

Hybrid Possibilistic Conditioning for Revision under Weighted Inputs

Salem Benferhat¹ and Célia da Costa Pereira² and Andrea G. B. Tettamanzi³

Abstract. We propose and investigate new operators in the possibilistic belief revision setting, obtained as different combinations of the conditioning operators on models and countermodels, as well as of how weighted inputs are interpreted. We obtain a family of eight operators that essentially obey the basic postulates of revision, with a few slight differences. These operators show an interesting variety of behaviors, making them suitable to representing changes in the beliefs of an agent in different contexts.

1 Introduction

Belief revision with uncertain input has been studied by Spohn [12], who has shown its close relationship with Jeffrey’s rule of conditioning [11] in probability theory, a generalization of Bayesian conditioning. Belief revision with uncertain inputs has to do with how an agent should revise its beliefs when told information either by a partially trusted source, or by a fully reliable source which provides a degree of certainty for information it reports. Possibilistic counterparts to the revision by uncertain inputs have been discussed in [6]. At the syntactic level, such a form of revision comes down to adding a formula to a belief base at a certain prescribed level of necessity. The problem is made difficult because the belief base must be modified so that the added formula maintains its prescribed priority, that is, it is neither implicitly inhibited by higher priority formulas that contradict it, nor pushed to higher priority levels by formulas that imply it. A first step towards an efficient way for doing this has been proposed in [2]. In that paper, the authors pursue the study of revision with sure and uncertain inputs within the framework of possibility theory. They start from investigating efficient syntactic implementation schemes for both belief revision and contraction in possibilistic logic and they show their full agreement with semantics. The authors also took advantage of the connections of the ordinal conditional functions framework [12] with possibility theory.

To the best of our knowledge, almost all of the existing possibilistic-based approaches are based on homogeneous operators for revision: the same operator is used to model the revision of the possibility distribution on the models (the interpretations that are consistent with the incoming information) and on the countermodels (the interpretations that contradict the incoming information). However, since the conditioning of the models and of the countermodels are independent, one might as well decide to use different operators on models and countermodels, thus obtaining hybrid possibilistic conditioning operators for belief revision. Indeed, in some applications, it may happen that a given operator provides the desired

behavior on models, but a counterintuitive one on countermodels, or viceversa.

Here, we investigate, at the level of semantic representation, the properties of such hybrid operators in comparison to standard “homogeneous” operators. The results are interesting: the new operators essentially obey the same postulates as the homogeneous operators, with a few slight differences. Taking the way weighted inputs are interpreted as an additional parameter of a possibilistic conditioning operator, we end up with eight operators (homogeneous and hybrid) showing an interesting variety of behaviors, which make them suitable to representing the changes in the agent beliefs in a variety of application contexts. More precisely, these operators may be classified according to their behavior: a first class of operators can be used when a completely trusted source told the agent that the certainty of a formula ϕ is a , ($N(\phi) = a$). In such a case, even if ϕ was believed to a higher degree, its degree of belief after revision will be exactly a . A second class of operators treats incoming information (ϕ, a) as an indication that ϕ is true provided by a source partially trusted to degree a . A third class of operators reacts to repeated confirmations of the same input ϕ by different sources by adopting belief ϕ with the same degree as the degree of trust (or reliability) of the most reliable among those sources. A fourth class of operators treats any confirmation of ϕ , even by little reliable sources, as additional evidence in favor of ϕ , thus increasing the degree of belief in ϕ . A fifth class of operators raises the possibility of all worlds that are compatible with incoming information. Finally, a sixth class of operators raises to 1 just the models of ϕ that were deemed most possible before the arrival of the input.

The paper is organized as follows: the representation of epistemic states with possibility theory is briefly presented in Section 2; Section 3 introduces iterated revision by possibilistic conditioning, starting from the simple scenario where the inputs are totally reliable, then extending it to cases in which the inputs are weighted, for which eight operators are presented. Section 4 concludes and suggests possible extensions of this work.

2 Possibilistic representations of epistemic states

Let L be a finite propositional language; \vdash denotes the classical consequence relation. Ω is the set of classical interpretations or worlds, and, given a formula $\phi \in L$, let $[\phi]$ denote the set of classical models of ϕ . Often epistemic states (or cognitive states), viewed as a set of beliefs about the real world (based on the available information), are represented by a total pre-order either on Ω , or on the set of formulas as in epistemic entrenchment relation presented in [8]. These orderings reflect the strength of the various beliefs maintained by an agent.

¹ Université d’Artois, France, email: benferhat@cril.univ-artois.fr

² Université de Nice Sophia Antipolis, France, email: celia.pereira@unice.fr

³ Università degli Studi di Milano, Italy, email: andrea.tettamanzi@unimi.it

In this section, we describe the representation of epistemic states in possibilistic logic at semantic level, where priorities are encoded by reals in the interval $[0,1]$. An equivalent representation of epistemic states at the syntactic level is possible (see, for instance, [2]), but is left out of the scope of this paper for the sake of simplicity.

