A Synthetic View of Belief Revision with Uncertain Inputs in the Framework of Possibility Theory*

Didier DUBOIS and Henri PRADE

Institut de Recherche en Informatique de Toulouse (IRIT)
Université Paul Sabatier, 118 route de Narbonne
31062 Toulouse Cedex 4, France

Abstract

This paper discusses belief revision under uncertain inputs in the framework of possibility theory. This framework is flexible enough to account for numerical and ordinal revision procedures. It is emphasized that revision under uncertain inputs can be understood in two different ways depending on whether the input is viewed as a constraint to be enforced, or as an unreliable piece of information. Two revision rules are proposed to implement these forms of revision. It is shown that M.A. Williams' transmutations, originally defined in the setting of Spohn's functions, can be captured in possibility theory, as well as Boutilier's natural revision. The use of conditioning greatly simplifies the description of these belief change operations. Lastly, preliminary results on implementing revision rules at the syntactic level are given.

1 - Introduction

Belief revision in the sense of Alchourrón, Gärdenfors and Makinson (1985) (also called AGM revision) presupposes the existence of a so-called epistemic entrenchment relation which rank-orders the formulas in a belief base to be revised (Gärdenfors, 1988). It has been pointed out that this epistemic entrenchment is nothing but a qualitative necessity measure and that faithful extensions of Gärdenfors' notions of expansion, contraction and revision can be defined in the framework of possibility theory and possibilistic logic; see (Dubois and Prade, 1991a, 1992). In this setting, belief change can be either discussed in terms of a distribution (here a possibility distribution) defined on the possible worlds (as in probability theory and in the work of Grove (1988)), or in terms of a knowledge base in possibilistic logic (made of

* This paper is a revised and expanded version of a paper entitled "Belief revision with uncertain inputs in the possibilistic setting", published in the Proceedings of the 12th Conference on Uncertainty in Artificial Intelligence, Portland, Oregon, Aug. 1-4, 1996 (E. Horvitz, F. Jensen, eds.), Morgan & Kaufmann, San Francisco, CA, pp. 236-243.

classical propositional formulas weighted by lower bounds of necessity measures) from which a possibility distribution can be defined.

With this view in mind, revising a layered knowledge base by the introduction of a new piece of information turns out to be representable by an appropriate conditioning of a possibility distribution. This conditioning operation remains in agreement with AGM postulates (after Alchourrón, Gärdenfors and Makinson) for revision (Gärdenfors, 1988). However in the spirit of conditioning the input information is taken as a sure fact (with maximal level of certainty), while in the spirit of belief revision theory, the input should only belong to the revised belief set, possibly at the lower level. The latter is called natural revision by Boutilier (1993). Conditions under which possibility theory captures natural revision are given.

In this paper the extension of possibilistic conditioning to uncertain inputs is extensively studied, along the lines previously investigated in Dubois and Prade (1991b, 1993). It is pointed out that belief revision with uncertain inputs can be expressed in terms of mixtures: convex sum/product mixtures in the probabilistic setting, and max-min or max-product mixtures in the possibilistic setting (Dubois et al., 1996). There are two views of what an uncertain input may mean. The first view, already implemented in the probabilistic rule due to Jeffrey, is to consider that an uncertain input is a constraint that the final belief state must satisfy. It acts as a correction of the prior belief state. The other view is that of an unreliable input that is liable of being rejected. Two possibilistic revision rules are described that implement this distinction.

It is then shown in the paper that natural revision is in some cases equivalent to several Jeffrey-like revisions under uncertain inputs in the possibilistic setting. Moreover, possibilistic revision under uncertain inputs is also shown to exactly coincide with M.A. Williams' notion of adjustment. More generally, a unified view is provided, of various approaches developed by Boutilier (1993), Boutilier and Goldszmidt (1993), Williams (1994a, b; 1995), Goldszmidt and Pearl (1992), Darwiche and Pearl (1994).

These works lay bare two representations of conditioning, that can be expressed in possibility theory, one (based on minimum operation) which is purely ordinal and is at work in Williams' adjustments for instance, and another (based on product, thus requiring a richer scale) alike Dempster's rule of conditioning, which is equivalent to Spohn (1988) conditioning through a rescaling and is used by Pearl and his coauthors. Lastly, it is explained how to revise a possibilistic knowledge base by means of an uncertain input without making the underlying possibility distribution explicit. We start by briefly recalling how belief change under uncertain inputs is dealt with in the probabilistic setting.

