepts or reject the classical view, however, no one seems to question whether a proposition, in order to be knowledge, must at least be true. Contrapositively, it seems generally agreed that one cannot be said to “know” something that is false. Empirical knowledge, in particular, necessarily always conforms to the actual state of the empirical world.

This position was taken in the well-known treatise by Hintikka [19]. It has subsequently appeared throughout investigations along these lines within the field of AI. Works of this genre include numerous publications by Konolige, Fagin, Halpern, and others [cf. 4, 5, 6, 14, 15, 16, 20, 23, 26]. As with Voorbraak, these mostly employ the familiar possible-worlds semantics and distinguish between the two notions in terms of what worlds the agent either knows or believes to be possible. And most have not attempted to exploit the classical view discussed above in order to combine the two notions in a unified system, but rather deal only with each of knowledge and belief separately. An exception is the work by Kraus and Lehman [21], which accounts for both notions in a single formalism and additionally introduces a temporal component to express the persistence of beliefs in time.

With few exceptions, these these systems explicitly express the view that a proposition cannot be regarded as knowledge unless it is true by means of the axiom

$$\mathcal{K}P \rightarrow P$$

The fact that so many of the proposed formulations include this axiom emphasizes its widespread acceptance.

This of course is entirely natural as long as one is discussing an agent’s capacity for knowledge of objective truths. From an agent-oriented standpoint, however, it may be questioned whether this is the correct approach. For, in practice, there is never complete certainty that an agent’s view of the empirical world is true. This is especially the

case with computational agents, whose knowledge bases are always, except in very trivial environments, only partial descriptions of the agent’s operational domain.

Indeed even in those cases where a correspondence between perceptions and actual facts seem inviolate, the agent’s comprehension of his external state of affairs is never wholly complete. There is always the possibility that extenuating circumstances can render his perceptions invalid. For example, if John observes that Mary is wearing a red hat, and this is corroborated by John’s friend Bob, then one is inclined to affirm that Mary’s hat is in fact red and that the perceptions are correct. This does not rule out the (admittedly remote) possibility, however, that both John and Bob were beset with temporary color blindness, and that Mary’s hat was in fact green. Hence, from the agent’s vantage point, his personal convictions of knowledge regarding the external world are, at best, only beliefs, albeit ones possibly held without equivocation.

A view somewhat in this vein was expressed by Moses and Shoham [26] in a formulation of belief as defeasible knowledge. Here the idea is that an agent can be held to “know” something until such time as one discovers countervailing information, after which the subject item becomes recognized as merely a belief. Since this defines belief in terms of knowledge, it appears to effectively reverse the classical relation expressed above. It must be considered, however, whether defeasible knowledge is in fact knowledge in the earlier sense. For if a conviction of knowledge can be overturned at a later time, then the presumed “knowledge” was not originally a complete and correct depiction of the external world. Hence it always was at best a belief, even though a belief held with firm conviction.

The argument, then, is that an agent’s assertion that he knows a proposition  $P$  is more exactly an assertion that he is convinced that  $P$  is true, and it does not rule out the possibility that his conviction is in error. Thus an agent’s assertions of knowledge are at best “justified beliefs.” In this respect, we return to a semblance of the classical view, but with an important exception: we do not claim that these beliefs need be “true.”

This opens the door for nonmonotonicity in that, if an agent is convinced that he knows something to be true, then he can be expected to act accordingly until such time as he may discover evidence to the contrary. When this occurs then he will need to reassess his position and reclassify the proposition as a formerly held belief.

Now when one separates the agent’s convictions from the question of whether these correctly correspond to the agent’s external world, one is then dealing only with *subjective* conditions of knowledge and belief. In other words, in the subjective realm of the agent’s own thinking processes, the agent can be convinced that a certain proposition is true, even when it is not.

This is important from an agent-oriented perspective because whether the agent’s convictions of knowledge are in fact true does not matter in terms of his own behavior; it matters only that he is acting *as if* they are true. Note also that this does not exclude beliefs that are held with absolute conviction, but which the agent nonetheless refers to as “belief” rather than “knowledge” (e.g., a belief in God). Here it is merely a matter of terminology, since an agent having such strong convictions will conduct his actions as if these were in fact items of true knowledge. On the other hand, if the agent acknowledges that the basis of his action is at best a belief, and not necessarily confirmed knowledge, then perhaps his actions will be undertaken with some trepidation, i.e, allowing for the possibility that these actions may be in error.

For purposes of replicating natural human reasoning on a computer, it is therefore of interest to develop models that capture this sort of reasoning in a well-defined formalism. The study of formalisms that deal with the notion of belief independently of the notion of knowledge is known as *doxastic logic*. The Lenzen paper [22] additionally surveys work in this area as of its publication date and introduces a new formulation where belief is defined in terms of probabilities. Works along these lines within AI include that of Konolige [20] and Fagin and Halpern [4, 5]. The former defines belief in terms of some notions of “belief set” and “deduction structure,” somewhat along the same lines as the derivation path and dynamic reasoning systems described in the following. The latter use an adaptation of the familiar possible-worlds semantics.

In addition one must mention the “autoepistemic logic” devised by Moore [24, 25]. This constitutes an early effort to formalize an agent’s reasoning about his own knowledge and beliefs. Moore’s logic is obtained by augmenting the classical predicate logic with a modal operator  $L$ , standing for the (subjective) *knows* or *believes*. That system was shown to capture several key aspects of the everyday reasoning with these ideas.

Another approach, somewhat along the same lines, has been proposed by Thijsse [35, 36]. This employs a “dual-leveled” formalism wherein the partial logic comprises an “internal” logic whose propositions are built up in the usual way from propositional variables using the connectives for negation and conjunction. These propositions are then modified with belief operators to become formulas of an “external” logic that obeys the laws of classical logic. As will be seen, this bears similarity with the approach taken here.

A well-known alternative treatment of belief that should be mentioned is Dempster-Shafer theory [33]. This has been studied by Halpern and Fagin [14] in comparison with still another interpretation based on generalized probabilities. Dempster-Shafer theory is concerned with delineating correct rules for combining evidence, where this evidence supports beliefs about the empirical world. Thus it seeks to specify how people *should* reason about their beliefs, not how they actually do.