2.1 Semantic representation of epistemic states

At the semantic level, an epistemic state is represented by a possibility distribution π , which is a mapping from Ω to the interval $[0,1]$. $\pi(\omega)$ represents the degree of compatibility of ω with the available information (or beliefs) about the real world. $\pi(\omega) = 0$ means that the interpretation ω is impossible, and $\pi(\omega) = 1$ means that nothing prevents ω from being the real world. The interpretations such that $\pi(\omega) = 1$ are considered as normal, expected. When $\pi(\omega) > \pi(\omega')$, ω is a preferred candidate to ω' for being the real state of the world. The less $\pi(\omega)$ the more abnormal ω is. A possibility distribution π is said to be normalized if $\exists \omega \in \Omega$, such that $\pi(\omega) = 1$.

Given a possibility distribution π , we can define two different measures on formulas of the language:

- the possibility degree $\Pi_\pi(\phi) = \max\{\pi(\omega) : \omega \in [\phi]\}$ which evaluates the extent to which ϕ is consistent with the available information expressed by π .
- the necessity degree $N_\pi(\phi) = 1 - \Pi(\neg\phi)$ which evaluates the extent to which ϕ is entailed by the available information.

When there is no ambiguity, we simply write $\Pi(\phi)$ (resp. $N(\phi)$) instead of $\Pi_\pi(\phi)$ (resp. $N_\pi(\phi)$). Note that $\Pi(\phi)$ is evaluated from the assumption that the situation where ϕ is true is as normal as can be. The duality equation $N(\phi) = 1 - \Pi(\neg\phi)$ extends the one existing in classical logic, where a formula is entailed from a set of classical formulas if and only if its negation is inconsistent with this set.

Lastly, given a possibility distribution π , the semantic determination of the belief set (corresponding to the agent's current beliefs) denoted by $BS(\pi)$, is obtained by considering all formulas which are more plausible than their negation, namely

$$BS(\pi) = \{\phi : \Pi(\phi) > \Pi(\neg\phi)\}.$$

Namely, $BS(\pi)$ is a classical base whose models are the interpretations having the highest degrees in π . When π is normalized, models of $BS(\pi)$ are interpretations which are completely possible, namely $[BS(\pi)] = \{\omega : \pi(\omega) = 1\}$. The formula ϕ belongs to $BS(\pi)$ when ϕ holds in all the most normal situations (hence ϕ is expected, or accepted as being true).

Example 1 Let π be defined as follows:

ω	$\pi(\omega)$
qr	1
$q\neg r$	1
$\neg qr$	$.7$
$\neg q\neg r$	$.5$

$BS(\pi)$ will contain formula q , because $\Pi(q) = 1 > \Pi(\neg q) = 0.7$, and also $q \vee r$, because $\Pi(q \vee r) = 1 > \Pi(\neg(q \vee r)) = 0.5$, but, for instance, not $q \wedge r$, because $\Pi(q \wedge r) = \Pi(\neg(q \wedge r)) = 1$.

3 Iterated semantic revision in possibilistic logic

The choice of a revision method partially depends on the status of the input information. We first consider revising with a totally reliable input, then we discuss the revision with an uncertain input. In the

case of uncertain information, the input is of the form (ϕ, a) which means that the classical formula ϕ should be believed to a degree of certainty a exactly. Here, uncertain input is treated according to the two views existing in the literature: (i) as a constraint which is enforced (as proposed in [2]) and (ii) by taking it into account only if it leads to a strengthening of the certainty (as proposed in [6]). Besides, for each view we consider two new possible combinations of the operators min and product: (1) the min operator for representing the models and the product for representing the countermodels (min/product) and (2) the *product* for representing the models and the min for representing the countermodels (product/min). We will denote by π the possibility distribution representing the epistemic state before the arrival of input (ϕ, a) and by π' the possibility distribution revised according to input (ϕ, a) . Accordingly, we will denote by N and Π the necessity and possibility measures induced by π and by N' and Π' the necessity and possibility measure induced by π' .

3.1 Revision with a totally reliable input

In the case of revision with a totally reliable (or certain, sure) input ϕ , it is assumed that all interpretations ω that falsify ϕ are declared impossible ($\pi(\omega) = 0$). This is performed by means of a conditioning device which transforms a possibility distribution π and a new and totally reliable information ϕ into a new possibility distribution denoted by $\pi' = \pi(\cdot \mid \phi)$. As stated in [2], natural properties for π' are the following AGM postulates [1], translated into a possibilistic setting:

- A₁**: π' should be normalized;
- A₂**: $\forall \omega \notin [\phi], \pi'(\omega) = 0$;
- A₃**: $\forall \omega, \omega' \in [\phi], \pi(\omega) > \pi(\omega')$ iff $\pi'(\omega) > \pi'(\omega')$;
- A₄**: if $N(\phi) > 0$, then $\forall \omega \in [\phi] : \pi(\omega) = \pi'(\omega)$;
- A₅**: if $\pi(\omega) = 0$, then $\pi'(\omega) = 0$.

