A Syntactic Possibilistic Belief Change Operator: Theory and empirical study

Célia da Costa Pereira, Andrea G. B. Tettamanzi

► To cite this version:

Célia da Costa Pereira, Andrea G. B. Tettamanzi. A Syntactic Possibilistic Belief Change Operator: Theory and empirical study. Web Intelligence and Agent Systems, 2014, 12 (2), pp.155-169. 10.3233/WIA-140290 . hal-00983200

HAL Id: hal-00983200

<https://hal.science/hal-00983200>

Submitted on 26 Apr 2014

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

A Syntactic Possibilistic Belief Change Operator: Theory and empirical study

Célia da Costa Pereira*, and
Andrea G. B. Tettamanzi

*Université Nice Sophia Antipolis, I3S, UMR 7271,
06900 Sophia Antipolis, France
E-mail: celia.pereira@unice.fr,
andrea.tettamanzi@unice.fr*

Abstract. We propose a syntactic possibilistic belief-change operator, which operates on a belief base of necessity-valued formulas. Such a base may be regarded as a finite and compact encoding of a possibility distribution over a possibly infinite set of interpretations. The proposed operator is designed so that it behaves like a semantic possibilistic belief-change operator for BDI agents recently proposed in the literature. The equivalence of the semantic and syntactic operators is then proved. Experimental results are presented. The aim of these experiments is to demonstrate that the cost of belief revision (expressed in terms of the number of entailment checks required) as well as the size of the belief base do not explode as the number of new pieces of information (formulas) supplied increases.

Keywords: BDI Agents, Belief Change, Possibility Theory.

1. Introduction

This article is an extended version of [9].

When representing uncertainty in logics, it is often supposed that there exists a set of possible worlds (models) which are candidates to be the real world. The uncertainty is due to the fact that we do not know which among those worlds corresponds to the real state. Two different situations can be considered:

- we cannot distinguish which world has a greater chance to be the real state than another — the non-deterministic case;
- some information is available leading us to suppose that not all these worlds are equally possible to be the real world state.

These two cases are differently treated depending on the logic used to represent the uncertainty. If we use probability theory, in the first case, all the worlds are associated to an equal probability while, in the second case, we need to have some evidence allowing us to associate different probabilities to the different worlds depending on their “distance” with respect to the real one. If we use possibility theory, the first case corresponds to total ignorance, and the second case may be solved by considering the plausibility order associated to the worlds.

Several approaches have been proposed for representing uncertainty in the BDI (Beliefs Desires and Intentions) components of an agent. The main goal of such approaches is to extend the traditional BDI model of agency to make it more suitable to represent real situations in which uncertainty is omnipresent. There are essentially two ways for dealing with uncertainty. By using probabilistic-based models which are suitable when information about past experiences is available, and by using possibilistic-based models which in their turn are more suitable when we lack statistical data but a notion of order on the events is available. This latter uncertainty model for representation has demonstrated its usefulness in representing uncertainty on beliefs since the seminal work by Dubois and Prade [13]. In that work, the authors pointed out the close relationships between the theory of belief revision developed by Gärdenfors based on the notion of epistemic entrenchment, and possibility theory applied to automated reasoning under uncertainty.

Other proposals have followed. For example, Dragoni and Giorgini [11] presented a model for belief revision in which they integrate symbolic and numer-

*Corresponding author. E-mail: celia.pereira@unice.fr.

*Corresponding author. E-mail: celia.pereira@unice.fr.

ical methods into a BDI architecture. Casali and colleagues [7] proposed a general model for graded BDI agents and an architecture for modeling the agent's graded mental attitudes. Like Casali, Blee and colleagues [6] introduce levels in all the mentalistic notions of BDI, as well as using numeric, possibilistic-type functions in its semantics. We have presented an integrated theoretical framework, grounded in possibility theory, to account for all the aspects involved in representing and changing beliefs for cognitive agents [8]. In that framework, graded beliefs are represented by means of a possibility distribution over interpretations, and a belief-change operator is proposed which obeys a possibilistic formulation of the AGM revision rationality postulates $K*1-K*8$ [16], and is a generalization of the possibilistic conditioning operator of Dubois and colleagues [14]. Such operator is one of the members of a family of possibilistic conditioning operators studied in [5].

Although our framework looks interesting and promising, the use of a possibility distribution to represent beliefs, while allowing us to model most of the intuitive properties in an elegant and natural way, poses computational problems. Indeed, for a propositional language, the interpretations are 2 to the power the number of atomic propositions; with more than a few dozen atomic propositions, a direct representation of a possibility distribution would require more memory space than available on most state-of-the-art computers. For more expressive languages, the set of interpretations may even be infinite.

Obviously, equivalent, but less demanding, strategies to encode and manipulate beliefs should be devised if such an approach is to be adopted in a realistic setting. In this paper, we propose one strategy to work around the direct representation of beliefs as a possibility distribution, which consists of adopting a *syntactic* representation, whereby the beliefs of an agent are represented by means of a fuzzy set of formulas in the language of choice. However, to achieve full equivalence with the direct, *semantic* representation, the belief-change operator has to be reformulated in terms of the syntactic representation.

To this aim, we devise a syntactic belief-change operator that works on a belief base, i.e., a fuzzy set of formulas in the language of choice, where membership degrees are regarded as necessity degrees, and transforms it to account for the arrival of new information from a partially trusted source. Then we prove that the syntactic operator is equivalent to the semantic operator. Experimental results are also presented. They

demonstrate that the cost of belief revision (expressed in terms of the number of entailment checks required) as well as the size of the belief base do not explode as the number of new pieces of information (formulas) supplied increases.

An alternative and very popular approach to representing beliefs and revising them with new evidence is Philippe Smets transferable belief model [20], which is based on the Dempster-Shafer theory of evidence, which in turn may be regarded as a generalization of the theory of probability. As pointed out by [12], the main problem with the Dempster-Shafer approach is its computational complexity. One should generate a frame of $2^{|\Omega|}$ elements, where Ω is the space of events (or, equivalently, the set of all interpretations)! This can be contrasted to the (semantic) possibilistic and a simple probabilistic representation, which “only” require to explicitly represent a frame of $|\Omega|$ elements.

The reason why a direct probabilistic representation of beliefs is not a viable alternative to the framework we propose is that, unlike the transferable belief model, it fails to distinguish *stochastic uncertainty*, resulting from a system behaving in a random way, from *epistemic uncertainty* (or ignorance), resulting from a lack of knowledge about the system. Such confusion may lead to incorrect results [20].

Furthermore, probabilistic revision does not satisfy the AGM postulates, especially because it fails to address the notion of minimal change [12].

The paper is organized as follows: Section 2 briefly reviews background notions required to follow the paper. Section 3 discusses the semantic and the syntactic representation, showing their equivalence. Section 4 reviews the semantic belief-change operator, then proposes a syntactic belief-change operators, and proves its equivalence. Section 5 provides an empirical study of the behavior of the proposed operator. Finally, Section 6 concludes.

2. Preliminaries

In this section, we briefly review the essential background and definitions on fuzzy set theory and possibility theory.

2.1. Fuzzy Sets

Fuzzy sets [21] are a generalization of classical (crisp) sets obtained by replacing the characteristic function of a set A , χ_A , which takes up values in $\{0, 1\}$

($\chi_A(x) = 1$ iff $x \in A$, $\chi_A(x) = 0$ otherwise) with a *membership function* μ_A , which can take up any value in $[0, 1]$. The value $\mu_A(x)$ or, more simply, $A(x)$ is the membership degree of element x in A , i.e., the degree to which x belongs in A .

A fuzzy set is completely defined by its membership function. Therefore, it is useful to define a few terms describing various features of this function, summarized in Figure 1. Given a fuzzy set A , its *core* is the (conventional) set of all elements x such that $A(x) = 1$; its *support*, $\text{supp}(A)$, is the set of all x such that $A(x) > 0$. A fuzzy set is *normal* if its core is nonempty. The set of all elements x of A such that $A(x) \geq \alpha$, for a given $\alpha \in (0, 1]$, is called the α -cut of A , denoted A_α . Sometimes it is convenient to define also the notion of a *strict* α -cut, $A_{>\alpha}$, as the set of all elements x of A such that $A(x) > \alpha$.