The literature is nowadays replete with writings related to agent-oriented systems. In particular is the version of agent-oriented programming described by Shoham [34]. This bears similarity with the notion of a “dynamic reasoning system” referred to in the following section 4. An interesting collection of essays addressing issues in the same genre as this paper is the volume by Ford, Glymour, and Hayes [8].

### 3 The system $\Sigma$ of subjective epistemic reasoning

The present system  $\Sigma$  has much the same aim as the system of “rational (introspective belief” devised by Voorbraak [37], and it might have been similarly named. The term “epistemic” has been adopted here, however, to emphasize that the system is intended as a model of the agent’s awareness of its own internal conditions of knowledge and belief. In other words, when a formula of the form  $\mathcal{K}P$  is posited in the system, this is taken as representing the agent’s affirming that  $P$  is being held as an item of it’s knowledge, and when a formula of the form  $\mathcal{B}P$  is posited, this is taken as representing the agent’s affirming that  $P$  is being held as a belief. Moreover, when deductions are made from the agent’s current set of knowledge and beliefs, these are viewed as conscious acts made by the agent, so also are items of which the agent is aware. Taking awareness as synonymous with conscious knowledge, this warrants the system’s being referred to as epistemic.

Whether  $\Sigma$  correctly captures natural reasoning with linguistic expressions of knowledge and belief is of course difficult to assess, but the system clearly has intuitive appeal. This is borne out by the collection of propositions it validates. In particular, these include an interpretation of the logical *or* and *and* in terms of the arithmetical *max* and *min*, in the context of the given set of linguistic belief degrees. To illustrate, suppose an agent *fairly confidently believe* that the economy is healthy, and it *strongly disbelieves* that there can be any useful number system in which  $1 = 0$ . Then the agent has a natural tendency to determine its degree of belief that at least one of these assertions holds as being the greater of the two degrees, i.e., it *fairly confidently believe* that either one or the other is true. Similarly, given the same conditions and degrees of belief, the agent is inclined to determine its degree of belief that both hold simultaneously as being the greater of two degrees, i.e., it *strongly disbelieves* that both can simultaneously be true. In what follows it will be seen that  $\Sigma$  provides a formal language suitable for expressing such propositions and a semantics that validates them. The system will also be seen to bear other intuitively reasonable properties, particularly the interrelationships between its various logical connectives and belief modifiers.

The core of  $\Sigma$  is patterned after a logic developed previously for reasoning with graduated linguistic expressions of likelihood [31, 32]. This also employed the dual-leveled style of formalism presented in [30]. These logics owe an obvious debt to the literature on fuzzy logic. First, their use of the arithmetic *max* and *min* to define the logical *or* and *and* derives from Zadeh’s first paper on the subject [38]. Second, these systems borrow Zadeh’s notion of a “linguistic variable” [cf. 39]. Specifically, *Belief* is here treated as a variable of this type. Last, these works have been inspired by the same aims as expressed by Zadeh to create methods for “computing with words” [40, 41] and, more recently, to devise a “computational theory of perceptions” [42]. While my own works have adopted a more symbolic approach—in contrast with Zadeh’s emphasis on semantic inference—they are nonetheless directed toward the same goals.

### 3.1 Formal language

Let us adopt the following notations:

$\mathcal{B}_4$	for	<i>unequivocally believes</i>
$\mathcal{B}_3$	for	<i>strongly believes</i>
$\mathcal{B}_2$	for	<i>fairly confidently believes</i>
$\mathcal{B}_1$	for	<i>somewhat believes</i>
$\mathcal{B}_0$	for	<i>neither believes nor disbelieves</i>
$\mathcal{B}_{-1}$	for	<i>somewhat disbelieves</i>
$\mathcal{B}_{-2}$	for	<i>fairly confidently disbelieves</i>
$\mathcal{B}_{-3}$	for	<i>strongly disbelieves</i>
$\mathcal{B}_{-4}$	for	<i>unequivocally disbelieves</i>

These are taken to express various grades of strength in the agent’s disposition towards believing in a given proposition. In particular,  $\mathcal{B}_0$  expresses uncertainty, either in the sense that the agent has information to both support and refute the proposition in question, or in the sense that the agent suffers from a complete lack of information regarding the truth

of falsehood of the proposition. Distinguishing between these two varieties of uncertainty requires a richer system than the one being considered here.

Both the number of these notations and their interpretations are, of course, to a large extent arbitrary. In order for the following formalism to maintain its coherence, it is only necessary that the subscripting of the qualifiers be symmetrical about 0. Thus the above selection should be viewed as being merely for purposes of illustrating a general methodology. The semantics (Section 3), however, implicitly also requires that  $\mathcal{B}_0$  represents a position of complete equivocation (uncertainty) and that the topmost and bottommost qualifiers represent positions of complete conviction (certainty).

The language for  $\Sigma$  is defined as follows. As *symbols* we shall have: the *propositional variables*  $p_1, p_2, \dots$ , the *logical connectives*  $\neg, \vee, \dot{\neg}, \dot{\vee}$ , the left and right *parentheses* “(” and “)”, and the *belief operators*  $\mathcal{B}_4, \dots, \mathcal{B}_{-4}$ . The *formulas* for  $\Sigma$  will be the combined members of the following:

$$\begin{aligned} F_1 &= \{P \mid P \text{ is } p_i, i = 0, 1, \dots\} \\ F_2 &= F_1 \cup \{\neg P, (P \vee Q) \mid P, Q \in F_1 \cup F_2\}^3 \\ F_3 &= \{\mathcal{B}_i P \mid P \in F_2, i = -4, \dots, 4\} \\ F_4 &= F_3 \cup \{\dot{\neg} P, (P \dot{\vee} Q) \mid P, Q \in F_3 \cup F_4\} \end{aligned}$$

In the section below, the formulas in  $F_2$  are interpreted as comprising a multivalent logic, while the belief formulas in  $F_4$  are interpreted as obeying the rules of the classical bivalent logic.