A₁ means that the new epistemic state is consistent. **A₂** confirms that ϕ is a sure piece of information. **A₃** means that the new possibility distribution should not alter the previous relative order between models of ϕ . **A₄** means that, when $N(\phi) > 0$ (ϕ is *a priori* accepted), then revision does not affect π . **A₅** stipulates that impossible worlds remain impossible after conditioning. Then it can be verified that any revision of the belief set $BS(\pi)$ by ϕ , leading to $BS(\pi(\cdot \mid \phi))$ with $\pi(\cdot \mid \phi)$ obeying **A₁**–**A₅**, satisfies all AGM postulates.

The previous properties **A₁**–**A₅** do not guarantee a unique definition of conditioning. Moreover, the effect of axiom **A₂** may result in a sub-normalized possibility distribution. Restoring the normalization, so as to satisfy **A₁**, can be done, in principle, by choosing any continuous t-norm $*$ such that $x * x = 0$ if and only if $x = 0$, and defining, when $\Pi(\phi) > 0$,

$$\pi(\omega \mid \phi) = \begin{cases} \Pi(\phi) \Rightarrow \pi(\omega), & \text{if } \omega \models \phi, \\ 0 & \text{otherwise,} \end{cases}$$

where \Rightarrow denotes the residuum of t-norm $*$ [5]. However, we will focus here on the idempotent t-norm (i.e., min) and the product t-norm, just because these two basic operations have been widely used in a belief-revision context, thereby obtaining two different types of conditioning [7]:

- In an ordinal setting, we assign maximal possibility to the best models of ϕ , then we get:

$$\pi(\omega \mid_m \phi) = \begin{cases} 1, & \text{if } \pi(\omega) = \Pi(\phi) \text{ and } \omega \models \phi, \\ \pi(\omega), & \text{if } \pi(\omega) < \Pi(\phi) \text{ and } \omega \models \phi, \\ 0 & \text{if } \omega \not\models \phi. \end{cases}$$

This is the definition of *minimum-based conditioning*.

- In a numerical setting, we proportionally rescale all models of ϕ upwards:

$$\pi(\omega \mid \cdot, \phi) = \begin{cases} \frac{\pi(\omega)}{\Pi(\phi)}, & \text{if } \omega \models \phi, \\ 0, & \text{otherwise.} \end{cases}$$

This is the definition of *product-based conditioning*.

These two revision methods satisfy an equation of the form

$$\forall \omega, \quad \pi(\omega) = \pi(\omega \mid \phi) * \Pi(\phi),$$

which is similar to Bayesian conditioning, where $*$ may stand for min and the product respectively. The rule based on the product is much closer to genuine Bayesian conditioning than the qualitative conditioning defined from the minimum which is purely based on comparing levels; product-based conditioning requires more of the structure of the unit interval. Besides, when $\Pi(\phi) = 0$, $\pi(\omega \mid_m \phi) = \pi(\omega \mid \cdot, \phi) = 1, \forall \omega$, by convention.

Example 2 Let us revise the possibility distribution π given in Example 1 by the information that q is certainly false. If we use minimum-based conditioning we get:

ω	$\pi(\omega \mid_m \neg q)$
$\neg q r$	1
$\neg q \neg r$.5
qr	0
$q \neg r$	0

However, if we use the product-based conditioning, we get:

ω	$\pi(\omega \mid \cdot, \neg q)$
$\neg q r$	1
$\neg q \neg r$	5/7
qr	0
$q \neg r$	0

3.2 Revision with an uncertain input

We shall now consider the revision of π by some uncertain input information of the form (ϕ, a) into a new epistemic state denoted by $\pi' = \pi(\omega \mid (\phi, a))$. The input (ϕ, a) may be interpreted, and therefore treated, according to two slightly different views:

1. as a constraint which forces π' to satisfy

$$N'(\phi) = a, \quad (\text{i.e., } \Pi'(\phi) = 1 \text{ and } \Pi'(\neg\phi) = 1 - a); \quad (1)$$

this is the view taken in [2];

2. as information from a partially trusted source, which is taken into account only if it leads to a strengthening of the certainty; in other words, it forces π' to satisfy

$$N'(\phi) = \max\{N(\phi), a\}, \quad (2)$$

i.e., the previously held degree of belief in ϕ is not lowered just because a less trusted sources confirms it; this is the view taken in [6].

Both views have their intuitive justification in some context and there is no reason to privilege one or the other *a priori*.