Fig. 1. Core, support, and α -cuts of a set A of the real line.

A convenient notational convention we will adopt throughout this paper for fuzzy sets, when the universe of discourse is discrete is, for fuzzy set A ,

$$A = \frac{A(x_1)}{x_1} + \frac{A(x_2)}{x_2} + \dots = \sum_i \frac{A(x_i)}{x_i}. \quad (1)$$

This notation is nothing more than a formal device and the fractions do not have to be interpreted as divisions but just as ordered pairs, while the $+$ does not stand for algebraic sum but rather for a function-theoretic union.

The usual set-theoretic operations of union, intersection, and complement can be defined as a generalization of their counterparts on classical sets by introducing two families of operators, called triangular norms and triangular co-norms. In practice, it is usual to employ the min norm for intersection and the max co-norm for union. Given two fuzzy sets A and B , and

an element x ,

$$(A \cup B)(x) = \max\{A(x), B(x)\}; \quad (2)$$

$$(A \cap B)(x) = \min\{A(x), B(x)\}; \quad (3)$$

$$\bar{A}(x) = 1 - A(x). \quad (4)$$

Finally, given two fuzzy sets A and B , $A \subseteq B$ if and only if, for all element x , $A(x) \leq B(x)$.

2.2. Possibility Theory

The membership function of a fuzzy set describes the more or less possible and mutually exclusive values of one (or more) variable(s). Such a function can then be seen as a possibility distribution [22]. Indeed, if F designates the fuzzy set of possible values of a variable X , $\pi_X = \mu_F$ is called the possibility distribution associated to X . The identity $\mu_F(v) = \pi_X(v)$ means that the membership degree of v to F is equal to the possibility degree of X being equal to v when all we know about X is that its value is in F . A possibility distribution for which there exists a completely possible value ($\exists v_0; \pi(v_0) = 1$) is said to be *normalized*.

Definition 1 (Possibility and Necessity Measures) A possibility distribution π induces a possibility measure and its dual necessity measure, denoted by Π and N respectively. Both measures apply to a crisp set A and are defined as follows:

$$\Pi(A) = \max_{s \in A} \pi(s); \quad (5)$$

$$N(A) = 1 - \Pi(\bar{A}) = \min_{s \in \bar{A}} \{1 - \pi(s)\}. \quad (6)$$

In words, the possibility measure of set A corresponds to the greatest of the possibilities associated to its elements; conversely, the necessity measure of A is equivalent to the impossibility of its complement $\bar{A}$.

3. Representing Graded Beliefs

3.1. Language and Interpretations

Information manipulated by a cognitive agent must be somehow represented. For the sake of simplicity, we use here perhaps the simplest symbolic representation, in the form of a classical propositional language.

Definition 2 (Language) Let $\mathcal{A}$ be a finite¹ set of atomic propositions and let $\mathcal{L}$ be the propositional language such that $\mathcal{A} \cup \{\top, \perp\} \subseteq \mathcal{L}$, and, $\forall \phi, \psi \in \mathcal{L}$, $\neg\phi \in \mathcal{L}$, $\phi \wedge \psi \in \mathcal{L}$, $\phi \vee \psi \in \mathcal{L}$.

As usual, one may define additional logical connectives and consider them as useful shorthands for combinations of connectives of $\mathcal{L}$, e.g., $\phi \supset \psi \equiv \neg\phi \vee \psi$.

We will denote by $\Omega = \{0, 1\}^{\mathcal{A}}$ the set of all interpretations on $\mathcal{A}$. An interpretation $\mathcal{I} \in \Omega$ is a function $\mathcal{I} : \mathcal{A} \rightarrow \{0, 1\}$ assigning a truth value $p^{\mathcal{I}}$ to every atomic proposition $p \in \mathcal{A}$ and, by extension, a truth value $\phi^{\mathcal{I}}$ to all formulas $\phi \in \mathcal{L}$.

Definition 3 The notation $[\phi]$ denotes the set of all models (namely, interpretations satisfying ϕ) of a formula $\phi \in \mathcal{L}$:

$$[\phi] = \{\mathcal{I} \in \Omega : \mathcal{I} \models \phi\}.$$

Likewise, if $S \subseteq \mathcal{L}$ is a set of formulas,

$$[S] = \{\mathcal{I} \in \Omega : \forall \phi \in S, \mathcal{I} \models \phi\} = \bigcap_{\phi \in S} [\phi].$$

3.2. Syntactic and Semantic Representations

Two alternative ways for representing graded beliefs can be obtained by following two distinct lines of reasoning.

The first line of reasoning starts from the logical idea of a belief set, like the one used in the AGM theory of revision [1]. In the literature on non-graded belief revision, the agent's belief state may be represented in different ways. As belief sets, which are sets of sentences closed under logical consequence, like in [1], or, alternatively, as belief bases, which are sets of sentences that are not logically closed [18] and contain the sentences that represent the explicit beliefs from which all the other beliefs can be derived. The representation based on belief sets has some advantages, but is not computationally adequate. On the other hand, a belief base may be regarded as a finite and compact representation of a belief set, and is, therefore, much more suited to computation.

As for the representation of graded beliefs, one can allow some formulas of the base to be believed only

to a given degree, thus obtaining a fuzzy belief base, which is a fuzzy set B of formulas. The degree to which a given formula $\phi \in \mathcal{L}$ is believed can be calculated as

$$B(\phi) = \max\{\alpha : B_\alpha \vdash \phi\}. \quad (7)$$

This is the *syntactic* representation of graded beliefs.

Alternatively, one may regard a *belief* as a necessity degree induced by a normalized possibility distribution π on the possible worlds $\mathcal{I} \in \Omega$ [4]:

$$\pi : \Omega \rightarrow [0, 1]; \quad (8)$$

where $\pi(\mathcal{I})$ is the possibility degree of interpretation $\mathcal{I}$. In this case, the degree to which a given formula $\phi \in \mathcal{L}$ is believed can be calculated as

$$B(\phi) = N([\phi]) = 1 - \max_{\mathcal{I} \not\models \phi} \pi(\mathcal{I}), \quad (9)$$

where N is the necessity measure induced by π . This is the *semantic* representation of graded beliefs.

This latter was our choice in the original proposal of the theoretical framework [8].

3.3. Equivalence

A direct consequence of a fundamental result on the equivalence between sets of necessity-values formulas and possibility distributions [15], is that the syntactic and the semantic representations of graded beliefs are equivalent. Therefore, they may be used interchangeably as convenience demands.

This means that, given a fuzzy belief base B such that, for all α , B_α is consistent, one can construct a possibility distribution π such that, for all $\phi \in \mathcal{L}$, $N([\phi]) = \max\{\alpha : B_\alpha \vdash \phi\}$, where B_α is the α -cut of fuzzy belief base B .

In particular, π may be defined as follows: for all $\mathcal{I} \in \Omega$,

$$\pi(\mathcal{I}) = \min_{\phi : \mathcal{I} \not\models \phi} \{1 - B(\phi)\} = 1 - \max_{\phi : \mathcal{I} \not\models \phi} B(\phi), \quad (10)$$

or, equivalently,

$$\pi(\mathcal{I}) = 1 - \max\{\alpha : B_\alpha \vdash \neg\phi_{\mathcal{I}}\}, \quad (11)$$

where $\phi_{\mathcal{I}}$ denotes the *minterm* of $\mathcal{I}$, i.e., the formula satisfied by $\mathcal{I}$ only.

¹Like in [3], we adopt the restriction to the finite case in order to use standard definitions of possibilistic logic. Extensions of possibilistic logic to the infinite case are discussed for example in [10].

Notice that π is normalized. Indeed, since, by hypothesis, for all α , B_α is consistent, there exists an interpretation $\mathcal{I}^* \in \Omega$, such that, for all $\alpha \in (0, 1]$, $\mathcal{I}^* \models B_\alpha$; therefore, $\pi(\mathcal{I}^*) = 1$, because no formula ϕ exists such that $\mathcal{I}^* \not\models \phi$ and $B(\phi) > 0$.