2 - The Probabilistic Setting for Belief Revision

Bayes rule of conditioning by a sure event, has been understood by some researchers as a belief revision operation that, given a prior probability P and a sure input A, yields a posterior probability P' = P(.|A) (see e.g., Gärdenfors, 1988). In this view, the input A is viewed as a

constraint P'(A) = 1. It has been extended by Jeffrey (1965) to the case of uncertain inputs under the form of a partition $\{A_1, A_2, \ldots, A_n\}$ of Ω, where each event A_i is attached probability α_i (with $\sum\alpha_i = 1$ and $P'(A_i) > 0$). The result of the revision is then

$$P'(B) = P(B \mid \{(A_i,\alpha_i)\}_{i=1,n}) = \sum_{i=1,n} \alpha_i P(B \mid A_i). \quad (1)$$

In the case of a binary partition, i.e., an uncertain input A with probability α, it gives:

$$P(B \mid (A,\alpha)) = \alpha P(B \mid A) + (1 - \alpha)P(B \mid \bar{A}). \quad (2)$$

This revision rule (1)-(2) has been justified by P.M. Williams (1980) on the basis of minimizing the informational distance I(P,P') under the constraints $P'(A_i) = \alpha_i$ for i = 1,n. (1) can also be justified at the formal level by the fact that the only way of combining the conditional probabilities $P(B \mid A_i)$ in an eventwise manner (i.e., using the same combination law for all events B) so as to recover a probability measure is to use a linear weighted combination such as (1) (Lehrer and Wagner, 1981). Note that the uncertain input is really viewed as a constraint that forces the probability measure to bear certain values on a partition of Ω. The input assigns new probability values to prescribed events forming a partition of Ω. Jeffrey's rule ensures that probabilities do not change in relative value for situations within each partition element A_i. In the above approach, the coefficient α_i is interpreted as the sure claim that the probability of input A_i is α_i; and leads to a correction of the prior probability. Especially it is a genuine revision process since the a priori probability and the input are at the same level, for instance they are both generic knowledge, or both uncertain pieces of evidence.

However, consider the case where the probability P refers to a population and the input $\{(A_i,\alpha_i)\}_{i=1,n}$ means that an event has occurred that informs about the current situation, and α_i is the probability that this particular event is A_i. This situation is one of unreliable input. The input is really a piece of uncertain evidence and (1) then just computes the expected value of the conditional probability $P(B \mid A)$ where A is a random event whose realizations belong to the partition $\{A_1, \ldots, A_n\}$. This random event A is then a genuine unreliable observation, α_i being the probability that A_i is *the* true input (and not the true probability of A_i to be enforced as in the previous paragraph). So, $\alpha_i = P''(A = A_i)$. In other words, the coefficients represent a *complete* probability assignment P'' on the family of subsets 2^Ω, and $\alpha_i = P''(\{A_i\})$, while in Jeffrey's rule, $\{(A_i,\alpha_i)\}_{i=1,n}$ is an *incompletely described* probability measure P' on Ω and $\alpha_i = P'(A_i)$. When $\alpha_i = 1$, the conditional probability $P(B \mid A_i)$ represents the result of focusing P on the reference class A_i, and not the revision of the generic knowledge encoded by P. Focusing on a reference class for the purpose of deriving a plausible conclusion pertaining on a case at hand is more in the spirit of many probabilists such as De Finetti, than the spirit of probability kinematics after Jeffrey. It may look strange that the two views, corresponding to a revision and an uncertain focusing respectively, coincide in their implementation. However the requirement that $\{A_1, \ldots, A_n\}$ forms a partition does not look compulsory when the input is viewed as an uncertain (random) observation rather than a constraint on probability values (since it must just be a probability assignment on 2^Ω). The partition is compulsory only when the uncertain input

is a constraint in order to ensure that the result verifies the condition $P(A_i \mid \{(A_i,\alpha_i), i = 1,n\}) = \alpha_i$ (since $\sum_{i=1,n} \alpha_i = 1$). The unreliable input (A,α) interpreted as "event A has probably been observed (but maybe nothing has been observed)", corresponds to the probability assignment $\{(A,\alpha), (\Omega, 1 - \alpha)\}$ and not $\{(A,\alpha), (\bar{A}, 1 - \alpha)\}$ as for Jeffrey's rule. Indeed $P(\{A\}) = \alpha$ does not imply that $P(\{\bar{A}\}) = 1 - \alpha$.

3 - The Possibilistic Setting for Belief Revision

The distinction between unreliable inputs and constraints on the uncertainty pervading the result of a belief change operation exists regardless of the considered representation framework for an agent's belief. The reason why it is useful to cast it in the setting of possibility theory is that the latter enables belief revision operations to be expressed that are exactly AGM revisions, but the terminology of possibility theory is very similar to the one of probability theory. This hybrid nature of possibility theory makes it possible to point out the role of conditioning as a particular revision tool, and to envisage the question of uncertain inputs in the AGM setting.