Clearly, properties defined for revision are all suitable for revising with uncertain input, with the exception of \mathbf{A}_2 , which is no longer appropriate since $\Pi'(\neg\phi) \neq 0$ for $a < 1$. \mathbf{A}_2 is replaced by the following two axioms:

$$\mathbf{A}'_2: \Pi'(\phi) = 1, \Pi'(\neg\phi) \leq 1 - a;$$

$$\mathbf{A}''_2: \forall \omega, \omega' \notin [\phi], \text{ if } \pi(\omega) \geq \pi(\omega') \text{ then } \pi'(\omega) \geq \pi'(\omega').$$

\mathbf{A}''_2 preserves the relative order between countermodels of ϕ , but in a weaker sense than in axiom \mathbf{A}_3 for the models of ϕ . Note that there is no further constraint which relates models of ϕ and countermodels of ϕ in the new epistemic state.

\mathbf{A}'_2 is general, in the sense that it covers both views of the uncertain input; however, for View 1, it might be replaced by a stricter version

$$\mathbf{A}'_{2=}: \Pi'(\phi) = 1, \Pi'(\neg\phi) = 1 - a.$$

\mathbf{A}_3 and \mathbf{A}''_2 suggest that revising with uncertain input can be achieved using two parallel changes with a sure input: first, a conditioning on ϕ and one on $\neg\phi$. Then, in order to satisfy \mathbf{A}'_2 , the distribution $\pi(\cdot \mid \neg\phi)$ is “denormalized”, so as to satisfy $\Pi'(\neg\phi) = 1 - a$. Therefore, revising with uncertain information can be achieved using the following definition:

$$\pi(\omega \mid (\phi, a)) = \begin{cases} \pi(\omega \mid \phi), & \text{if } \omega \models \phi; \\ (1 - a) * \pi(\omega \mid \neg\phi), & \text{otherwise,} \end{cases} \quad (3)$$

where $*$ is either min or the product, \cdot , depending on whether conditioning is based on the product or the minimum operator.

When $*$ = product (resp. min) the possibilistic revision is called *product-based* (resp. *minimum-based*) conditioning with an uncertain input, denoted $\pi(\omega \mid \cdot, (\phi, a))$, (resp. $\pi(\omega \mid_m (\phi, a))$).

One important thing to remark is that conditioning is performed on models ($\omega : \omega \models \phi$) and countermodels ($\omega : \omega \not\models \phi$) independently. Therefore, nothing forbids, in principle, applying one $*$ operator to models and another $*$ operator on countermodels. Indeed, in some applications, it may happen that one $*$ operator provides the desired behavior on models, but a counterintuitive one on countermodels, or viceversa. For example, minimum-based conditioning lowers the possibility of all countermodels greater than $1 - a$, while leaving untouched the others, which might be regarded as intuitively correct when modeling belief revision in a cognitive agent; however, on the other hand, it only raises the possibility of the most possible models, whereas one might find that it would be more desirable that the possibility of *all* worlds compatible with incoming information should increase proportionally, which is the behavior provided by product-based conditioning. The independence of conditioning on models and countermodels allows us to try different combinations of operators to obtain exactly the desired behavior.

According to the two interpretations of uncertain inputs (namely, as a constraint or as partially trusted information), two families of possibilistic conditioning operators may be defined. For the sake of clarity, we will replace the generic conditioning symbol “ \mid ” by two distinct specific symbols, namely “ \downarrow ” for Family 1 and “ \uparrow ” for Family 2. Furthermore, we will distinguish minimum-based and product-based conditioning by the symbols m and \cdot , added to the conditioning symbol as superscripts, to indicate their use for models, or as subscripts, to indicate their use for countermodels.

Family 1: The *minimum-based conditioning* is defined,

- for $\omega \models \phi$ (models), as

$$\pi(\omega \downarrow^m (\phi, a)) = \begin{cases} 1, & \pi(\omega) = \Pi(\phi); \\ \pi(\omega), & \pi(\omega) < \Pi(\phi); \end{cases} \quad (4)$$

- for $\omega \not\models \phi$ (countermodels), as

$$\pi(\omega \downarrow_m (\phi, a)) = \begin{cases} 1 - a, & \text{if } \pi(\omega) = \Pi(\neg\phi) \\ & \text{or } \pi(\omega) > 1 - a; \\ \pi(\omega), & \text{otherwise.} \end{cases} \quad (5)$$

The *product-based conditioning* is defined,

- for $\omega \models \phi$ (models), as

$$\pi(\omega \downarrow^* (\phi, a)) = \begin{cases} \frac{\pi(\omega)}{\Pi(\phi)}, & \Pi(\phi) > 0; \\ 1, & \Pi(\phi) = 0; \end{cases} \quad (6)$$