3.4. Properties of Graded Beliefs

Straightforward consequences of the properties of possibility and necessity measures are that $\mathbf{B}(\phi) > 0 \Rightarrow \mathbf{B}(\neg\phi) = 0$, this means that if the agent somehow believes ϕ then it cannot believe $\neg\phi$ at all; and, for all $\phi, \psi \in \mathcal{L}$,

$$\mathbf{B}(\top) = 1, \quad (12)$$

$$\mathbf{B}(\perp) = 0, \quad (13)$$

$$\mathbf{B}(\phi \wedge \psi) = \min\{\mathbf{B}(\phi), \mathbf{B}(\psi)\}, \quad (14)$$

$$\mathbf{B}(\phi \vee \psi) \geq \max\{\mathbf{B}(\phi), \mathbf{B}(\psi)\}. \quad (15)$$

Another straightforward property is that, for all $\phi \in \mathcal{L}$, $\mathbf{B}(\phi) \geq B(\phi)$.

The idea of a syntactic representation of beliefs is that it should be *parsimonious*, and ideally minimal with respect to (fuzzy) set inclusion, i.e., the belief base does not need to include explicitly formulas whose degree of belief can be derived from it. In particular, this means that if formula ϕ is already a member of B with degree $B(\phi)$, then, if $\alpha \leq B(\phi)$,

$$B \cup \frac{\alpha}{\phi} = B.$$

In general, if $\alpha \leq \mathbf{B}(\phi)$, which is equivalent to saying that $B_\alpha \vdash \phi$, adding $\frac{\alpha}{\phi}$ to B would be redundant, and may be dispensed with. Since all tautologies are always believed to degree 1, adding a tautology to a belief base is always redundant, and may thus be avoided.

4. Belief Change

In this section we begin by briefly reviewing the semantic belief-change operator proposed in [8], then we propose a syntactic belief-change operator and prove its equivalence with the semantic operator.

4.1. Semantic Operator

Agents update their possibility distribution π in light of new information $\phi \in \mathcal{L}$ coming from a source trusted to a certain extent $\tau \in [0, 1]$ by means of the following belief change operator [8], which is formulated in terms of the semantic representation of beliefs.

Definition 4 (Belief Change Operator) *The possibility distribution π' which induces the new belief set $\mathbf{B}'$ after receiving information ϕ from a source trusted to degree τ is computed from possibility distribution π relevant to the previous belief set $\mathbf{B}$ ($\mathbf{B}' = \mathbf{B} * \frac{\tau}{\phi}$, $\pi' = \pi * \frac{\tau}{\phi}$) as follows: for all interpretation $\mathcal{I}$,*

$$\pi'(\mathcal{I}) = \begin{cases} \frac{\pi(\mathcal{I})}{1 - \mathbf{B}(\neg\phi)}, & \text{if } \mathcal{I} \models \phi \text{ and } \mathbf{B}(\neg\phi) < 1; \\ 1, & \text{if } \mathcal{I} \models \phi \text{ and } \mathbf{B}(\neg\phi) = 1; \\ \min\{\pi(\mathcal{I}), 1 - \tau\}, & \text{if } \mathcal{I} \not\models \phi. \end{cases} \quad (16)$$

The first case of Equation 16 is formally equivalent to the Goguen implication, which is the residuum $\Pi([\phi]) \Rightarrow \pi(\mathcal{I})$ of the t-norm used to condition the possibility distribution on the models of ϕ , here the product². The condition $\mathbf{B}(\neg\phi) < 1$ is equivalent to $\exists \mathcal{I}' : \mathcal{I}' \models \phi \Rightarrow \pi(\mathcal{I}') > 0$, i.e., $\Pi([\phi]) > 0$; likewise, the condition $\mathbf{B}(\neg\phi) = 1$ is equivalent to $\Pi([\phi]) = 0$, which implies $\pi(\mathcal{I}) = 0 \forall \mathcal{I} \models \phi$. Therefore, The second case in Equation 16 provides for the *revision* of beliefs that contradict ϕ . In general, the operator treats new information ϕ in the negative sense: being told ϕ denies the possibility of world situations where ϕ is false (third case of Equation 16). The possibility of world situations where ϕ is true may only increase due to the first case in equation 16 or revision (second case of Equation 16). If information from a fully trusted source contradicts an existing proposition that is fully believed, then revising with the above operator leads the agent to believe the more recent information and give up the oldest to restore consistency.

The belief change operator $*$ of Definition 4 is hybrid in that it uses different t-norms to condition models and countermodels of ϕ . Such hybrid operator have been studied in [5]. In particular, this operator uses the product t-norm to condition models and the min t-norm to condition countermodels. The motivation for using this particular combination is that using the min for countermodels is suitable for representing weary or conservative agents, whereas using product for models is suitable for representing open-minded or insecure agents who, if they receive new information which contradicts one among their old beliefs, start to ques-

²As one would expect, it is possible to generalize this definition to any other t-norm. However, this is not the focus of this article: here, we are interested in reformulating this particular definition in syntactic terms.

tion also their other beliefs. Of course, we are not claiming that this operator makes sense in all scenarios: other operators may be more appropriate to deal with other situations.

It has been proved [8] that the belief change operator $*$ of Definition 4 obeys a possibilistic formulation of the AGM revision rationality postulates $K*1-K*8$ [16].

It is easy to verify that the $*$ operator is a generalization of the possibilistic conditioning operator of Dubois and colleagues [14].

After recalling that the *expansion* of a crisp set of formulas K with a formula $\phi \in \mathcal{L}$ is $K + \phi = \{\psi : K \cup \{\phi\} \vdash \psi\}$, let us define the expansion of a fuzzy set of formulas $\mathbf{B}$ with a formula $\phi \in \mathcal{L}$ from a source trusted to degree τ , for all $\psi \in \mathcal{L}$, as

$$\left(\mathbf{B} + \frac{\tau}{\phi}\right)(\psi) = \max \left\{ \alpha : \left(\mathbf{B} \cup \frac{\tau}{\phi}\right)_{\alpha} \vdash \psi \right\}. \quad (17)$$

In terms of possibility distribution, this corresponds to

$$\left(\pi + \frac{\tau}{\phi}\right)(\mathcal{I}) = \min\{\pi(\mathcal{I}), \phi^{\mathcal{I}} + (1 - \phi^{\mathcal{I}})(1 - \tau)\}. \quad (18)$$

4.2. Syntactic Operator

First of all, it is reasonable to require that a syntactic belief change operator produces a new belief base $B' = B + \frac{\tau}{\phi}$, starting from B , by using only formulas that are in $\text{supp}(B)$ and ϕ .

Four examples of increasing difficulty will help us to better frame the problem of how to express the semantic belief-change operator in terms of the syntactic representation.

Example 1 Let $B = \frac{1}{p} + \frac{0.2}{p \supset q}$, and let us calculate $B' = B * \frac{0.6}{\neg q}$. The α -cuts of B are

$$B_{\alpha} = \begin{cases} \{p\}, & \text{for } 0.2 < \alpha \leq 1; \\ \{p, p \supset q\}, & \text{for } 0 < \alpha \leq 0.2. \end{cases} \quad (19)$$

Since we only know the semantic belief-change operator, we have to transform the beliefs represented by B into the corresponding possibility distribution π , by using Equation 11. The set of interpretations Ω contains the following four interpretations, listed with their cor-

responding minterms:

$$\begin{aligned} \mathcal{I}_0 &= \{p \mapsto 0, q \mapsto 0\}, & \phi_{\mathcal{I}_0} &= \neg p \wedge \neg q, \\ \mathcal{I}_1 &= \{p \mapsto 0, q \mapsto 1\}, & \phi_{\mathcal{I}_1} &= \neg p \wedge q, \\ \mathcal{I}_2 &= \{p \mapsto 1, q \mapsto 0\}, & \phi_{\mathcal{I}_2} &= p \wedge \neg q, \\ \mathcal{I}_3 &= \{p \mapsto 1, q \mapsto 1\}, & \phi_{\mathcal{I}_3} &= p \wedge q. \end{aligned}$$