3.1 - Basic Notions

It makes sense to view the possibilistic approach as a slight improvement of the pure logical approach to knowledge representation from the point of view of expressiveness. Instead of viewing a belief state as a flat set of mutually exclusive situations, one adds a complete partial ordering on top, according to which some situations are considered as more plausible than others. A cognitive state can then be modelled by a possibility distribution π, that is, a mapping from Ω to a totally ordered set V containing a greatest element (denoted 1) and a least element (denoted 0), typically the unit interval V = [0,1]. However any finite, or infinite and bounded, chain will do as well. The advantage of using the plausibility scale V is that it makes it easier to compare cognitive states. The idea of representing a cognitive state via a plausibility ordering on a set of situations ("possible worlds") is also developed by Grove (1988) and systematically used by Boutilier in his works (e.g., Boutilier, 1993), although he uses the language of modal logic.

Similarly to the probabilistic case, a possibility distribution generates a set function Π called a possibility measure (Zadeh, 1978) defined by (for simplicity V = [0,1])

$$\Pi(A) = \max_{\omega \in A} \pi(\omega) \tag{3}$$

and satisfying $\Pi(A \cup B) = \max(\Pi(A),\Pi(B))$ as a basic axiom. Except if the contrary is explicitly stated, π will be supposed to be normalized, i.e., $\exists \omega_0$, $\pi(\omega_0) = 1$ (ω_0 is not necessarily unique). The degree of certainty of A is measured by means of the dual necessity function $N(A) = 1 - \Pi(\bar{A})$. Moreover, if Ω is the set of interpretations of a propositional

language $\mathcal{L}$, let us denote by [p] the set of models of a proposition $p \in \mathcal{L}$. Then, any possibility distribution π leads to a belief set K = {p, such that N([p]) > 0} i.e., a deductively closed sets of logical formulae, where $p \in K$ means that p is accepted by the agent holding a cognitive state described by π. The models of K form the core of π, $C(\pi) = \{\omega, \pi(\omega) = 1\}$ since $p \in K$ iff N([p]) > 0 iff [p] contains $C(\pi)$. However K can also be viewed as a completely preordered belief set according to the levels N([p]). If N([p]) = 1, it means that p will never be questioned. In Gärdenfors theory only tautologies are of that kind. In possibility theory, other formulas can be protected in this way. Apart from this point of divergence, the ordering of K so-obtained is an epistemic entrenchment ordering in the sense of Gärdenfors (Dubois and Prade, 1991a).

3.2 - Revision and Conditioning

Revision with a sure input in possibility theory is performed by means of a conditioning device similar to the probabilistic one. A conditional possibility is a distinguished solution to an equation of the form (Cox, 1946):

$$\forall B, \prod(A \cap B) = \prod(B \mid A) * \prod(A). \tag{4}$$

that is similar to Bayesian conditioning, together with $N(B \mid A) = 1 - \prod(\bar{B} \mid A)$. Possible choices for $*$ are min and the product (the latter makes sense only in the numerical settings). For $*$ = min this equation may have more than one solution $\prod(B \mid A)$. The least specific solution to (4) is chosen, (i.e., the solution with the greatest possibility degrees in agreement with the constraint (4)). This choice follows a minimal commitment principle specifying that every situation remains possible unless explicitly ruled out. The possibility distribution underlying the conditional possibility measure $\prod(\cdot \mid A)$ is defined by (Dubois and Prade, 1986)

$$\begin{aligned} \pi(\omega \mid A) &= 1 \text{ if } \pi(\omega) = \prod(A) > 0, \omega \in A, \\ &= 0 \text{ if } \omega \notin A, \\ &= \pi(\omega) \text{ if } \pi(\omega) < \prod(A), \omega \in A. \end{aligned} \tag{5}$$

Then $\prod(B \mid A) = \prod(A \cap B)$ if $\prod(A \cap B) < \prod(A)$ and $\prod(B \mid A) = 1$ if $\prod(A \cap B) = \prod(A) > 0$. If $\prod(A) = 0$, then $\pi(\cdot \mid A) = 0$ that is, a fully inconsistent possibility distribution results. In this approach, conditioning on an impossible event results in destroying π. This is a matter of convention : Alternatively we might define $\pi(\cdot \mid A) = 1$, if $\prod(A) = 0$, that is, the non-informed possibility distribution is obtained. Note that the existence of situations ω such that $\pi(\omega) = \prod(A)$ is no longer guaranteed in the infinite case, and is generally added as an extra condition to the possibility distribution; it is referred to in the literature in terms of *well-ranked* orderings, and leads to adding specific postulates for revision (Williams, 1994a). Moreover, considering the belief set K induced by π, and A = [q] for some input formula q, it is not difficult to check that the revised belief set K^*_q induced by $\pi(\omega \mid A)$ is K^*_q = {p, such that N([p] | A) > 0}. The

set of models of K^*_q is $\{\omega, \pi(\omega) = \prod(A)\}$, which proves that possibilistic conditioning is fully coherent with an AGM revision when $\prod(A) > 0$. Indeed K^*_q can also be defined as $\{p, N(\neg q \vee p) > N(\neg q)\}$. Note that if $\prod(A) = 1$, then $\pi(\cdot \mid A) = \min(\pi, \mu_A)$ where μ_A is the characteristic function of A. It corresponds to the idea of an expansion, for then K^*_q is the set of consequences of $K \cup \{q\}$ (Dubois and Prade, 1992).