- for $\omega \not\models \phi$ (countermodels), as

$$\pi(\omega \downarrow_*. (\phi, a)) = \begin{cases} (1 - a) \frac{\pi(\omega)}{\Pi(\neg\phi)}, & \Pi(\neg\phi) > 0; \\ 1 - a, & \Pi(\neg\phi) = 0. \end{cases} \quad (7)$$

Family 2: The *minimum-based conditioning* is defined,

- for $\omega \models \phi$ (models), as

$$\pi(\omega \uparrow^m (\phi, a)) = \begin{cases} 1, & \pi(\omega) = \Pi(\phi); \\ \pi(\omega), & \pi(\omega) < \Pi(\phi); \end{cases} \quad (8)$$

- for $\omega \not\models \phi$ (countermodels), as

$$\pi(\omega \uparrow_m (\phi, a)) = \min\{1 - a, \pi(\omega)\}. \quad (9)$$

The *product-based conditioning* is defined,

- for $\omega \models \phi$ (models), as

$$\pi(\omega \uparrow^* (\phi, a)) = \begin{cases} \frac{\pi(\omega)}{\Pi(\phi)}, & \Pi(\phi) > 0; \\ 1, & \Pi(\phi) = 0; \end{cases} \quad (10)$$

- for $\omega \not\models \phi$ (countermodels), as

$$\pi(\omega \uparrow_*. (\phi, a)) = (1 - a)\pi(\omega). \quad (11)$$

From the above definitions, it is clear that the new ranking on models of ϕ is simply obtained using conditioning with a sure input.

For Family 1 conditioning operators, the new ranking of countermodels of ϕ depends on the relative position of the *a priori* certainty of ϕ , and the prescribed posterior certainty of ϕ :

- If $N(\phi) \leq a$ and when $*$ = min, all countermodels that were originally more plausible than $1 - a$, are forced to level $1 - a$, which means that some strict ordering between countermodels of ϕ may be lost. When $*$ = product, all plausibility levels are proportionally shifted down (to the level $1 - a$).
- If $N(\phi) > a$ the best countermodels of ϕ are raised to level $1 - a$. Moreover, when $*$ = product, the plausibility levels of other countermodels are proportionally shifted up (to level $1 - a$).

For Family 2 conditioning operators, the new ranking of countermodels of ϕ depends on the relative position of the *a priori* certainty of ϕ , and the degree a to which the source of input ϕ is to be trusted:

- As it is the case with Family 1, if $N(\phi) \leq a$ and when $*$ = min, all countermodels that were originally more plausible than $1 - a$, are forced to level $1 - a$. When $*$ = product, all plausibility levels are proportionally shifted down (to the level $1 - a$).

- However, if $N(\phi) > a$ the best countermodels of ϕ are left untouched or even, when $*$ = product, proportionally shifted down by a factor $1 - a$.

Note that, in both families, when $a = 1$, we recover conditioning by a totally reliable input.

When $*$ = product, a stronger version of \mathbf{A}_2'' holds whereby the order of countermodels of ϕ is fully preserved, hence it satisfies:

A₆: $\forall \omega_1, \omega_2 \notin [\phi], \pi(\omega_1) \leq \pi(\omega_2)$ iff $\pi'(\omega_1) \leq \pi'(\omega_2)$.

Moreover, if $N(\phi) \leq a$, we can check that the following two postulates are also satisfied:

A₇: If $\omega_1 \models \phi$ and $\omega_2 \models \neg\phi$, then $\pi(\omega_1) < \pi(\omega_2)$ implies $\pi'(\omega_1) < \pi'(\omega_2)$.

A₈: If $\omega_1 \models \phi$ and $\omega_2 \models \neg\phi$, then $\pi(\omega_1) \leq \pi(\omega_2)$ implies $\pi'(\omega_1) \leq \pi'(\omega_2)$.

Example 3 Let us again consider the possibility distribution π of Example 1. Let $(q \vee r, .2)$ be the uncertain input. Note that $N_\pi(q \vee r) = .5$, and hence taking into account the input should decrease our belief in the information $q \vee r$. Using minimum-based conditioning, we get:

ω	$\pi(\omega \mid_m (q \vee r, .2))$
$q r$	1
$q \neg r$	1
$\neg q r$	$.7$
$\neg q \neg r$	$.8$

In this example, the product-based conditioning leads to the same result. Note that the main difference with conditioning with sure input is that countermodels of ϕ are no longer impossible.

In Example 3, the uncertain input is treated as a constraint which is enforced; therefore, a Family 1 operator is used. If the input had been treated as information from a partially trusted source, no revision would have taken place.

3.3 Generalized Conditioning Operators

We now introduce the family of generalized conditioning operators, arising from all possible combinations of

- the view according to which the input is treated;
- the conditioning operator to be applied to models;
- the conditioning operator to be applied to countermodels.

Therefore, we have eight operator definitions, as summarized in Table 1, for all ω , ϕ and a .