Therefore,

$$\begin{aligned} \pi(\mathcal{I}_0) &= 1 - \max\{\alpha : B_{\alpha} \vdash p \vee q\} = 0, \\ \pi(\mathcal{I}_1) &= 1 - \max\{\alpha : B_{\alpha} \vdash q \supset p\} = 0, \\ \pi(\mathcal{I}_2) &= 1 - \max\{\alpha : B_{\alpha} \vdash p \supset q\} = 0.8, \\ \pi(\mathcal{I}_3) &= 1 - \max\{\alpha : B_{\alpha} \vdash \neg p \vee \neg q\} = 1. \end{aligned}$$

Now, we apply the semantic belief-change operator to obtain $\pi' = \pi * \frac{0.6}{\neg q}$, by keeping in mind that $\mathbf{B}(q) = 1 - 0.8 = 0.2$:

$$\begin{aligned} \pi'(\mathcal{I}_0) &= \frac{\pi(\mathcal{I}_0)}{1 - \mathbf{B}(q)} = \frac{0}{0.8} = 0, \\ \pi'(\mathcal{I}_1) &= \min\{\pi(\mathcal{I}_1), 0.4\} = 0, \\ \pi'(\mathcal{I}_2) &= \frac{\pi(\mathcal{I}_2)}{1 - \mathbf{B}(q)} = \frac{0.8}{0.8} = 1, \\ \pi'(\mathcal{I}_3) &= \min\{\pi(\mathcal{I}_3), 0.4\} = 0.4. \end{aligned}$$

From possibility distribution π' we may compute the degree to which all relevant formulas are now believed:

$$\begin{aligned} \mathbf{B}'(p) &= 1 - \max_{\mathcal{I} \models p} \pi'(\mathcal{I}) = 1, \\ \mathbf{B}'(p \supset q) &= 1 - \max_{\mathcal{I} \models p \supset q} \pi'(\mathcal{I}) = 0, \\ \mathbf{B}'(\neg q) &= 1 - \max_{\mathcal{I} \models \neg q} \pi'(\mathcal{I}) = 0.6. \end{aligned}$$

Therefore, we may conclude that a reasonable candidate for B' should be the fuzzy set $\frac{1}{p} + \frac{0.6}{\neg q}$. Indeed, it is easy to verify that its corresponding possibility distribution is π' , as expected: the α -cuts of B' are

$$B'_{\alpha} = \begin{cases} \{p\}, & \text{for } 0.6 < \alpha \leq 1, \\ \{p, \neg q\}, & \text{for } 0 < \alpha \leq 0.6, \end{cases} \quad (20)$$

and applying Equation 11 to compute the corresponding possibility distribution yields π' .

What this example tells us is that, in this case, the belief base has been changed by removing formula $p \supset q$, which contradicts the new information $\neg q$ and was believed to degree 0.2 only, and by adding the new formula $\neg q$ with membership degree $\tau = 0.6$. ★

It is interesting to observe what happens when the initial belief base contains formulas that have nothing to do with incoming information. The following example is a variation of the previous one, where another formula using atom r , completely independent of both p and $p \supset q$, is introduced into the initial belief base.

Example 2 Let $B = \frac{1}{p} + \frac{0.2}{p \supset q} + \frac{0.3}{r}$, and let us calculate, just like in the previous example, $B' = B * \frac{0.6}{\neg q}$.

The α -cuts of B are now

$$B_\alpha = \begin{cases} \{p\}, & \text{for } 0.3 < \alpha \leq 1; \\ \{p, r\}, & \text{for } 0.2 < \alpha \leq 0.3; \\ \{p, p \supset q\}, & \text{for } 0 < \alpha \leq 0.2. \end{cases} \quad (21)$$

The set of interpretations Ω contains the following eight interpretations, listed with their corresponding minterms:

$$\begin{aligned} \mathcal{I}_0 &= \{p \mapsto 0, q \mapsto 0, r \mapsto 0\}, & \phi_{\mathcal{I}_0} &= \neg p \wedge \neg q \wedge \neg r, \\ \mathcal{I}_1 &= \{p \mapsto 0, q \mapsto 1, r \mapsto 0\}, & \phi_{\mathcal{I}_1} &= \neg p \wedge q \wedge \neg r, \\ \mathcal{I}_2 &= \{p \mapsto 1, q \mapsto 0, r \mapsto 0\}, & \phi_{\mathcal{I}_2} &= p \wedge \neg q \wedge \neg r, \\ \mathcal{I}_3 &= \{p \mapsto 1, q \mapsto 1, r \mapsto 0\}, & \phi_{\mathcal{I}_3} &= p \wedge q \wedge \neg r, \\ \mathcal{I}_4 &= \{p \mapsto 0, q \mapsto 0, r \mapsto 1\}, & \phi_{\mathcal{I}_4} &= \neg p \wedge \neg q \wedge r, \\ \mathcal{I}_5 &= \{p \mapsto 0, q \mapsto 1, r \mapsto 1\}, & \phi_{\mathcal{I}_5} &= \neg p \wedge q \wedge r, \\ \mathcal{I}_6 &= \{p \mapsto 1, q \mapsto 0, r \mapsto 1\}, & \phi_{\mathcal{I}_6} &= p \wedge \neg q \wedge r, \\ \mathcal{I}_7 &= \{p \mapsto 1, q \mapsto 1, r \mapsto 1\}, & \phi_{\mathcal{I}_7} &= p \wedge q \wedge r. \end{aligned}$$

Applying Equation 11 yields the possibility distribution π corresponding to belief base B :

$$\begin{aligned} \pi(\mathcal{I}_0) &= 1 - \max\{\alpha : B_\alpha \vdash p \vee q \vee r\} = 0, \\ \pi(\mathcal{I}_1) &= 1 - \max\{\alpha : B_\alpha \vdash p \vee \neg q \vee r\} = 0, \\ \pi(\mathcal{I}_2) &= 1 - \max\{\alpha : B_\alpha \vdash \neg p \vee q \vee r\} = 0.7, \\ \pi(\mathcal{I}_3) &= 1 - \max\{\alpha : B_\alpha \vdash \neg p \vee \neg q \vee r\} = 0.7, \\ \pi(\mathcal{I}_4) &= 1 - \max\{\alpha : B_\alpha \vdash p \vee q \vee \neg r\} = 0, \\ \pi(\mathcal{I}_5) &= 1 - \max\{\alpha : B_\alpha \vdash p \vee \neg q \vee \neg r\} = 0, \\ \pi(\mathcal{I}_6) &= 1 - \max\{\alpha : B_\alpha \vdash \neg p \vee q \vee \neg r\} = 0.8, \\ \pi(\mathcal{I}_7) &= 1 - \max\{\alpha : B_\alpha \vdash \neg p \vee \neg q \vee \neg r\} = 1. \end{aligned}$$

Now, we apply the semantic belief-change operator to obtain $\pi' = \pi * \frac{0.6}{\neg q}$, by keeping in mind that $\mathbf{B}(q) = 1 - 0.8 = 0.2$:

$$\begin{aligned} \pi'(\mathcal{I}_0) &= \frac{\pi(\mathcal{I}_0)}{1 - \mathbf{B}(q)} = \frac{0}{0.8} = 0, \\ \pi'(\mathcal{I}_1) &= \min\{\pi(\mathcal{I}_1), 0.4\} = 0, \\ \pi'(\mathcal{I}_2) &= \frac{\pi(\mathcal{I}_2)}{1 - \mathbf{B}(q)} = \frac{0.7}{0.8} = 0.875, \\ \pi'(\mathcal{I}_3) &= \min\{\pi(\mathcal{I}_3), 0.4\} = 0.4, \\ \pi'(\mathcal{I}_4) &= \frac{\pi(\mathcal{I}_4)}{1 - \mathbf{B}(q)} = \frac{0}{0.8} = 0, \\ \pi'(\mathcal{I}_5) &= \min\{\pi(\mathcal{I}_5), 0.4\} = 0, \\ \pi'(\mathcal{I}_6) &= \frac{\pi(\mathcal{I}_6)}{1 - \mathbf{B}(q)} = \frac{0.8}{0.8} = 1, \\ \pi'(\mathcal{I}_7) &= \min\{\pi(\mathcal{I}_7), 0.4\} = 0.4. \end{aligned}$$

From π' , we compute the degree to which all relevant formulas are now believed:

$$\begin{aligned} \mathbf{B}'(p) &= 1 - \max_{\mathcal{I} \models p} \pi'(\mathcal{I}) = 1, \\ \mathbf{B}'(p \supset q) &= 1 - \max_{\mathcal{I} \models p \supset q} \pi'(\mathcal{I}) = 0, \\ \mathbf{B}'(r) &= 1 - \max_{\mathcal{I} \models r} \pi'(\mathcal{I}) = 0.125, \\ \mathbf{B}'(\neg q) &= 1 - \max_{\mathcal{I} \models \neg q} \pi'(\mathcal{I}) = 0.6. \end{aligned}$$

Therefore, we may conclude that B' must be the fuzzy set $\frac{1}{p} + \frac{0.6}{\neg q} + \frac{0.125}{r}$. We may verify that its corresponding possibility distribution is π' .