Further connectives can be introduced in the usual manner as abbreviations:

$$\begin{aligned} (P \wedge Q) &\text{ for } \neg(\neg P \vee \neg Q) \\ (P \dot{\wedge} Q) &\text{ for } \dot{\neg}(\dot{\neg} P \dot{\vee} \dot{\neg} Q) \\ (P \dot{\rightarrow} Q) &\text{ for } (\dot{\neg} P \dot{\vee} Q) \\ (P \dot{\leftrightarrow} Q) &\text{ for } ((P \dot{\rightarrow} Q) \dot{\wedge} (Q \dot{\rightarrow} P)) \end{aligned}$$

In addition, one can introduce connectives AL and AM, standing for *at least* and *at most* by

$$\begin{aligned} AL\mathcal{B}_i P &\text{ for } \mathcal{B}_i P \dot{\vee} \dots \dot{\vee} \mathcal{B}_4 P \\ AM\mathcal{B}_i P &\text{ for } \mathcal{B}_{-4} P \dot{\vee} \dots \dot{\vee} \mathcal{B}_i P \end{aligned}$$

These definitions capture the usual meanings of the respective terms in application to the given belief operators.

The operator  $\mathcal{K}$  for *knows* is defined by taking

$$\mathcal{K}P \text{ for } \mathcal{B}_4 P$$

This is in keeping with the foregoing discussion inasmuch as it defines knowledge as unequivocal belief (conviction of truth). This distinguishes it from the  $\mathcal{K}$  appearing in previous works.

---

<sup>3</sup> This notation abbreviates the usual inductive definition, in this case the smallest class of formulas containing  $F_1$  together with all formulas that can be built up from formulas in  $F_1$  in the two prescribed ways.

It is natural to introduce unsubscripted letters  $\mathcal{B}$ ,  $\mathcal{D}$ , and  $\mathcal{U}$  for *believes*, *disbelieves*, and *is uncertain about* by

$$\begin{aligned} \mathcal{B}P & \text{ for } \mathcal{B}_1P \dot{\vee} \mathcal{B}_2P \dot{\vee} \mathcal{B}_3P \dot{\vee} \mathcal{B}_4P \\ \mathcal{D}P & \text{ for } \mathcal{B}_{-1}P \dot{\vee} \mathcal{B}_{-2}P \dot{\vee} \mathcal{B}_{-3}P \dot{\vee} \mathcal{B}_{-4}P \\ \mathcal{U}P & \text{ for } \mathcal{B}_0P \end{aligned}$$

Parentheses will be dropped when the intended grouping is clear, based on the usual precedence rules; associativity is assumed to be to the right. Formulas without belief modifiers are *first-* or *lower-level* formulas, and those with belief modifiers are *second-* or *upper-level*. Examples of formulas expressing meaningful statements appear in the sections below.

One may note that the present formalism explicitly rules out “mixed” formulas such as

$$P \dot{\rightarrow} \mathcal{B}_3P$$

The reason for this is that to assert a proposition  $P$  in any logical formalism is normally taken to posit  $P$  as being “true,” and this implicitly amounts to positing objective truth. Since the present system is aimed to model only assertions of subjective truth, such unqualified assertions are therefore not appropriate. One does have a subjective analog of  $P$  in  $\mathcal{K}P$ , however, and formulas of the form

$$\mathcal{K}P \dot{\rightarrow} \mathcal{B}_3P$$

are included in the above language. Such formulas will not be validated by the semantics for  $\Sigma$ , however, since here the various grades of belief are regarded as semantically distinct.

The language of  $\Sigma$  also rules out nested modification, expressing such propositions as that the agent “knows that it knows,” or “believes that it believes,” or “knows that it believes,” and so on. Accordingly this leaves out the “introspective” axioms considered in prior works. This omission is intentional, simply to keep the present semantics from becoming too complex. Such may be considered, if this is deemed useful, however, in a future work.

### 3.2 Semantics

As was mentioned, *Belief* is here treated as a *linguistic variable* in the sense of [39]. Specifically, this consists of (i) a set of primitive linguistic terms (propositional variables, connectives, belief modifiers, etc.), (ii) a grammar that specifies how to build up more complex expressions by using various rules of combination, and (iii) a mapping that associates each such expression with a well-defined meaning in some interpretation domain. Items (i) and (ii) have been provided above. Item (iii) is given by the definitions below.

The semantics for  $\Sigma$  consists of some *belief mappings*  $\beta$ , each of which associates each lower-level formula with a number in  $[0, 1]$ , and some *truth valuations*  $v$ , each of which associates each upper-level formula with a *truth value*,  $T$  or  $F$ . As will be seen, every valuation  $v$  is defined in terms of some given mapping  $\beta$ , and is uniquely determined by that  $\beta$ .

A *belief mapping* is a function  $\beta : \mathbb{F}_2 \rightarrow [0, 1]$  satisfying, for any propositional variable  $P$ ,

$$\beta(P) \in [0, 1]$$

and for any lower-level  $P$  and  $Q$ ,

$$\beta(\neg P) = 1 - \beta(P)$$

$$\beta(P \vee Q) = \max[\beta(P), \beta(Q)]$$

From these definitions we have

$$\begin{aligned} \beta(P \wedge Q) &= \beta(\neg(\neg P \vee \neg Q)) \\ &= 1 - \max[1 - \beta(P), 1 - \beta(Q)] \\ &= \min[\beta(P), \beta(Q)] \end{aligned}$$

For the *truth valuations*  $v$ , first let

$$\begin{aligned} \iota_4 &= [1, 1] \quad (\text{singleton } 1) \\ \iota_3 &= [\frac{6}{7}, 1) \\ \iota_2 &= [\frac{5}{7}, \frac{6}{7}) \\ \iota_1 &= [\frac{4}{7}, \frac{5}{7}) \\ \iota_0 &= (\frac{3}{7}, \frac{4}{7}) \\ \iota_{-1} &= (\frac{2}{7}, \frac{3}{7}] \\ \iota_{-2} &= (\frac{1}{7}, \frac{2}{7}] \\ \iota_{-3} &= (0, \frac{1}{7}] \\ \iota_{-4} &= [0, 0] \quad (\text{singleton } 0) \end{aligned}$$