The above discussion of conditional possibility, using $* = \min$, makes sense in a purely qualitative setting. In a quantitative setting, $*$ = product may sound more reasonable and the expression corresponding to (4) is

$$\forall B, \prod(B \mid A) = \frac{\prod(A \cap B)}{\prod(A)} \tag{6}$$

provided that $\prod(A) \neq 0$. This is formally Dempster rule of conditioning, specialized to possibility measures, i.e., consonant plausibility measures in the sense of Shafer. The corresponding revised possibility distribution is

$$\pi(\omega \mid A) = \frac{\pi(\omega)}{\prod(A)}, \forall\ \omega \in A; \pi(\omega \mid A) = 0 \text{ otherwise.} \tag{7}$$

Note that for both forms of conditioning, $N(A) = 1 \Rightarrow \pi(\cdot \mid A) = \pi$ (no revision takes place when the input information is already known with certainty).

The idea of AGM style revision is to reach a consistent cognitive state even if the input information contradicts the a priori cognitive state described by a belief set. Possibilistic conditioning verifies the following counterparts of the AGM axioms of belief revision (Dubois and Prade, 1992), letting π^*_A denote a possibility distribution obtained by revising π with input A according to the AGM postulates :

- $(\prod^*_1)$ for any subset $A \subseteq \Omega$, π^*_A represents a cognitive state (stability)
- $(\prod^*_2)$ $N^*_A(A) = 1$ (priority to the new information)
- $(\prod^*_3)$ $\pi^*_A \geq \min(\pi, \mu_A)$ (revising does not give more specific results than expanding)
- $(\prod^*_4)$ if $\prod(A) = 1$ then $\pi^*_A \leq \min(\pi, \mu_A)$ (if A is not rejected by π, revision reduces to expansion)
- $(\prod^*_5)$ $\pi^*_A = 0$ if and only if $\prod(A) = 0$
- $(\prod^*_6)$ $A = B \Rightarrow \pi^*_A = \pi^*_B$
- $(\prod^*_7)$ $\pi^*_{A \cap B} \geq \min(\pi^*_A, \mu_B)$
- $(\prod^*_8)$ if $\prod(B \mid A) = 1$ then $\pi^*_{A \cap B} \leq \min(\pi^*_A, \mu_B)$

It makes sense to consider possibilistic conditioning as an AGM-like revision and let $\pi^*_A = \pi(\cdot \mid A)$ with any the two definitions of conditioning. Note that if A and B are not disjoint, it can be verified that $\pi(\cdot\ (\cdot \mid A) \mid B) = \pi(\cdot \mid A \cap B)$ when iterating the revisions. However if $A \cap B = \emptyset$ then $\pi(\cdot\ (\mid A) \mid B)$ is the absurd possibility distribution.

Moreover the qualitative conditioning (5) embodies a principle of minimal change in the sense of the Hamming distance between π and π' defined by $H(\pi,\pi') = \sum_{\omega\in\Omega} |\pi(\omega) - \pi'(\omega)|$ (Ω finite). It makes sense for a finite totally ordered possibility scale V, just mapping the levels to integers. It has been proved that $H(\pi; \pi(\cdot \mid A))$ is minimal under the constraint $N(A \mid A) = 1$ as long as there is a single situation ω_A where $\pi(\omega_A) = \prod(A)$ (Dubois and Prade, 1992). When there is more than one situation ω where $\pi(\omega) = \prod(A)$, the principle of minimal change leads to as many revision functions by selecting one situation ω_A where $\pi(\omega_A) = \prod(A)$ and letting $\pi^*_A(\omega_A) = 1$ and $\pi^*_A(\omega_A) = \prod(A)$ for other most possible situations in A (this is possibilistic conditioning following Ramer (1989)). In that case $\pi(\cdot \mid A)$ is the disjunction of these minimal change revisions. A minimal change information-theoretic justification for the numerical, product-based conditioning rule could be envisaged in terms of preserving relative levels of plausibility.