What this example tells us is that formulas like r , that are believed to a larger degree than the negation of new information, must be “weakened” to a certain extent by the arrival of new information that contradicts current beliefs. The amount of such weakening appears to be inversely proportional to the degree to which new information contradicts current beliefs. In the case of r ,

$$\begin{aligned} B'(r) &= 1 - \frac{1 - B(r)}{1 - \mathbf{B}(q)} = 1 - \frac{1 - 0.3}{1 - 0.2} \\ &= 1 - 0.875 = 0.125. \end{aligned}$$

★

Things get even more complicated with the following example, which demonstrates the need of including “novel” formulas created as a disjunction of previously held beliefs and incoming information that contradicts current beliefs, to replace formulas removed from the initial belief base.

Example 3 Let $B = \frac{0.6}{\neg p} + \frac{0.4}{q}$, and let us calculate $B' = B * \frac{0.2}{p}$. The α -cuts of B are

$$B_\alpha = \begin{cases} \emptyset, & \text{for } 0.6 < \alpha \leq 1; \\ \{\neg p\}, & \text{for } 0.4 < \alpha \leq 0.6; \\ \{\neg p, q\}, & \text{for } 0 < \alpha \leq 0.4. \end{cases} \quad (22)$$

The set of interpretations Ω contains the following four interpretations, listed with their corresponding minterms:

$$\begin{aligned} \mathcal{I}_0 &= \{p \mapsto 0, q \mapsto 0\}, & \phi_{\mathcal{I}_0} &= \neg p \wedge \neg q, \\ \mathcal{I}_1 &= \{p \mapsto 0, q \mapsto 1\}, & \phi_{\mathcal{I}_1} &= \neg p \wedge q, \\ \mathcal{I}_2 &= \{p \mapsto 1, q \mapsto 0\}, & \phi_{\mathcal{I}_2} &= p \wedge \neg q, \\ \mathcal{I}_3 &= \{p \mapsto 1, q \mapsto 1\}, & \phi_{\mathcal{I}_3} &= p \wedge q. \end{aligned}$$

Applying Equation 11 yields the possibility distribution π corresponding to belief base B :

$$\begin{aligned}\pi(\mathcal{I}_0) &= 1 - \max\{\alpha : B_\alpha \vdash p \vee q\} = 0.6, \\ \pi(\mathcal{I}_1) &= 1 - \max\{\alpha : B_\alpha \vdash q \supset p\} = 1, \\ \pi(\mathcal{I}_2) &= 1 - \max\{\alpha : B_\alpha \vdash p \supset q\} = 0.4, \\ \pi(\mathcal{I}_3) &= 1 - \max\{\alpha : B_\alpha \vdash \neg p \vee \neg q\} = 0.4.\end{aligned}$$

Now, we apply the semantic belief-change operator to obtain $\pi' = \pi * \frac{0.2}{p}$, by keeping in mind that $\mathbf{B}(\neg p) = 0.6$:

$$\begin{aligned}\pi'(\mathcal{I}_0) &= \min\{\pi(\mathcal{I}_0), 0.8\} = 0.6, \\ \pi'(\mathcal{I}_1) &= \min\{\pi(\mathcal{I}_1), 0.8\} = 0.8, \\ \pi'(\mathcal{I}_2) &= \frac{\pi(\mathcal{I}_2)}{1 - \mathbf{B}(\neg p)} = \frac{0.4}{0.4} = 1, \\ \pi'(\mathcal{I}_3) &= \frac{\pi(\mathcal{I}_3)}{1 - \mathbf{B}(\neg p)} = \frac{0.4}{0.4} = 1.\end{aligned}$$

From π' , we may compute:

$$\begin{aligned}\mathbf{B}'(p) &= 1 - \max_{\mathcal{I} \not\models p} \pi'(\mathcal{I}) = 0.2, \\ \mathbf{B}'(\neg p) &= 1 - \max_{\mathcal{I} \models \neg p} \pi'(\mathcal{I}) = 0, \\ \mathbf{B}'(q) &= 1 - \max_{\mathcal{I} \not\models q} \pi'(\mathcal{I}) = 0.\end{aligned}$$

However, belief base B' cannot be $\frac{0.2}{p}$, as one would expect based on the two previous examples, because that would give

$$1 - \max\{\alpha : B_\alpha \vdash p \vee q\} = 0.8 \neq \pi'(\mathcal{I}_0).$$

The only way of obtaining the correct result by using only formulas that either were in the initial base or in incoming information is by observing that $\mathbf{B}'(p \vee q) = 0.4$ and thus by letting $B' = \frac{0.2}{p} + \frac{0.4}{p \vee q}$.

In other words, formula q that had to be given up because it was not believed more than $\neg p$ must be replaced by its disjunction with incoming information p , giving $p \vee q$, with the same degree of membership as q in the initial base.

Why did not this feature show up in the previous examples? The reason is that the disjunction with incoming information would give a tautology, which is always fully believed: its inclusion in B' would be redundant and could thus be dispensed with. ★

From the last example, it is not completely clear whether a formula of the form $\psi \vee \phi$, where ϕ is the incoming formula, must be added to replace formulas that got deleted only, or it must be added for all formulas ψ in the initial base. The following example helps us to clarify this issue.

Example 4 Let us use the same Ω as in the previous example, and, given

$$B = \frac{0.9}{p \supset q} + \frac{0.7}{p \vee q} + \frac{0.7}{q} + \frac{0.25}{\neg(p \wedge q)} + \frac{0.25}{p \oplus q} + \frac{0.25}{\neg p} + \frac{0.25}{\neg p \wedge q},$$

where $p \oplus q \equiv (\neg p \wedge q) \vee (p \wedge \neg q)$ represents the “exclusive or” logical operator, let us calculate $B' = B * \frac{0.6}{p}$.

The possibility distribution π corresponding to belief base B may be determined as usual, by applying Equation 11:

$$\begin{aligned}\pi(\mathcal{I}_0) &= 1 - \max\{\alpha : B_\alpha \vdash p \vee q\} = 0.3, \\ \pi(\mathcal{I}_1) &= 1 - \max\{\alpha : B_\alpha \vdash q \supset p\} = 1, \\ \pi(\mathcal{I}_2) &= 1 - \max\{\alpha : B_\alpha \vdash p \supset q\} = 0.1, \\ \pi(\mathcal{I}_3) &= 1 - \max\{\alpha : B_\alpha \vdash \neg p \vee \neg q\} = 0.75.\end{aligned}$$

Now, $\mathbf{B}(\neg p) = 0.25$ and the semantic belief change yields

$$\begin{aligned}\pi'(\mathcal{I}_0) &= \min\{\pi(\mathcal{I}_0), 0.4\} = 0.3, \\ \pi'(\mathcal{I}_1) &= \min\{\pi(\mathcal{I}_1), 0.4\} = 0.4, \\ \pi'(\mathcal{I}_2) &= \frac{\pi(\mathcal{I}_2)}{1 - \mathbf{B}(\neg p)} = \frac{0.1}{0.75} = 0.1333, \\ \pi'(\mathcal{I}_3) &= \frac{\pi(\mathcal{I}_3)}{1 - \mathbf{B}(\neg p)} = \frac{0.75}{0.75} = 1,\end{aligned}$$

from which we must infer that

$$B' = \frac{0.8667}{p \supset q} + \frac{0.7}{p \vee q} + \frac{0.6}{p}.$$

Notice that B' must be obtained by adding formulas $\psi \vee p$ for all the formulas originally in B to the contracted belief base. What happens then is that all of those $\psi \vee p$ either reduce to $p \vee q$ or to a tautology:

$$\begin{aligned}(p \supset q) \vee p &= \neg p \vee q \vee p = \top \vee q = \top, \\ (p \vee q) \vee p &= p \vee q, \\ \neg(p \wedge q) \vee p &= \neg p \vee \neg q \vee p = \top \vee \neg q = \top, \\ (p \oplus q) \vee p &= (\neg p \wedge q) \vee (p \wedge \neg q) \vee p = p \vee q, \\ \neg p \vee p &= \top, \\ (\neg p \wedge q) \vee p &= (\neg p \vee p) \wedge (p \vee q) = \top \wedge (p \vee q) \\ &= p \vee q.\end{aligned}$$

Furthermore, there is no need to include $\frac{0.6}{q}$ in B' , since q may be deduced from p and $p \supset q$, both members of $B'_{0.6}$.