Then, given any belief mapping  $\beta$ , define  $v : F_4 \rightarrow [T, F]$  by, for all  $i = -4, \dots, 4$ ,

$$\begin{aligned} v(\mathcal{B}_i P) &= T \text{ iff } \beta(P) \in \iota_i \\ v(\dot{\neg} P) &= T \text{ iff } v(P) = F \\ v(P \dot{\vee} Q) &= T \text{ iff either } v(P) = T \text{ or } v(Q) = T \end{aligned}$$

Given these definitions, it is easily established that

$$\begin{aligned} v(P \wedge Q) &= T \text{ iff } v(P) = T \text{ and } v(Q) = T \\ v(P \dot{\neg} Q) &= T \text{ iff } v(P) = F \text{ or } v(Q) = T \\ v(P \dot{\leftrightarrow} Q) &= T \text{ iff } v(P \dot{\neg} Q) = T \text{ and } v(Q \dot{\neg} P) = T \end{aligned}$$

Thus all the upper-level logical connectives behave in the normal ways. It may also be verified that

$$\begin{aligned} v(\text{AL}\mathcal{B}_i P) &= T \text{ iff } \beta(P) \in \iota_j \text{ for some } j \geq i \\ v(\text{AM}\mathcal{B}_i P) &= T \text{ iff } \beta(P) \in \iota_j \text{ for some } j \leq i \\ v(\mathcal{B}P) &= T \text{ iff } v(\mathcal{B}_1 P) = T \text{ or } \dots \text{ or } v(\mathcal{B}_4 P) = T \\ v(\mathcal{D}P) &= T \text{ iff } v(\mathcal{B}_{-1} P) = T \text{ or } \dots \text{ or } v(\mathcal{B}_{-4} P) = T \end{aligned}$$

An upper-level formula  $P$  will be a *tautology* with respect to this semantics, if  $v(P) = T$  for all valuations  $v$ . Since every valuation  $v$  is defined in terms of an underlying belief

mapping  $\beta$ , this means that  $P$  evaluates to  $T$  for all possible assignments of belief values to its constituent propositional variables.

### 3.3 Basic system properties

It follows from the foregoing that the upper-level formulas obey the rules of classical propositional calculus (CPC) as described, for example, in [17]. To wit, all formulas of the forms

$$\begin{aligned} &(P \dot{\rightarrow} (Q \dot{\rightarrow} P)) \\ &((P \dot{\rightarrow} (Q \dot{\rightarrow} R)) \dot{\rightarrow} ((P \dot{\rightarrow} Q) \dot{\rightarrow} (P \dot{\rightarrow} R))) \\ &((\dot{\rightarrow} P \dot{\rightarrow} \dot{\rightarrow} Q) \dot{\rightarrow} (Q \dot{\rightarrow} P)) \end{aligned}$$

are tautologies, and the rule of *Modus Ponens*,

$$\text{From } P \text{ and } (P \dot{\rightarrow} Q) \text{ infer } Q$$

is *valid* in the sense that, whenever one has  $v(P) = T$  and  $v(P \dot{\rightarrow} Q) = T$ , one also has  $v(Q) = T$ .

These facts may be established via the usual methods. To illustrate, for the first of the above, suppose that for some valuation  $v$  we have

$$v(P \dot{\rightarrow} (Q \dot{\rightarrow} P)) = F$$

Then, by the definition of  $\dot{\rightarrow}$ ,

$$v(P) = T \text{ and } v(Q \dot{\rightarrow} P) = F$$

where the latter implies, again by the definition of  $\dot{\rightarrow}$ ,

$$v(Q) = T \text{ and } v(P) = F.$$

But then we have that  $v(P) = T$  and  $v(P) = F$ , which is a contradiction. Hence there can be no  $v$  that makes the given formula  $F$ , which means that it must be a tautology.

By the ‘‘Adequacy Theorem’’ proven in [17] it is known that the above formulas and rule comprise a complete axiomatization of CPC. Thus, given that these items are all valid in the present semantics, it follows that this semantics validates all classical tautologies that are expressible as upper-level formulas of  $\Sigma$ .

More pertinent to the present work is that the following are valid for all  $i$  and  $j$ :

$$\mathcal{B}_i P \wedge \mathcal{B}_j Q \dot{\rightarrow} \mathcal{B}_{\min[i,j]}(P \wedge Q) \tag{1}$$

and

$$\mathcal{B}_i P \wedge \mathcal{B}_j Q \dot{\rightarrow} \mathcal{B}_{\max[i,j]}(P \vee Q) \tag{2}$$

These express the kind of reasoning discussed in Section 1. To see this, let  $P$  be a propositional variable standing for ‘‘The economy is healthy,’’ let  $Q$  be a different propositional variable standing for ‘‘There can be a useful number system in which  $1 = 0$ ,’’ let  $i = 2$ , and let  $j = -3$ . Then (1) expresses

‘‘If the agent *fairly confidently believes* that the economy is healthy, and the agent *strongly disbelieves* that there can be a useful number system in which  $1 = 0$ , then it *strongly disbelieves* that both are true.’’

Given the same assumptions, (2) expresses

“If the agent *fairly confidently believes* that the economy is healthy, and the agent *strongly disbelieves* that there can be a useful number system in which  $1 = 0$ , then it *fairly confidently believes* that at least one of these are true.”

Formula (1) may be established in the following manner. Suppose that

$$v(\mathcal{B}_i P \dot{\wedge} \mathcal{B}_j Q \dot{\rightarrow} \mathcal{B}_{\min[i,j]}(P \wedge Q)) = F$$

Then, by the definition of  $\dot{\rightarrow}$ ,

$$v(\mathcal{B}_i P \dot{\wedge} \mathcal{B}_j Q) = T \tag{3}$$

and

$$v(\mathcal{B}_{\min[i,j]}(P \wedge Q)) = F \tag{4}$$

From (3), by the definition of  $\dot{\wedge}$ ,

$$v(\mathcal{B}_i P) = T \text{ and } v(\mathcal{B}_j Q) = T$$

which, by the definition of  $v$ , means

$$\beta(P) \in \iota_i \text{ and } \beta(Q) \in \iota_j$$

giving

$$\beta(P \wedge Q) = \min[\beta(P), \beta(Q)] \in \iota_{\min[i,j]}$$

by the definition of  $\beta$ . But from (4), by the definition of  $v$ ,

$$\beta(P \wedge Q) \notin \iota_{\min[i,j]}$$

This is a contradiction; hence (1) is a tautology. Formula (2) may be validated similarly.