Proposition All the operators defined in Table 1 satisfy axioms $\mathbf{A}_1, \mathbf{A}_2', \mathbf{A}_2'', \mathbf{A}_3, \mathbf{A}_4, \mathbf{A}_5, \mathbf{A}_7$, and \mathbf{A}_8 .

Additionally, all four \downarrow operators satisfy also axiom $\mathbf{A}_{2=}$, and operators $\downarrow^m, \downarrow_*, \uparrow_m^m, \uparrow^m, \uparrow_m^*$, and \uparrow_* satisfy axiom \mathbf{A}_6 .

Proof: Omitted due to lack of space.

For the reader's convenience, Table 2 provides a summary of the axioms satisfied by each operator.

Based on the axioms satisfied by each operator, we may notice that Family 2 constitutes a homogeneous cluster of conditioning operators, whereas Table 2 suggests that Family 1 should be divided

Table 1. Definitions of the eight generalized conditioning operators.

	$\omega \models \phi$	$\omega \not\models \phi$
$\pi(\omega \downarrow_m^m(\phi, a))$	1, if $\pi(\omega) = \Pi(\phi)$ $\pi(\omega)$, if $\pi(\omega) < \Pi(\phi)$	$1 - a$, if $\pi(\omega) = \Pi(\neg\phi)$ or $\pi(\omega) > 1 - a$ $\pi(\omega)$, otherwise
$\pi(\omega \downarrow_*^m(\phi, a))$	1, if $\pi(\omega) = \Pi(\phi)$ $\pi(\omega)$, if $\pi(\omega) < \Pi(\phi)$	$(1 - a) \frac{\pi(\omega)}{\Pi(\neg\phi)}$, if $\Pi(\neg\phi) > 0$ $1 - a$, if $\Pi(\neg\phi) = 0$
$\pi(\omega \downarrow_m^*(\phi, a))$	$\frac{\pi(\omega)}{\Pi(\phi)}$, if $\Pi(\phi) > 0$ 1, if $\Pi(\phi) = 0$	$1 - a$, if $\pi(\omega) = \Pi(\neg\phi)$ or $\pi(\omega) > 1 - a$ $\pi(\omega)$, otherwise
$\pi(\omega \downarrow_*^*(\phi, a))$	$\frac{\pi(\omega)}{\Pi(\phi)}$, if $\Pi(\phi) > 0$ 1, if $\Pi(\phi) = 0$	$(1 - a) \frac{\pi(\omega)}{\Pi(\neg\phi)}$, if $\Pi(\neg\phi) > 0$ $1 - a$, if $\Pi(\neg\phi) = 0$
$\pi(\omega \uparrow_m^m(\phi, a))$	1, if $\pi(\omega) = \Pi(\phi)$ $\pi(\omega)$, if $\pi(\omega) < \Pi(\phi)$	$\min\{1 - a, \pi(\omega)\}$
$\pi(\omega \uparrow_*^m(\phi, a))$	1, if $\pi(\omega) = \Pi(\phi)$ $\pi(\omega)$, if $\pi(\omega) < \Pi(\phi)$	$(1 - a)\pi(\omega)$
$\pi(\omega \uparrow_m^*(\phi, a))$	$\frac{\pi(\omega)}{\Pi(\phi)}$, if $\Pi(\phi) > 0$ 1, if $\Pi(\phi) = 0$	$\min\{1 - a, \pi(\omega)\}$
$\pi(\omega \uparrow_*^*(\phi, a))$	$\frac{\pi(\omega)}{\Pi(\phi)}$, if $\Pi(\phi) > 0$ 1, if $\Pi(\phi) = 0$	$(1 - a)\pi(\omega)$

Table 2. Summary of the axioms satisfied by each operator.

	\downarrow_m^m	\downarrow_*^m	\downarrow_m^*	\downarrow_*^*	\uparrow_m^m	\uparrow_*^m	\uparrow_m^*	\uparrow_*^*
A₁	•	•	•	•	•	•	•	•
A₂	•	•	•	•	•	•	•	•
A₂^h	•	•	•	•				
A₂ⁿ	•	•	•	•	•	•	•	•
A₃	•	•	•	•	•	•	•	•
A₄	•	•	•	•	•	•	•	•
A₅	•	•	•	•	•	•	•	•
A₆		•		•	•	•	•	•
A₇	•	•	•	•	•	•	•	•
A₈	•	•	•	•	•	•	•	•

into two homogeneous sub-families, which we might call Family 1.1, comprising \downarrow_m^m and \downarrow_*^m , and Family 1.2, comprising \downarrow_m^* and \downarrow_*^* . Family 1.1's specificity is that it does not always fully preserve the order of countermodels of ϕ , thus causing some information loss. Family 1.2, on the other hand, fully preserves the order of countermodels of ϕ , like Family 2, while forcing the necessity of ϕ to a .