Had not $\frac{0.7}{p \vee q}$ been included in B' , the term $\frac{0.6}{p \vee q}$, resulting from $B'(p \vee q) = 1 - \frac{0.3}{0.75}$, would have been there instead, giving

$$1 - \max\{\alpha : B_\alpha \vdash p \vee q\} = 0.4 \neq \pi'(\mathcal{I}_0).$$

This means that $\psi \vee \phi$ must be included in B' for all formulas ψ in B . ★

One of the oldest ideas in belief revision theory, due to Isaac Levi, is that changing your beliefs given some new information should proceed in two steps: first, give up old beliefs contradicted by new information; second, add the new information [19].

The syntactic expansion operator $+$, equivalent to the semantic expansion operator of Equation 18 may be defined as

$$B + \frac{\tau}{\phi} = B \cup \frac{\tau}{\phi}. \quad (23)$$

Notice that, if ϕ is already in B at least to degree τ , $B \cup \frac{\tau}{\phi} = B$.

As for the syntactic contraction operator, it is not obvious how it should be defined. A first guess at it might be what we may call *simple* contraction,

$$B - \phi = B \cap B_{>\mathbf{B}(\phi)}, \quad (24)$$

i.e., B with all formulas belonging to it up to the degree to which ϕ is believed removed. However, for reasons that should be clear from our discussion of the examples, we prefer to resort to the following *normalized* contraction operator: for all $\psi \in \mathcal{L}$,

$$(B \dot{-} \phi)(\psi) = \begin{cases} 0, & \text{if } B(\psi) \leq \mathbf{B}(\phi), \\ 1 - \frac{1-B(\psi)}{1-\mathbf{B}(\phi)}, & \text{if } B(\psi) > \mathbf{B}(\phi), \end{cases} \quad (25)$$

which is almost the same as Equation 24, except that it redistributes between 0 and 1 the degree of membership of all the formulas that do not get removed.

A syntactic belief-change operator may now be defined as a two-stage operator, which first contracts the belief base by the negation of the information communicated by a partially trusted source, and then expands it with such new information.

Definition 5 (Belief Change Operator) *The belief base B' which induces the new belief set $\mathbf{B}'$ after re-*

Fig. 2. A flow chart of the syntactic belief change operator, showing how $B(\psi)$, given as input, is changed into $B'(\psi)$.

ceiving information ϕ from a source trusted to degree τ is computed from belief base B relevant to the previous belief set $\mathbf{B}$ ($\mathbf{B}' = \mathbf{B} * \frac{\tau}{\phi}$, $B' = B * \frac{\tau}{\phi}$) as follows:

$$B' = (B \dot{-} \neg\phi) + \frac{\tau}{\phi} + \sum_{\psi} \frac{B(\psi)}{\phi \vee \psi}. \quad (26)$$

An interesting feature of the above definition is that Equation 26 may be regarded as a fuzzy generalization of the well-known Levi identity $K * \phi = (K - \neg\phi) + \phi$.

Definition 5 may be translated into the following procedure:

1. remove from belief base B all formulas ψ such that $B(\psi) \leq \mathbf{B}(\neg\phi)$;
2. redistribute the degrees of membership of all extant formulas ψ , now ranging in $(\mathbf{B}(\neg\phi), 1]$, by remapping them to the full $(0, 1]$ interval: $B'(\psi) = 1 - \frac{1-B(\psi)}{1-\mathbf{B}(\neg\phi)}$;
3. add formula ϕ with membership degree τ to the resulting contracted belief base;
4. finally, for all formula ψ , add $\phi \vee \psi$ with membership degree $B(\psi)$; of course, only formulas with a non-zero membership degree in the initial belief base will have to be considered.

An alternative way of writing Equation 26, which emphasizes the way individual membership degrees

change, is, for all $\psi \in \mathcal{L}$,

$$B'(\psi) = \begin{cases} \max\{\tau, B(\phi)\}, & \psi = \phi, \\ \max\{B(\xi), B(\phi)\}, & \psi = \phi \vee \xi, \\ 0, & \psi \neq \phi, \psi \neq \phi \vee \xi, \\ & B(\psi) \leq \mathbf{B}(\neg\phi), \\ 1 - \frac{1 - B(\psi)}{1 - \mathbf{B}(\neg\phi)}, & \psi \neq \phi, \psi \neq \phi \vee \xi, \\ & B(\psi) > \mathbf{B}(\neg\phi). \end{cases} \quad (27)$$

Notice the striking formal symmetry between this equation and Equation 16. The flow chart shown in Figure 2 makes it clear how the membership values of formulas in the belief base change as the consequence of applying Equation 27.

In the special case where $\mathbf{B}(\neg\phi) = 1$ and $\tau = 1$, $B * \frac{\tau}{\phi} = \phi$. When $\mathbf{B}(\neg\phi) = 0$, $B * \frac{\tau}{\phi} = B + \frac{\tau}{\phi}$, i.e., belief change reduces to simple expansion.

4.3. Equivalence

Equivalence between the semantic and syntactic operators will be proved by showing that the syntactic operator applied to a belief base B corresponding to possibility distribution π yields a revised belief base B' whose corresponding possibility distribution is exactly the distribution π' one would obtain by applying the semantic operator to distribution π . This is summarized in the following commutative diagram:

$$\begin{array}{ccc} \pi *_{\text{sem}} \frac{\tau}{\phi} & \rightarrow & \pi' \\ \uparrow & & \uparrow \\ B *_{\text{syn}} \frac{\tau}{\phi} & \rightarrow & B' \end{array} \quad (28)$$

where $*_{\text{sem}}$ symbolizes the semantic belief-change operator of Definition 4 and $*_{\text{syn}}$ stands for the syntactic belief-change operator of Definition 5.

Theorem 1 *Let B be a belief base and π_B its corresponding possibility distribution. For all $\phi \in \mathcal{L}$ and $\tau \in (0, 1]$, let $B' = B * \frac{\tau}{\phi}$ and $\pi_{B'}$ be the possibility distribution corresponding to B' . Then,*

$$\pi_{B'} = \pi_B * \frac{\tau}{\phi}. \quad (29)$$

Proof: According to Equation 10, for all $\mathcal{I}$,

$$\pi_B(\mathcal{I}) = 1 - \max_{\psi: \mathcal{I} \not\models \psi} B(\psi). \quad (30)$$

For convenience of notation, let $\pi'_B = \pi_B * \frac{\tau}{\phi}$. Applying the semantic belief-change operator to π_B yields, by Definition 4, for all $\mathcal{I}$,

$$\pi'_B(\mathcal{I}) = \begin{cases} \frac{1 - \max_{\psi: \mathcal{I} \not\models \psi} B(\psi)}{1 - \mathbf{B}(\neg\phi)}, & \mathcal{I} \models \phi, \\ & \mathbf{B}(\neg\phi) < 1; \\ 1, & \mathcal{I} \models \phi, \\ & \mathbf{B}(\neg\phi) = 1; \\ \min \left\{ 1 - \max_{\psi: \mathcal{I} \not\models \psi} B(\psi), 1 - \tau \right\}, & \mathcal{I} \not\models \phi. \end{cases} \quad (31)$$