One can develop a kind of converse proposition for each of (1) and (2) as the formulas

$$\mathcal{B}_i(P \wedge Q) \dot{\rightarrow} (\mathcal{B}_i P \dot{\wedge} \text{AL} \mathcal{B}_i Q) \dot{\vee} (\mathcal{B}_i Q \dot{\wedge} \text{AL} \mathcal{B}_i P) \tag{5}$$

$$\mathcal{B}_i(P \vee Q) \dot{\rightarrow} (\mathcal{B}_i P \dot{\wedge} \text{AM} \mathcal{B}_i Q) \dot{\vee} (\mathcal{B}_i Q \dot{\wedge} \text{AM} \mathcal{B}_i P) \tag{6}$$

The meanings expressed by these are intuitively plausible, given the usual meanings of “at least,” and “at most.” Formula (5) can be shown a tautology as follows. Suppose

$$v(\mathcal{B}_i(P \wedge Q)) = T \tag{7}$$

Then, by definition of the valuations  $v$ , it is sufficient to show

$$v(\mathcal{B}_i P \dot{\wedge} \text{AL} \mathcal{B}_i Q) \dot{\vee} (\mathcal{B}_i Q \dot{\wedge} \text{AL} \mathcal{B}_i P) = T$$

which, by the definition of  $\dot{\vee}$ , amounts to showing that either

$$v(\mathcal{B}_i P \dot{\wedge} \text{AL} \mathcal{B}_i Q) = T \tag{8}$$

or

$$v(\mathcal{B}_i Q \dot{\wedge} \text{AL} \mathcal{B}_i P) = T \tag{9}$$

From (7), by the property of  $\wedge$  cited in Section 3,

$$\min(\beta(P), \beta(Q)) \in \iota_i$$

This implies that either

$$\beta(P) \in \iota_i \text{ and } \beta(Q) \in \iota_j, \text{ for some } j \geq i, \quad (10)$$

or

$$\beta(Q) \in \iota_i \text{ and } \beta(P) \in \iota_j, \text{ for some } j \geq i, \quad (11)$$

But then, by the properties of  $\wedge$  and AL cited in Section 3, (10) implies (8), and (11) implies (9). Formula (6) may be verified similarly.

Other tautologies of this system are

$$\mathcal{B}_i \neg P \leftrightarrow \mathcal{B}_{-i} P \quad (12)$$

$$\dot{\neg} \mathcal{B}_i P \leftrightarrow \dot{\vee}_{-4 \leq j \leq 4, j \neq i} \mathcal{B}_j P \quad (13)$$

The former shows that the lower-level negation interacts with the belief operators in intuitively plausible ways. To illustrate, if  $i = 2$ , (12) says that “the agent *fairly confidently believes not-P*” means the same as “the agent *fairly confidently disbelieves P*,” and if  $i = -3$ , (12) says that “the agent *strongly disbelieves not-P*” means the same as “the agent *strongly believes P*.” The latter formula, on the other hand, does not reflect an naturally occurring property of negation, but rather expresses only the logical consequence of the various belief operators having been defined to be semantically mutually exclusive.

To verify (12), observe that

$$1 - \beta(P) \in \iota_i \text{ iff } \beta(P) \in \iota_{-i}$$

which, by the definition of  $\beta$ , gives

$$\beta(\neg P) \in \iota_i \text{ iff } \beta(P) \in \iota_{-i}$$

which in turn, by the definition of  $v$ , gives

$$v(\mathcal{B}_i \neg P) = T \text{ iff } v(\mathcal{B}_{-i} P) = T$$

The latter leads to (12) by analysis of the definition of  $\leftrightarrow$ . Verifying (13) is more tedious but straightforward.

Two immediate consequences of (12) are that

$$\mathcal{K} \neg P \leftrightarrow \mathcal{B}_{-4} P$$

$$\mathcal{B}_{-4} \neg P \leftrightarrow \mathcal{K} P$$

are tautologies. These express further intuitively plausible consequences of the definitions, to wit, the agent *knows* that *not-P* if and only if he *unequivocally disbelieves P*, and he *unequivocally disbelieves not-P* if and only if he *knows P*. Three tautologies that can be verified along the same lines as (13) are

$$\dot{\neg} \mathcal{B} P \leftrightarrow \mathcal{D} P \dot{\vee} \mathcal{U} P$$

$$\dot{\neg} \mathcal{D} P \leftrightarrow \mathcal{B} P \dot{\vee} \mathcal{U} P$$

$$\dot{\neg} \mathcal{U} P \leftrightarrow \mathcal{B} P \dot{\vee} \mathcal{D} P$$

A particularly useful result is that, with respect to the connectives  $\neg$ ,  $\vee$ , and  $\wedge$ , the lower-level of  $\Sigma$  imposes a structure on the interval  $[0, 1]$  that has at least the richness of

	4	3	2	1	0	-1	-2	-3	-4
4	4, 4	4, 3	4, 2	4, 1	4, 0	4, -1	4, -2	4, -3	4, -4
3	4, 3	3, 2	2, 1	1, 0	0, -1	-1, -2	-2, -3	-3, -4	Any
2	4, 2	2, 1	1, 0	0, -1	-1, -2	-2, -3	-3, -4	Any	Any
1	4, 1	1, 0	0, -1	-1, -2	-2, -3	-3, -4	Any	Any	Any
0	4, 0	0, -1	-1, -2	-2, -3	-3, -4	Any	Any	Any	Any
-1	4, -1	-1, -2	-2, -3	-3, -4	Any	Any	Any	Any	Any
-2	4, -2	-2, -3	-3, -4	Any	Any	Any	Any	Any	Any
-3	4, -3	-3, -4	Any	Any	Any	Any	Any	Any	Any
-4	4, -4	Any	Any	Any	Any	Any	Any	Any	Any