3.3 - Boutilier's Natural Revision and Possibility Theory

However there is a difference between conditioning and revision in the AGM style, in the scope of revising the plausibility ordering that guides the revision. Let π be a possibility distribution on a set of interpretations and K the belief set generated by π in the corresponding propositional language. Since $p \in K$ corresponds to $N([p]) > 0$, the success postulate $q \in K^*_q$ should only require that $N^*_q([q]) > 0$, while $N(A \mid A) = 1$ holds for the conditioning, in accordance to conditional probability ($P(A \mid A) = 1$). So axiom ($\prod^*_2$) sounds like a very strong translation of the AGM success postulate. Following Boutilier (1993), we can consider that the success postulate is better expressed in terms of a revised necessity measure N^*_A by $N^*_A(A) > 0$; this is called natural revision by Boutilier. A natural revision by input A only comes down to assigning to the most plausible situations in A a degree of possibility higher than to other situations (this is enough to ensure that $N^*_A(A) > 0$) while retaining the same ordering of situations as before revision, including *for situations outside A*. It implies that, after revision, some situations where A is not true may remain more plausible than situations where A is true. This feature does not fit the idea of revision via conditioning whereby, in the revised state, situations where the input information is false are deemed impossible (as is the case with probabilistic revision). A natural revision obeys the principle of minimal change, since only the set of most plausible worlds in A are moved to become the overall most plausible states.

To describe this elementary change in possibilistic terms requires the use of non-normalized possibility distributions. Indeed if $\prod(\bar{A}) > \prod(A)$, then $\prod(\bar{A}) = 1$ holds if the associated possibility distribution π is normalized. Besides, it must be the case that $\prod^*_A(\bar{A}) = \prod(\bar{A})$ (since the levels of possibility of $\bar{A}$-worlds remains unchanged) while $\prod^*_A(A) > \prod^*_A(\bar{A})) = 1$. This is not allowed by the usual conventions of possibility theory. Describing this type of change in possibility theory becomes easy if the possibility scale is not upper-bounded. Otherwise, it would require a "denormalisation" of the possibility distribution π, by shifting all possibility levels down (while preserving the ordering of interpretations), followed

by a shift of all best A-worlds up to 1, so that the latter step be meaningful. One way of implementing this method on the unit interval, with π normalized, is as follows:

$$\pi^*_A(\omega) = 1 \text{ if } \omega \in \{\omega \in A, \pi(\omega) = \prod(A)\} \text{ and } C(\pi) \cap A = \emptyset$$
$$\pi^*_A(\omega) = \pi(\omega) = 1 \text{ if } C(\pi) \cap A$$
$$\pi^*_A(\omega) = \pi(\omega) \text{ if } \omega \notin C(\pi) \text{ and } \omega \notin \{\omega \in A, \pi(\omega) = \prod(A)\} \text{ if } C(\pi) \cap A = \emptyset$$
and $$\pi^*_A(\omega) = \lambda \text{ where } 1 - N(C(\pi)) < \lambda < 1 \text{ if } \omega \in \bar{A} \cap C(\pi).^1 \quad (8)$$

There are three situations:

i) A contains $C(\pi)$: then $N(A) > 0$, and A is believed a priori. Then $\bar{A} \cap C(\pi) = \emptyset$ and $\pi^*_A(\omega) = \pi(\omega)$. No change occurs. Natural revision does not affect the possibility distribution at all, while conditioning leads to a non-trivial expansion, generally (since $\pi(\cdot \mid A) \neq \pi$).

ii) $A \cap C(\pi) = \emptyset$. Then all the most plausible situations in A are moved up to 1 and all most plausible situations are moved down to plausibility λ.

iii) if $\bar{A} \cap C(\pi) \neq \emptyset$ and $A \cap C(\pi) \neq \emptyset$, then $N(A) = N(\bar{A}) = 0$, and the natural revision comes down to a non trivial expansion of the belief set K. It affects the possibility distribution π just by moving the $\bar{A}$-worlds in $C(\pi)$ to the next possibility level down, that is (8) becomes:

$$\pi^+_A(\omega) = \pi(\omega) \text{ if } \omega \notin C(\pi) \text{ or } \omega \in C(\pi) \cap A$$
and $$\pi^+_A(\omega) = \lambda \text{ where } 1 - N(C(\pi)) < \lambda < \prod(A) = 1 \text{ otherwise.} \quad (9)$$

The study of revision operators in the setting of possibility theory, whether conditioning or natural revisions, points out that while the AGM theory gives a recipe to belief sets, possibilistic revision operations also revise the epistemic ordering of situations. Clearly iterated revision then becomes possible, as also pointed out by Williams (1994b). A noticeable property of natural revisions is that $(\pi^*_A)^*_B$ remains non trivial even if $A \cap B = \emptyset$.