On the other hand, according to Equation 10, for all $\mathcal{I}$,

$$\pi_{B'}(\mathcal{I}) = 1 - \max_{\psi: \mathcal{I} \not\models \psi} B'(\psi). \quad (32)$$

Therefore, in order to prove that $\pi'_B = \pi_{B'}$, we must prove the following three theses, corresponding to each of the three conditions of Equation 31: for all $\mathcal{I}$,

1. if $\mathcal{I} \models \phi$ and $\mathbf{B}(\neg\phi) < 1$,

$$1 - \max_{\psi: \mathcal{I} \not\models \psi} B'(\psi) = \frac{1 - \max_{\psi: \mathcal{I} \not\models \psi} B(\psi)}{1 - \mathbf{B}(\neg\phi)}; \quad (33)$$

2. if $\mathcal{I} \models \phi$ and $\mathbf{B}(\neg\phi) = 1$,

$$\max_{\psi: \mathcal{I} \not\models \psi} B'(\psi) = 0; \quad (34)$$

3. finally, if $\mathcal{I} \not\models \phi$,

$$1 - \max_{\psi: \mathcal{I} \not\models \psi} B'(\psi) = \min \left\{ 1 - \max_{\psi: \mathcal{I} \not\models \psi} B(\psi), 1 - \tau \right\}. \quad (35)$$

To prove Thesis 1, it suffices to substitute the left-hand side of Equation 33 with its definition:

$$\begin{aligned} 1 - \max_{\psi: \mathcal{I} \not\models \psi} B'(\psi) &= 1 - \max_{\psi: \mathcal{I} \not\models \psi} \left\{ 1 - \frac{1 - B(\psi)}{1 - \mathbf{B}(\neg\phi)} \right\} = \\ &= \min_{\psi: \mathcal{I} \not\models \psi} \frac{1 - B(\psi)}{1 - \mathbf{B}(\neg\phi)} = \\ &= \frac{1}{1 - \mathbf{B}(\neg\phi)} \min_{\psi: \mathcal{I} \not\models \psi} \{1 - B(\psi)\} = \\ &= \frac{1 - \max_{\psi: \mathcal{I} \not\models \psi} B(\psi)}{1 - \mathbf{B}(\neg\phi)}. \end{aligned}$$

To prove Thesis 2, we simply observe that since, by hypothesis, $\mathbf{B}(\neg\phi) = 1$, $B(\psi) \leq \mathbf{B}(\neg\phi)$ for all ψ ;

therefore, for all $\psi \neq \phi$, $B'(\psi) = 0$ (third case of Equation 27).

As for Thesis 3, Equation 35 may be rewritten as

$$1 - \max_{\psi: \mathcal{I} \not\models \psi} B'(\psi) = 1 - \max \left\{ \max_{\psi: \mathcal{I} \not\models \psi} B(\psi), \tau \right\}.$$

Therefore, we have to prove that, if $\mathcal{I} \not\models \phi$,

$$\max_{\psi: \mathcal{I} \not\models \psi} B'(\psi) = \max_{\psi: \mathcal{I} \not\models \psi} \{B(\psi), \tau\}. \quad (36)$$

Notice that ϕ is among the formulas not satisfied by $\mathcal{I}$ and, by the first case of Equation 27, $\max_{\psi: \mathcal{I} \not\models \psi} B'(\psi) \geq B'(\phi) \geq \tau$ and $\max_{\psi: \mathcal{I} \not\models \psi} B(\psi) \geq B(\phi)$.

Let us concentrate on the most believed formula ψ^* among the formulas not satisfied by $\mathcal{I}$,

$$\psi^* = \arg \max_{\psi: \mathcal{I} \not\models \psi} \{B(\psi), \tau\}.$$

To begin with, we remark that ψ^* might be ϕ . However, if that were the case, we know that $B'(\phi) = \max\{B(\phi), \tau\}$ (first case of Equation 27).

Now, if $\psi^* \neq \phi$, either $B(\psi^*) \leq B(\neg\phi)$, in which case it is removed from the belief base, or $B(\psi^*) > B(\neg\phi)$, in which case it is kept, but its degree of membership is decreased by the normalized contraction operator (Step 2 of the procedure). In any case, $B'(\psi^*) \leq B(\psi^*)$ and, *a fortiori*, $B'(\psi^*) \leq \max\{B(\psi^*), \tau\}$. However, Step 4 of the procedure or, equivalently, the second case of Equation 27 guarantees that a new formula, namely $\psi^* \vee \phi$, is then introduced into the new belief base B' , with the same membership degree as ψ^* in the original base B :

$$B'(\psi^* \vee \phi) = B(\psi^*).$$

Furthermore, it must be $B(\psi^*) \geq \tau$, otherwise we could choose $\psi^* = \phi$ (see remark above). Therefore,

$$B'(\psi^* \vee \phi) = \max\{B(\psi^*), \tau\}.$$

Since the same happens to all formulas ψ not satisfied by $\mathcal{I}$, it must be that $B'(\psi^* \vee \phi) = \max_{\psi: \mathcal{I} \not\models \psi} B'(\psi)$.

Therefore, Equation 36 must necessarily hold. $\square$

Fig. 3. A screenshot of KOBDIG's user interface: the user is entering input $\frac{0.6}{p}$ to revise the initial belief base B of Example 4.2.

Fig. 4. A screenshot of KOBDIG's user interface: the belief tab of the agent's mental state pane shows the result of revising belief base B of Example 4.2 by $\frac{0.6}{p}$.

5. Experiments

A Java implementation of the syntactic belief change operator described in Section 4 is available in the open-source KOBDIG project, hosted on SourceForge³.

Figure 3 and 4 show two screenshots of KOBDIG's user interface while computing Example 4.2 above.

The Java code of method `updateBeliefs` of class `Agent` in KOBDIG, which implements the belief-change operator, is given in Figure 5. Of all the operations performed by this method, the most computation-intensive is by far the calculation of the degree of belief in a formula, carried out by method `PossibilisticFactBase.necessity`, which implements Equation 7 by checking, for each $\alpha \in \Lambda(B)$, where $\Lambda(B)$ is the level set of B , i.e., the set of non-null membership degrees α such that at least

³Project URL: <http://sourceforge.net/p/kobdig/>.

```

public void updateBeliefs(Fact fact, TruthDegree trust) {
    // First of all, compute the degree to which the incoming
    // fact contradicts the agent's current beliefs:
    TruthDegree contradiction = beliefs.necessity(fact.negated());

    PossibilisticFactBase newBeliefs = new PossibilisticFactBase();

    // Insert into the revised base only those facts whose degree
    // of membership is greater than the "contradiction":
    Iterator<Fact> i = beliefs.factIterator();
    while(i.hasNext()) {
        Fact psi = i.next();
        TruthDegree t = beliefs.membership(psi);
        if(!contradiction.isAtLeastAsTrueAs(t)) {
            // Redistribute the degrees of membership:
            TruthDegree tcond = new TruthDegree(1.0 -
                t.negated().doubleValue() /
                contradiction.negated().doubleValue());
            newBeliefs.tell(psi, tcond);
        }
    }

    // For each fact in the base, add a new fact whose formula is
    // the disjunction of the existing fact and the incoming fact:
    i = beliefs.factIterator();
    while(i.hasNext()) {
        Fact psi = i.next();
        TruthDegree t = beliefs.necessity(psi);
        Fact disjunction = new Fact(new PropositionalFormula(
            Operator.OR, psi.formula(), fact.formula()));
        if(!newBeliefs.necessity(disjunction).isAtLeastAsTrueAs(t))
            newBeliefs.tell(disjunction, t);
    }

    // Add the incoming fact to the base
    // with the same membership degree as its degree of trust:
    if(!newBeliefs.necessity(fact).isAtLeastAsTrueAs(trust))
        newBeliefs.tell(fact, trust);

    // Done! Now, replace the original beliefs:
    beliefs = newBeliefs;
    // Even if we checked not to include redundant facts,
    // depending on the order they were inserted, some
    // may still be there. Remove them:
    beliefs.simplify();
}

```

Fig. 5. Java code of the belief-change algorithm.

one element belongs in B to degree α , whether B_α entails the formula. Checking whether a set of formulas entails a given formula is a logical reasoning problem which may be reduced to the *satisfiability problem*, whose computational complexity varies depending on the specific logic considered, but is in general quite high. For instance, satisfiability in propositional logic (also known as Boolean satisfiability) is NP-complete [17]; concept satisfiability in description logics goes from polynomial to NEXPTIME-complete [2]. Therefore, counting the number of entailment checks that have to be performed in order to apply the proposed syntactic belief change operator provides a faithful estimate of its computational resource requirements.