Table 1: Chaining with the Łukasiewicz  $\rightarrow$ .

a De Morgan lattice (an in-depth analysis of such algebras may be found in [31, p. 38]). This is established by verifying that the following are tautologies.

$$\begin{aligned}
\mathcal{B}_i P &\leftrightarrow \mathcal{B}_i \neg \neg P && \text{(involution)} \\
\mathcal{B}_i (P \vee Q) &\leftrightarrow \mathcal{B}_i (Q \vee P) && \text{(commutativity)} \\
\mathcal{B}_i (P \wedge Q) &\leftrightarrow \mathcal{B}_i (Q \wedge P) \\
\mathcal{B}_i ((P \vee Q) \vee R) &\leftrightarrow \mathcal{B}_i (P \vee (Q \vee R)) && \text{(associativity)} \\
\mathcal{B}_i ((P \wedge Q) \wedge R) &\leftrightarrow \mathcal{B}_i (P \wedge (Q \wedge R)) \\
\mathcal{B}_i (P \vee (Q \wedge R)) &\leftrightarrow \mathcal{B}_i ((P \vee Q) \wedge (P \vee R)) && \text{(distributivity)} \\
\mathcal{B}_i (P \wedge (Q \vee R)) &\leftrightarrow \mathcal{B}_i ((P \wedge Q) \vee (P \wedge R)) \\
\mathcal{B}_i (P \vee (P \wedge Q)) &\leftrightarrow \mathcal{B}_i (P) && \text{(absorption)} \\
\mathcal{B}_i (P \wedge (P \vee Q)) &\leftrightarrow \mathcal{B}_i (P) \\
\mathcal{B}_i (P \vee Q) &\leftrightarrow \mathcal{B}_i \neg (\neg P \wedge \neg Q) && \text{(De Morgan's laws)} \\
\mathcal{B}_i (P \wedge Q) &\leftrightarrow \mathcal{B}_i \neg (\neg P \vee \neg Q)
\end{aligned}$$

Establishing most of these is routine. For the absorption and De Morgan's laws, consider the cases (i)  $\beta(P) \geq \beta(Q)$  and (ii)  $\beta(P) < \beta(Q)$ , and for the distributivity laws, consider the cases (i)  $\beta(P) \geq \beta(Q) \geq \beta(R)$ , (ii)  $\beta(P) \geq \beta(R) \geq \beta(Q)$ , (iii)  $\beta(Q) \geq \beta(P) \geq \beta(R)$ , (iv)  $\beta(Q) \geq \beta(R) \geq \beta(P)$ , (v)  $\beta(R) \geq \beta(P) \geq \beta(Q)$ , and (vi)  $\beta(R) \geq \beta(Q) \geq \beta(P)$ .

This shows that, as a reasoning system,  $\Sigma$  is quite strong. It is not as rich as a Boolean algebra, however, since it does not have unit and zero elements, i.e., there are no lower-level formulas that map identically to either 0 or 1.

### 3.4 Inference for the lower level

The literature on fuzzy logic is replete with varieties of multi-valued modes of inference. The paper [29] examined 10 such operators, and there are many others. Among these the most well-known is the Łukasiewicz  $\rightarrow$ , defined for the infinitary logic  $\aleph_1$  [cf. 27]. This

may be introduced into  $\Sigma$  in the following manner.<sup>4</sup> Expand the symbol set to include the connective  $\rightarrow$ , extend the definition of the set  $F_2$  to include formulas of the form  $(P \rightarrow Q)$ , and extend the mappings  $\beta$  by

$$\beta(P \rightarrow Q) = \min[1, 1 - \beta(P) + \beta(Q)]$$

Then the question arises as to the properties of this connective. In particular, it is of interest to consider the status of formulas expressing inference chaining:

$$\mathcal{B}_i(P \rightarrow Q) \wedge \mathcal{B}_j(Q \rightarrow R) \dot{\rightarrow} \mathcal{B}_k(P \rightarrow R) \quad (14)$$

The issue is, given choices for  $i$  and  $j$  from  $4, \dots, -4$ , for what corresponding  $k$  will the above formula be a tautology. It turns out that, as is typical of fuzzy reasoning systems, the precision of the consequences tends to degrade with successive applications of  $\rightarrow$ . For the present system, this means that, instead of formulas of the form (14), we must consider formulas of the more general form

$$\mathcal{B}_i(P \rightarrow Q) \wedge \mathcal{B}_j(Q \rightarrow R) \dot{\rightarrow} \mathcal{B}_k(P \rightarrow R) \dot{\vee} \dots \dot{\vee} \mathcal{B}_{k'}(P \rightarrow R) \quad (15)$$

The results of the analysis are shown in Table 1. Row labels are  $i$ 's, column labels are  $j$ 's, and the content of the  $i, j$  cell is the corresponding  $k, k'$  pair. For example, if  $i = 1$  (*somewhat believes*) and  $j = 4$  (*unequivocally believes*), then  $k = 4$  (*unequivocally believes*) and  $k' = 1$  (*somewhat believes*). The cells containing "Any" indicate  $i, j$  for which any choices of  $k, k'$  will result in a tautology. The entries in Table 1 may be verified by standard algebraic techniques.

## 4 Belief revision

It would be of interest to consider how  $\Sigma$  might be employed within the broader context of a system for belief revision. The subject of belief revision is concerned with devising formal representations of the manner in which an agent's collection of beliefs may change over time due to the acquisition of new information. This includes both the addition of new beliefs to the current belief set and the removal of old beliefs in the light of newer countervailing information. A seminal work in this area is Gärdenfors [9]. The subject has received increased attention in recent years, particularly because of its close relation to nonmonotonic reasoning (cf. Antoniou [1] for an overview). The collection [10] contains further developments along these lines.