The difference between conditioning and natural revision actually comes down to a difference in interpreting the input information. The possibility distribution π describes the normal course of things, pointing out the most usual situations, and the belief set selects the accepted default propositions, those which are true in all the usual situations. Conditioning by A means i) either that A is a piece of evidence about a case under study, and eliminating $\bar{A}$-worlds corresponds to focusing on a proper reference class, for the purpose of drawing plausible conclusion about the case at hand ii) or drastically revising the possibility distribution by a piece of generic knowledge stating that $\bar{A}$-worlds do not exist, that A should be fully certain in the revised belief set. There is the same ambiguity as in probability theory regarding Bayes conditioning. Natural revision by A means a revision of the generic knowledge by addition of a new constraint stating that "it is usual that A", resulting in accepting A. The latter is in the spirit

1 Note that λ exists since Ω is finite, so that $N(C(\pi)) > 0$.

of Jeffrey's rule and the input A is then viewed as a constraint on the uncertainty of A after the belief change. Then natural revision is a matter of handling uncertain inputs in a possibilistic setting.

3.4 - Contraction

The contraction of a possibility distribution with respect to $A \subseteq \Omega$ corresponds to forgetting that A is believed if A was previously in the belief set ($N(A) > 0 \Leftrightarrow \Pi(A) = 1 > \Pi(\bar{A})$). In such a case, the result π^-_A of the contraction must lead to a possibility measure Π^-_A such that $\Pi^-_A(A) = \Pi^-_A(\bar{A}) = 1$, i.e., complete ignorance about A. Intuitively if $\Pi(A) = \Pi(\bar{A}) = 1$ already, then we should have $\pi^-_A = \pi$. Besides if $\Pi(A) = 1 > \Pi(\bar{A})$ then we should have $\pi^-_A(\omega) = 1$ for some ω in $\bar{A}$, and especially for those ω such that $\Pi(\bar{A}) = \pi(\omega)$. If $\Pi(\bar{A}) = 1 > \Pi(A)$, i.e., $\bar{A}$ represents an accepted belief, π will be unchanged. It leads to (Dubois and Prade, 1992)

$$\begin{aligned} \pi^-_A(\omega) &= 1 \text{ if } \pi(\omega) = \Pi(\bar{A}),\ \omega \notin A \\ &= \pi(\omega) \text{ otherwise.} \end{aligned} \qquad (10)$$

By construction, π^-_A again corresponds to the idea of minimally changing π so as to forget A, when there is a unique $\omega \in \bar{A}$ such that $1 > \Pi(\bar{A}) = \pi(\omega)$. When there are several elements in $\{\omega \notin A, \pi(\omega) = \Pi(\bar{A})\}$, minimal change contractions correspond to letting $\pi^-_A(\omega) = 1$ for any selection of such situation, and π^-_A corresponds to considering the envelope of the minimal change solutions. If $\Pi(\bar{A}) = 0$, what is obtained is the fullmeet contraction (Gärdenfors, 1988). This contraction coincides exactly with a natural contraction in the sense of Boutilier and Goldzsmidt (1993).

In the classical case, Levi and Harper's identities (Gärdenfors, 1988) respectively define the revision by A as first a contraction forgetting $\bar{A}$ and then an expansion adding A, and the contraction by A as retaining the beliefs accepted both in the cognitive state π and its revision by $\bar{A}$. These results remain valid in the possibilistic setting, namely (Dubois and Prade, 1992):

$$\pi(\cdot \mid A) = (\pi^-_{\bar{A}})^+_A = \min(\mu_A, \pi^-_{\bar{A}});\ \pi^-_A = \max(\pi, \pi(\cdot \mid \bar{A})).$$

The left-hand identity can also be recovered changing conditioning into natural revision (8), provided that expansion is (9), that is, without eliminating the $\bar{A}$ worlds. In the right-hand identity natural revision (8) can be substituted to conditioning as well. An alternative contraction rule to (10) is:

$$\begin{aligned} \pi^-_A(\omega) &= \frac{\pi(\omega)}{\Pi(\bar{A})} \text{ if } \omega \in \bar{A} \\ &= \pi(\omega) \text{ otherwise,} \end{aligned}$$

that is the companion to the numerical Bayesian-like possibilistic revision rule. Again Levi's and Harper's identity hold between the two Bayesian-like rules of conditionning and contraction.

4 - Uncertain Inputs

In the following we shall consider a possibilistic cognitive state π, and some uncertain input information of the form (A,α) where α is a degree of necessity. The input is not modelled in the same way whether it is a constraint or an unreliable input.