Example 4 Let us again consider the possibility distribution π of Example 1 and let us revise it for inputs $(\neg q, .2)$, $(q \vee r, .4)$ and $(r, .6)$ using every single operator defined in Table 1. Notice that the former input requires a revision, whereas the latter two bring about an expansion of the belief base. We get the results shown in Table 3.

While essentially obeying the same postulates, with a few slight differences, the eight operators show an interesting variety of behaviors, which make them suitable to a variety of contexts. Given an application, it is highly likely that one can find among them the one that fits the requirements of a belief revision operator in that context. The following are some guidelines to help the reader to single out the particular operator that suits her needs:

- The four \downarrow operators treat incoming information (ϕ, a) as a fully reliable indication of the degree of necessity of ϕ ; in other words, they act as if a completely trusted source told the agent that

Table 3. Results of revising possibility distribution π of Example 1 for the inputs of Example 4 with every operator of Table 1.

$\omega \longrightarrow$	qr	$q\neg r$	$\neg qr$	$\neg q\neg r$
$\pi(\omega \downarrow_m^m(\neg q, .2))$	0.8	0.8	1	0.5
$\pi(\omega \downarrow_m^m(q \vee r, .4))$	1	1	0.7	0.6
$\pi(\omega \downarrow_m^m(r, .6))$	1	0.4	0.7	0.4
$\pi(\omega \downarrow_*^m(\neg q, .2))$	0.8	0.8	1	0.5
$\pi(\omega \downarrow_*^m(q \vee r, .4))$	1	1	0.7	0.6
$\pi(\omega \downarrow_*^m(r, .6))$	1	0.4	0.7	0.2
$\pi(\omega \downarrow_m^*(\neg q, .2))$	0.8	0.8	1	0.7143
$\pi(\omega \downarrow_m^*(q \vee r, .4))$	1	1	0.7	0.6
$\pi(\omega \downarrow_m^*(r, .6))$	1	0.4	0.7	0.4
$\pi(\omega \downarrow_*^*(\neg q, .2))$	0.8	0.8	1	0.7143
$\pi(\omega \downarrow_*^*(q \vee r, .4))$	1	1	0.7	0.6
$\pi(\omega \downarrow_*^*(r, .6))$	1	0.4	0.7	0.2
$\pi(\omega \uparrow_m^m(\neg q, .2))$	0.8	0.8	1	0.5
$\pi(\omega \uparrow_m^m(q \vee r, .4))$	1	1	0.7	0.5
$\pi(\omega \uparrow_m^m(r, .6))$	1	0.4	0.7	0.4
$\pi(\omega \uparrow_*^m(\neg q, .2))$	0.8	0.8	1	0.5
$\pi(\omega \uparrow_*^m(q \vee r, .4))$	1	1	0.7	0.3
$\pi(\omega \uparrow_*^m(r, .6))$	1	0.4	0.7	0.2
$\pi(\omega \uparrow_m^*(\neg q, .2))$	0.8	0.8	1	0.7143
$\pi(\omega \uparrow_m^*(q \vee r, .4))$	1	1	0.7	0.5
$\pi(\omega \uparrow_m^*(r, .6))$	1	0.4	0.7	0.4
$\pi(\omega \uparrow_*^*(\neg q, .2))$	0.8	0.8	1	0.7143
$\pi(\omega \uparrow_*^*(q \vee r, .4))$	1	1	0.7	0.3
$\pi(\omega \uparrow_*^*(r, .6))$	1	0.4	0.7	0.2

$N(\phi) = a$; therefore, even if ϕ was believed to a higher degree, its degree of belief after revision will be exactly a . In iterated revision, it is always the last input about ϕ that determines $N(\phi)$.

- The four \uparrow operators treat incoming information (ϕ, a) as an indication that ϕ is true provided by a source partially trusted to degree a ; therefore, the agent's degree of belief in ϕ will never decrease. By the way, this is the intuitive motivation for using symbol \uparrow for this family of operators.
- The two \uparrow_m operators react to repeated confirmations of the same input ϕ by different sources by adopting belief ϕ with the same degree as the degree of trust (or reliability) of the most reliable among those sources. We might say these operators are rather weary or conservative.
- The two \uparrow operators, instead, treat any confirmation of ϕ , even by little reliable sources, as additional evidence in favor of ϕ , thus increasing the degree of belief of ϕ . These two operators will be appropriate for modeling the behavior of a credulous agent.
- The four \downarrow operators re-normalize the possibility distribution π' proportionally, i.e., they raise the possibility of all worlds that are compatible with incoming information; in a sense, they model the behaviour of an open-minded agent who, upon being convinced of ϕ more than it was before, concedes that all worlds in which ϕ holds are now less unlikely.
- In contrast, the four \downarrow^m operators re-normalize π' in the most conservative way, by raising to 1 just the models of ϕ that were deemed most possible before the arrival of the input; they model the behavior of a more opinionated agent, who is not willing to give up any of its beliefs unless absolutely forced to do so by new evidence and, even then, only by the smallest amount possible.