In [30] there was defined the notion of a *dynamic reasoning system* (DRS), which explicitly portrays reasoning as an activity that takes place in time. This is obtained from the conventional notion of formal logical system by lending special semantic status to the concept of a derivation path (i.e., a proof). Introduction of new knowledge or beliefs into the path occurs in two ways: either new propositions are added in the form of axioms, or some propositions are derived from earlier ones by means of an inference rule. In either case, the action is regarded as occurring in a discrete time step, and the new proposition is labeled with a time stamp (an integer) indicating the step at which this occurred. Moreover, for propositions entered into the path as a result of rule applications, the label

---

<sup>4</sup> This adapts a similar discussion that appeared in [32] and corrects a technical error in the previous analysis.

additionally contains a record of which inference rule was used and which propositions were employed as premises. At any given time, the contents of the path is regarded as being the sum total of the agent’s knowledge and beliefs as of that time. This is to be contrasted with other systems of belief revision, which assume that the agent additionally knows all the logical consequences of the basic belief set. Such systems are said to exhibit “logical omniscience.” For an in-depth analysis of this issue, together with a manner of addressing it, see the paper by Fagin, Halpern, Moses, and Vardi [7].

The DRS framework allows one to portray belief revision as a form of back tracking along the lines of the “truth maintenance,” or “reason maintenance,” systems devised by Doyle [2, 3]. If at some point in time a contradictory proposition is entered into the path, then this triggers a process (possibly human assisted) of working backwards through the path, following the information stored in the labels, looking for “culprit” formulas that may be held responsible for having led to the contradiction. Some of these formulas are then removed from the set of currently held beliefs, so as to eliminate the contradiction, and one uses the formula labels once again to forward chain through all deductions that originated from these formulas and remove them from the current belief set as well. This process can obviously give rise to complexity issues, but it is nonetheless both theoretically feasible and finitary.

The notion of a DRS was proposed initially for encoding reasoning with a logic  $Q$  of *qualified syllogisms*, which, as mentioned, served as the model after which the current logic  $\Sigma$  has been patterned. Accordingly, it is not difficult to see how  $\Sigma$  could be cast in this framework and thereby provided with suitable mechanisms for belief revision. The same framework also provides a means for encoding rules for belief combination. For example, if at some point in the path it is determined that  $P$  is *strongly believed*, and at some other point it is determined that  $P$  is *somewhat disbelieved*, then at a later point a combination rule can be applied to deduce that, e.g.,  $P$  is only *somewhat believed*. It is planned to elaborate the details of these manners of belief management in a future work.

## 5 Concluding remarks

It is of interest that the various lower-level operators, including the Łukasiewicz  $\rightarrow$ , are closed with respect to the intervals  $\iota_4, \dots, \iota_{-4}$ . That is, for each  $i$ , there is a  $j$  such that

$$\text{if } \beta(P) \in \iota_i, \text{ then } \beta(\neg P) \in \iota_j$$

for each  $i, j$ , there is a  $k$  such that

$$\text{if } \beta(P) \in \iota_i \text{ and } \beta(Q) \in \iota_j, \text{ then } \beta(P \vee Q) \in \iota_k$$

and similarly, for each  $i, j$ , there is a  $k$  such that

$$\text{if } \beta(P) \in \iota_i \text{ and } \beta(Q) \in \iota_j, \text{ then } \beta(P \rightarrow Q) \in \iota_k$$

This means that, while the belief mappings are useful for providing an interpretation of the linguistic terms  $\mathcal{B}_i$  and for justifying the interpretations of the logical connectives, all reasoning within  $\Sigma$  can be carried out as operations performed directly on these terms. In the case of  $\vee$ , for example, one can set up a table analogous to Table 1, showing, for each  $i, j$ , what should be the corresponding  $k$  such that, if  $\mathcal{B}_i P$  and  $\mathcal{B}_j Q$  are true, then

$\mathcal{B}_k(P \vee Q)$  is true. In other words, the valuations  $v$  can be defined in terms of such tables, without reference to the belief mappings.

A consequence of this is that the semantics for  $\Sigma$  is not difficult to implement on a computer. Programs that evaluate propositions according to such tables would not differ much from those that evaluate formulas of classical propositional calculus according to the usual truth tables.

Another issue is that of a semantically complete axiomatization for  $\Sigma$ . There have been recent successes in axiomatizing what may be regarded as the core of fuzzy logic. Hähnle [18] presents one such formalization for use in logic circuit design. A major undertaking along these lines is that of Hajek [12] (cf. also Hajek and Godo [13]). That work explores the extent to which the elements of fuzzy logic can be expressed as a multivalent logic using only the language of first-order logic, and it provides semantic completeness results for all the relevant formalisms.

How to go about axiomatizing the dual-leveled types of formalisms studied here, however, is another matter. One such effort, capturing some aspects of fuzzy logic, is the formal theory of “semantic equivalence” developed in [28]. There a semantic completeness result made use of some deep theorems from abstract algebra. But it remains to be determined whether such methods will apply here. Indeed the kind of interaction between levels as expressed by lines (1) and (2) of Section 3.3 makes the system  $\Sigma$  somewhat unique—and the task of axiomatizing it even more challenging. Exploration of this issue has been reserved for a future work.

## References

- [1] G. Antoniou, *Nonmonotonic reasoning*, MIT Press, Cambridge, MA, 1997.
- [2] J. Doyle, A truth maintenance system, *Artificial Intelligence* 12(3) (1979) 231–272.
- [3] J. Doyle, Reason maintenance and belief revision: foundations vs. coherence theories, in P. Gärdenfors, ed., *Belief revision*, Cambridge University Press, New York, 1992, pp. 29–51.
- [4] R. Fagin, R. and J. Halpern, Belief, awareness, and limited reasoning, *Artificial Intelligence*, 34 (1988) 39–76.
- [5] R. Fagin and J. Halpern, Uncertainty, belief, and probability, *Computational Intelligence*, 6 (1991) 160–173.
- [6] R. Fagin, J. Halpern, Y. Moses, and M.Y. Vardi, *Reasoning about knowledge*, MIT Press, Cambridge, MA, 1995.
- [7] R. Fagin, J. Halpern, Y. Moses, and M.Y. Vardi, A nonstandard approach to the logical omniscience problem, *Artificial Intelligence*, 79 (1995) 203–240.
- [8] K.M. Ford, C. Glymour, and P.J. Hayes, eds., *Android Epistemology*, MIT Press, Cambridge, MA, 1995.
- [9] P. Gärdenfors, *Knowledge in flux: modeling the dynamics of epistemic states*, MIT Press/Bradford Books, Cambridge, MA, 1988.