4.1 - Uncertain Inputs as Constraints

In this case, the uncertain input (A,α) forces the revised cognitive state π' to satisfy $N'(A) = \alpha$ (i.e., $\prod'(A) = 1$ and $\prod'(\bar{A}) = 1 - \alpha$) and the following belief change rule is in the spirit of Jeffrey's rule and respects these constraints

$$\pi(\omega \mid (A,\alpha)) = \max(\pi(\omega \mid A), (1 - \alpha) * \pi(\omega \mid \bar{A})) \qquad (11)$$

where $*$ = min or product according to whether $\pi(\omega \mid A)$ is the ordinal or Bayesian-like revised possibility distribution. To see why, notice that it achieves two changes in parallel: a revision (conditionning) on A and another one on $\bar{A}$. Then, the distribution on $\bar{A}$ is denormalized, so as to satisfy $\prod'(\bar{A}) = 1 - \alpha$ for $\pi' = \pi(\cdot \mid (A,\alpha))$. When $*$ = product, all plausibility levels are proportionally shifted down, while only the greatest ones are decreased using minimum. The most plausible worlds in A become fully plausible, the most plausible situations in $\bar{A}$ are forced to level $1 - \alpha$ and, when $*$ = min all situations that were originally more plausible than $1 - \alpha$, if any are forced to level $1 - \alpha$ as well. This operation minimizes changes of the possibility levels of situations so as to accommodate the constraint $N'(A) = \alpha$. Moreover, firmly entrenched beliefs are left untouched:

Lemma 1: If $*$ = min then $N'(B) = \min(N(B \mid A), \max(\alpha, N(B \mid \bar{A}))$. ■

Proof: Just use the definition of necessity and the equation (11).

Lemma 2: If $*$ = min and $N(B) > \max(N(A), N(\bar{A}))$ then $N(B) = \min(N(B \mid A), N(B \mid \bar{A}))$,

Proof: Using the duality between N and $\prod$, and the max-decomposition of $\prod$, it can be checked that $N(B) > \max(N(A), N(\bar{A}))$ is equivalent to

$$\prod(\bar{B}) < \min(\prod(A),\prod(\bar{A})) \Leftrightarrow \max(\prod(A \cap \bar{B}), \prod(\bar{A} \cap \bar{B})) < \min(\max(\prod(A \cap B), \prod(A \cap \bar{B})), \max(\prod(\bar{A} \cap B), \prod(\bar{A} \cap \bar{B}))).$$

Letting $x = \prod(A \cap B)$, $y = \prod(\bar{A} \cap B)$, $z = \prod(A \cap \bar{B})$, and $t = \prod(\bar{A} \cap \bar{B})$, it reads

$$\begin{cases} z < \max(x,z) \\ z < \max(y,t) \\ t < \max(x,z) \\ t < \max(y,t) \end{cases}$$

which is equivalent to

$$\begin{cases} x > \max(z,t) \\ y > \max(z,t) \end{cases}$$

which means

$\prod(A \cap B) > \max(\prod(A \cap \bar{B}), \prod(\bar{A} \cap \bar{B}))$ and $\prod(\bar{A} \cap B) > \max(\prod(A \cap \bar{B}), \prod(\bar{A} \cap \bar{B}))$.

Hence $\prod(\bar{B} \mid A) = \prod(A \cap \bar{B}) < \prod(A \cap B) = \prod(A)$ and $\prod(\bar{B} \mid \bar{A}) = \prod(\bar{A} \cap \bar{B}) < \prod(\bar{A} \cap B) = \prod(\bar{A})$. Hence $N(B) = \min(N(A \cup B), N(\bar{A} \cup B)) = 1 - \max((\prod(A \cap \bar{B}), \prod(\bar{A} \cap \bar{B})) = 1 - \max(\prod(\bar{B} \mid A), \prod(\bar{B} \mid \bar{A})) = \min (N(B \mid A), N(B \mid \bar{A}))$.

Note that for $*$ = product, this result does not hold because, generally, one of the equalities $\prod(\bar{B} \mid A) = \prod(A \cap \bar{B})$ and $\prod(\bar{B} \mid \bar{A}) = \prod(\bar{A} \cap \bar{B})$ does not hold due to the proportional renormalization on A or its complement. Lemma 2 holds in this case only if $N(A) = N(\bar{A}) = 0$.

■

Proposition 1: If $*$ = min and if an event B is such that $N(B) > \max(N(A), N(\bar{A}), \alpha)$, then $N'(B) = N(B \mid (A,\alpha)) = N(B)$.

Proof: Note that $N'(B) = \min(N(B \mid A), \max(\alpha, N(B \mid \bar{A}))$, using Lemma 1. Using Lemma 2 if $N(B) = N(B \mid A) \leq N(B \mid \bar{A})$, we find $N'(B) = \min(N(B), \max(\alpha, N(B \mid \bar{A})) = N(B)$ and if $N(B) = N(B \mid \bar{A}) \leq N(B \mid A)$, we find $N'(B) = \min(N(B \mid A), \max(\alpha, N(B)) = N(B)$ again if $N(B) \geq \alpha$.