We study experimentally the performance of the proposed syntactic belief change operator by the following protocol:

1. given a propositional language $\mathcal{L}_n$ with n atomic proposition, $n = 1, 2, \dots$;

RND_WFF(β)

- With probability β :
 1. randomly choose $op \in \{\neg, \wedge, \vee\}$;
 2. if $op = \neg$, return $\neg \text{RND_WFF}(\beta)$;
else return $(\text{RND_WFF}(\beta) \text{ op } \text{RND_WFF}(\beta))$;
- else, return a random atomic proposition from $\mathcal{L}_n$.

Fig. 6. The algorithm used to generate random propositional formulas from $\mathcal{L}_n$.

2. start out with an empty belief base $B_0 = \emptyset$;
3. generate random formulas $\phi_i \in \mathcal{L}_n$, $i = 1, 2, \dots$, with the recursive algorithm shown in Figure 6, which takes a parameter $0 < \beta < 1$ allowing us to control the expected syntactic depth of the formula;
4. generate a random trust degree τ_i , $i = 1, 2, \dots$;
5. compute base B_i , $i = 1, 2, \dots$, as $B_i = B_{i-1} * \frac{\tau_i}{\phi_i}$, by applying the syntactic belief change operator.

Figure 7 shows how the number of entailment checks grows as a function of i for different values of n , while Figure 8 shows the same for $\|\text{supp}(B_i)\|$, i.e., the size of the belief base. The random formulas used to change beliefs were generated with $\beta = \frac{1}{2}$. Since the experiments are random in nature, ten independent trials have been performed for each value of n ; each box-and-whiskers plot in the two figures represents a distribution over those trials.

It is quite evident, from these diagrams, that both the number of entailment checks required to perform a belief change and the size of the belief base grow less than linearly with respect to the number of formulas acquired by an agent. In fact, the entailment checks appear to stabilize around an average after a number of belief changes, which increases with the size of the language.

It is interesting to notice that the number of entailment checks required to perform a belief change appears to be somehow related to $n\|\text{supp}(B)\|$. In fact, the data plotted in Figure 9 are fitted by the linear model

$$\# \text{checks} \simeq 2n\|\text{supp}(B)\|,$$

shown as a solid line in Figure 9.

This may be taken as evidence that adopting a syntactic possibilistic representation of beliefs together with the proposed syntactic operator does not increase

Fig. 7. Number of entailment checks over a number of subsequent applications i of the syntactic belief change operator for different sizes n of the propositional language.

Fig. 8. Cardinality of the belief base over a number of subsequent applications i of the syntactic belief change operator for different sizes n of the propositional language.

Fig. 9. Scatter plot of the number of entailment checks required to perform a belief change against $n ||\text{supp}(B)||$.

the complexity of belief change with respect to a classical setting where beliefs are not graded and uncertainty/degree of trust of information sources is not taken into account.

6. Conclusion

Starting from a possibilistic belief-change operator proposed in the literature which was defined in terms of a semantic representation of beliefs, we have devised its syntactic counterpart, i.e., a possibilistic belief-change operator that works on a syntactic representation of beliefs in the form of a fuzzy belief base, and that behaves exactly like the semantic operator. Furthermore, we have performed an empirical study of the syntactic operator which shows that the cost of performing a revision, as well as the size of the resulting belief base, remain tractable even after many pieces of information have been received.

The advantages of being able to work directly on a syntactic representation should be quite obvious: a semantic representation of beliefs in terms of a possibility distribution is well-suited to a theoretical treatment of the matter, but does not constitute a serious candidate for an implementation of any cognitive agent framework adopting possibility theory for reasoning about beliefs. Implementations with any ambition to tackle real-world applications will have necessarily to resort to a syntactic representation, and thus will have to use a belief-change operator that works directly on that representation.

References

- [1] C. E. Alchourrón, P. Gärdenfors, and D. Makinson. On the logic of theory change: Partial meet contraction and revision functions. *J. Symb. Log.*, 50(2):510–530, 1985.
- [2] Franz Baader, Diego Calvanese, Deborah McGuinness, Daniele Nardi, and Peter Patel-Schneider, editors. *The Description Logic Handbook: Theory, implementation and applications*. Cambridge, 2003. ISBN 0521781760.
- [3] S. Benferhat and S. Kaci. Logical representation and fusion of prioritized information based on guaranteed possibility measures: application to the distance-based merging of classical bases. *Artif. Intell.*, 148(1-2):291–333, 2003. ISSN 0004-3702.
- [4] S. Benferhat, D. Dubois, H. Prade, and M-A. Williams. A practical approach to revising prioritized knowledge bases. *Studia Logica*, 70(1):105–130, 2002.
- [5] S. Benferhat, C. da Costa Pereira, and A. Tettamanzi. Hybrid possibilistic conditioning for revision under weighted inputs. In *ECAI 2012*, pages 151–156, 2012.
- [6] J. Blee, D. Billington, and A. Sattar. Reasoning with levels of modalities in BDI logic. In *Proceedings of PRIMA'07*, pages 410–415. Springer-Verlag, 2009.
- [7] A. Casali, L. Godo, and C. Sierra. Graded BDI models for agent architectures. In *CLIMA V*, pages 18–33, 2004.
- [8] C. da Costa Pereira and A. Tettamanzi. An integrated possibilistic framework for goal generation in cognitive agents. In *Proceedings of AAMAS'10*, pages 1239–1246. IFAAMAS, 2010.
- [9] C. da Costa Pereira and A. Tettamanzi. A syntactic possibilistic belief change operator for cognitive agents. In *IAT 2011*, pages 38–45, 2011.
- [10] B. De Baets, E. Tsiporkova, and R. Mesiar. Conditioning in possibility theory with strict order norms. *Fuzzy Sets Syst.*, 106(2):221–229, 1999.
- [11] A. F. Dragoni and P. Giorgini. Belief revision through the belief-function formalism in a multi-agent environment. In *ATAL*, pages 103–115, 1996.
- [12] A.F. Dragoni. Belief revision under uncertainty in a multi-agent environment. In *Knowledge-Based Systems: Advanced Concepts, Techniques and Applications*, pages 191–215. World Scientific Publishers, 1997.
- [13] D. Dubois and H. Prade. Epistemic entrenchment and possibilistic logic. *Artif. Intell.*, 50(2):223–239, 1991.
- [14] D. Dubois and H. Prade. A synthetic view of belief revision with uncertain inputs in the framework of possibility theory. *International Journal of Approximate Reasoning*, 17:295–324, 1997.
- [15] D. Dubois, J. Lang, and H. Prade. Possibilistic logic. In Dov M. Gabbay, C. J. Hogger, and J. A. Robinson, editors, *Handbook of Logic in Artificial Intelligence and Logic Programming*, pages 439–513. Oxford University Press, 1994.
- [16] P. Gärdenfors. Belief revision: A vademecum. In *Meta-Programming in Logic*, pages 1–10. Springer, Berlin, 1992.
- [17] M. R. Garey and D. S. Johnson. *Computers and Intractability: A guide to the Theory of NP-Completeness*. Freeman, New York, 1979.
- [18] S. O. Hansson. Knowledge-level analysis of belief base operations. *Artif. Intell.*, 82(1-2):215–235, 1996.
- [19] I. Levi. Subjunctives, dispositions and chances. *Synthese*, 34: 423–455, 1977.

- [20] Philippe Smets and Robert Kennes. The transferable belief model. *Artificial Intelligence*, 66, 1994.
- [21] L. A. Zadeh. Fuzzy sets. *Information and Control*, 8:338–353, 1965.
- [22] L. A. Zadeh. Fuzzy sets as a basis for a theory of possibility. *Fuzzy Sets and Systems*, 1:3–28, 1978.