- [10] P. Gärdenfors, Belief revision: an introduction, in P. Gärdenfors, ed., *Belief revision*, Cambridge University Press, New York, 1992, pp. 1–28.
- [11] E.L. Gettier, Is justified true belief knowledge?, *Analysis*, 23 (1962-1963) 121–123.
- [12] P. Hájek, *Metamathematics of fuzzy logic*, Kluwer, Boston, MA, 1998.
- [13] P. Hájek and L. Godo, Deductive systems of fuzzy logic, *Tatra Mountains Math. Publ.*, 13 (1997) 35–66.
- [14] J.Y. Halpern and R. Fagin, Two views of belief: belief as generalized probability and belief as evidence, *Artificial Intelligence*, 54 (1992) 275–317.
- [15] J.Y. Halpern and D. A. McAllester, Likelihood, probability, and knowledge, *Computational Intelligence*, 5 (1989) 151–160.
- [16] J.Y. Halpern and Y. Moses, A guide to completeness and complexity of knowledge and belief, *Artificial Intelligence*, 54 (1992) 310–379.
- [17] A.G. Hamilton, *Logic for mathematicians, Revised Edition*, Cambridge University Press, 1988.
- [18] R. Hähnle, Proof theory of many-valued optimization-logic design: connections and interactions, *Soft Computing*, 01 (1997) 107–119.
- [19] J. Hintikka, *Knowledge and belief: an introduction to the logic of the two notions*, Cornell University Press, Ithaca, NY, 1962.
- [20] K. Konolige, *A deduction model of belief*, Morgan Kaufmann, Los Altos, CA, 1986.
- [21] S. Kraus and D. Lehmann, Knowledge, belief and time, *Theoretical Computer Science*, 58 (1988) 155–174.
- [22] W. Lenzen, Recent work in epistemic logic, *Acta Philosophica Fennica*, XXX(1), 1978.
- [23] F. Lin and Y. Shoham, A logic of justified assumptions, *Artificial Intelligence*, 57 (1992) 271–289.
- [24] R.C. Moore, Semantical considerations on nonmonotonic logic, *Artificial Intelligence*, 25 (1985) 75–94.
- [25] R.C. Moore, Autoepistemic logic, in: P. Smets, E.H. Mamdani, D. Dubois, and H. Prade, eds., *Non-Standard Logics for Automated Reasoning*, Academic Press, San Diego, CA, (1988) 105–136.
- [26] Y. Moses and Y. Shoham, Belief as defeasible knowledge, *Artificial Intelligence*, 62(2) (1993) 299–322.
- [27] N. Rescher, *Many-valued logic*, McGraw-Hill, NY, 1969.

- [28] D.G. Schwartz, Axioms for a theory of semantic equivalence, *Fuzzy Sets and Systems*, 21 (1987) 319–349.
- [29] D.G. Schwartz, Fuzzy inference in a formal theory of semantic equivalence, *Fuzzy Sets and Systems*, 31 (1989) 205–216.
- [30] D.G. Schwartz, Dynamic reasoning with qualified syllogisms, *Artificial Intelligence*, 93(1-2) (1997) 103–167.
- [31] D.G. Schwartz, Layman’s probability theory: a calculus for reasoning with linguistic likelihood, *Information Sciences*, 126(1-4) (2000) 71–82.
- [32] D.G. Schwartz, Łukasiewicz inference in Layman’s Probability Theory, *Proceedings of the North American Fuzzy Information Processing Society, NAFIPS’2000*, Atlanta, GA, July 13–15, 2000, pp. 420–423.
- [33] G. Shafer, *A mathematical theory of evidence*, Princeton University Press, Princeton, NJ, 1976.
- [34] Y. Shoham, Agent-oriented programming, *Artificial Intelligence*, 60(1) (1993) 51–92.
- [35] E. Thijsse, *Partial logic and knowledge representation*, Doctoral Dissertation, Tilburg University, Eburon Publishers, Delft, The Netherlands, 1992.
- [36] E. Thijsse, Combining partial and classical semantics: a hybrid approach to belief and awareness, in P. Doherty, ed., *Partiality, Modality, and Nonmonotonicity*, CSLI Publications, Center for the Study of Language and Information, Stanford, CA, 1996, pp. 223–249.
- [37] F. Voorbraak, *As far as I know: epistemic logic and uncertainty*, Doctoral Dissertation, Department of Philosophy, University of Utrecht, The Netherlands, 1993.
- [38] L.A. Zadeh, Fuzzy sets, *Information and Control*, 8 (1965) 338–353.
- [39] L.A. Zadeh, The concept of a linguistic variable and its application to approximate reasoning, part I, *Information Sciences* 8 (1975) 199–249; part II, 8, 301–357; part III, 9, 43–80. Reprinted in Yager, R.R., Ovchinnikov, S., Tong, R.M., and Nguyen, H.T., Eds., *Fuzzy Sets and Applications: Selected Papers by L.A. Zadeh*, John Wiley and Sons, New York, 1987, pp. 219–269, 271–327, 329–366.
- [40] L.A. Zadeh, Fuzzy logic = computing with words, *IEEE Transactions on Fuzzy Systems*, 4(2) (1996) 103–111.
- [41] L.A. Zadeh, From computing with numbers to computing with words—from manipulation of measurements to manipulation of perceptions, *IEEE Transactions on Circuits and Systems*, 45(1) (1999) 105–119.
- [42] L.A. Zadeh, A new direction in AI: toward a computational theory of perceptions, *AI Magazine*, 22(1) (2001) 73–84.