It is hard to give a general recipe suggesting the proper operator to use in each situation, because some of the differences among them are very subtle. A reasonable suggestion would be to determine experimentally which of the eight operators (or of a subset thereof, determined *a priori* based on some desired properties) is most suitable to a given application scenario. This is, after all, the usual way to proceed when it comes to choosing from a parametric family of operators, e.g., logical connectives or defuzzification operators in fuzzy logic.

4 Conclusion, Related, and Future Work

We have defined on the semantic level eight belief revision operators based on possibilistic conditioning showing an interesting variety of behaviors while obeying the basic postulates of belief revision.

The possibilistic conditioning operators we denote by \uparrow_m^m and \uparrow ; were proposed and characterized by Dubois and Prade [6]; the ones we denote by \downarrow_m^m and \downarrow ; were studied by Benferhat, Dubois, Prade, and Williams [2]. Finally, one of the hybrid operators, namely the one we denote by \uparrow_m , was proposed and characterized by da Costa Pereira and Tettamanzi [3] to model belief revision in BDI agents with partially trusted sources.

The next obvious step will be to work out the syntactic implementation of the operators studied, in agreement with the semantics, so that the revision of a belief base can be efficiently computed.

Several authors have proposed postulates for iterated belief revision which are added to the AGM postulates. Benferhat and colleagues describe in [2] the intuition behind the Darwiche and Pearl (DP) postulates [4] and conclude that possibilistic revision with uncertain input is more in the spirit of the DP postulates, except that in possibilistic revision there is no limitation on the input (ϕ, a) leading to a revision.

The proposed operators could be improved in order to deal with some “weaknesses” typical of the operators obeying the AGM and the DP postulates.

In [10], the authors propose an approach to deal with the problem of “drowning effect”. This problem is raised by the fact that after revising a belief base Σ with a totally reliable formula ϕ , the result of revision does not include the formulas whose weights are lower than the inconsistency degree of the new base. The problem consists then in the possible loss of too much information if the inconsistency degree is high. Like in [2], the operators proposed here do not deal with this problem. This is kept for future work.

The two families of conditioning operators studied in this paper consider the input (ϕ, a) as the constraint $N'(\phi) = a$ and $N'(\phi) \geq a$, respectively. In the same vein, one could also consider revising with (negative) uncertain inputs giving the constraint $N'(\phi) \leq a$, with $a < 1$. This is an interesting issue that has not been considered before, and we leave it for consideration in future work. To be sure, to address this issue one would need to modify Equation 2.

Another proposal that could be taken into account for future work is Jin and Thielscher's [9] Independence postulate for iterated belief revision, which aims at overcoming the “weakness” of the AGM and DP postulates which force an agent to delete everything it has previously learned upon reception of an input which contradicts its currently held beliefs.

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] 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).
- [3] C. da Costa Pereira and A. Tettamanzi, ‘An integrated possibilistic framework for goal generation in cognitive agents’, in *Proceedings of AAMAS'10*, pp. 1239–1246. IFAAMAS, (2010).
- [4] D. Darwiche and J. Pearl, ‘On the logic of iterated belief revision’, *Artificial Intelligence*, **89**, 1–29, (1997).
- [5] B. De Baets, E. Tsiporkova, and R. Mesiar, ‘Conditioning in possibility theory with strict order norms’, *Fuzzy Sets Syst.*, **106**(2), 221–229, (1999).
- [6] 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).
- [7] D. Dubois and H. Prade, ‘Possibility theory: qualitative and quantitative aspects’, in *Handbook of Defeasible Reasoning and Uncertainty Management Systems, Vol. 1: Quantified Representation of Uncertainty and Imprecision*, eds., D. Gabbay and P. Smets, 169–226, Kluwer Academic Press, (1998).
- [8] P. Gärdenfors, *Knowledge in Flux: Modeling the Dynamics of Epistemic States*, MIT Press, 1988.
- [9] Y. Jin and M. Thielscher, ‘Iterated belief revision, revised’, *Artif. Intell.*, **171**(1), 1–18, (2007).
- [10] G. Qi, ‘A semantic approach for iterated revision in possibilistic logic’, in *AAAI*, pp. 523–528, (2008).
- [11] G. Shafer, ‘Jeffrey's rule of conditioning’, *Philosophy of Science*, **48**, 337–362, (1981).
- [12] W. Spohn, ‘Ordinal conditional functions: A dynamic theory of epistemic states’, in *Causation in Decision, Belief Change, and Statistics*, eds., W. L. Harper and B. Skyrms, volume 2, 105–134, D. Reidel, Dordrecht, (1988).