When $\alpha = 1$, $\pi(\omega \mid (A,\alpha)) = \pi(\omega \mid A)$, that is the revision by conditioning is obtained, but when $\alpha = 0$, we obtain a possibility distribution less specific than π, such that $N(A) = N(\bar{A}) = 0$, which is a kind of contraction different from an AGM contraction in the sense of (10), and more drastic. Indeed with (11) the uncertain input (A,0) is interpreted as enforcing total ignorance on A, while contracting A in the AGM sense leaves a cognitive state that believes $\bar{A}$ untouched, as can be checked on (10). When $\alpha > 0$ and $*$ = min, the possibilistic Jeffrey-like rule exactly coincides with what Williams (1994b) calls an "adjustment" (see Subsection 4.2 below).

Rule (9) can be extended to a set of input constraints $\prod(A_i) = \lambda_i$, i = 1,n, where $\{A_i, i = 1,n\}$ forms a partition of Ω, such that $\max_{i=1,n} \lambda_i = 1$ (normalisation). It gives the following form of the Jeffrey-like rule where $*$ = minimum or product whether $\pi(\omega \mid A_i)$ is ordinal or numerical:

$$\pi(\omega \mid \{(A_i,\lambda_i)\}) = \max_i \lambda_i * \pi(\omega \mid A_i). \quad (12)$$

These qualitative counterpart of convex mixtures were introduced by Dubois and Prade(1990). They that also underlie qualitative decision theory (Dubois and Prade, 1995; Dubois et al., 1996).

Belief revision rule (11-12) for * = product has been proposed under a quite different appearance by Spohn (1988), using an ordinal conditional function κ valued on the set $\mathbb{N}$ of natural integers. An ordinal conditional function is defined from a set of integers $\{\kappa(\omega), \omega \in \Omega\}$ and is a set-function such that $\forall A \subseteq \Omega$, $\kappa(A) = \min\{\kappa(\omega) \mid \omega \in A\}$. $\kappa(\omega)$ can be viewed as a degree of impossibility of ω. Letting $\prod_\kappa(A) = 1 - N_\kappa(\bar{A}) = 2^{-\kappa(A)}$, it is easy to check that $\pi_\kappa(\omega)$ is equal to $2^{-\kappa(\omega)}$, where π_κ is the possibility distribution associated with $\prod_\kappa$. Spohn (1988) defines two conditioning concepts:

– $\forall \omega \in A$, $\kappa(\omega \mid A) = \kappa(\omega) - \kappa(A)$

– the (A,n)-conditionalization of κ (conditioning by an uncertain input $\kappa'(A) = n$)

$\kappa(\omega \mid (A,n)) = \kappa(\omega \mid A)$ if $\omega \in A$;
$\kappa(\omega \mid (A,n)) = n + \kappa(\omega \mid \bar{A})$ if $\omega \in \bar{A}$.

whose possibilistic counterparts are equations (7), and (11) with $*$ = product:

$$\pi_\kappa(\omega \mid A) = \frac{\pi_\kappa(\omega)}{\prod_\kappa(A)} \text{ if } \omega \in A;\ \pi_\kappa(\omega \mid A) = 0 \text{ otherwise;} \tag{13}$$

$$\pi_\kappa(\omega \mid (A,n)) = \frac{\pi_\kappa(\omega)}{\prod_\kappa(A)} \text{ if } \omega \in A; \tag{14}$$
$$= (1 - \alpha) \cdot \frac{\pi_\kappa(\omega)}{\prod_\kappa(A)} \text{ if } \omega \notin A$$

with $\alpha = 1 - 2^{-n}$. The counterpart of (11) can be extended to an input ordinal conditional function κ' defined on the partition $\{A_1, \ldots, A_n\}$: $\kappa(\omega \mid \kappa') = \kappa'(A_i) + \kappa(\omega \mid A_i)$, $\forall \omega \in A_i$, i = 1,n. This rule can be exactly mapped to the possibilistic belief change rule (12) where $*$ = product and $\lambda_i = 2^{-\kappa'(A_i)}$.

4.2 - Unreliable Inputs

If (A,α) is viewed as an unreliable input, it means that α is the degree of certainty that A is the actual input supplied by the external source; but maybe the source does not properly works and then not input is considered. The unreliable input is then represented by the weighted *nested* pair of subsets $F = \{(A,1), (\Omega, 1 - \alpha)\}$ where the weights denote degrees of possibility. The revised cognitive state $\pi(\cdot \mid F)$ is defined by a possibilistic mixture as

$$\pi(\omega \mid F) = \max(\pi(\omega \mid A), \pi(\omega) * (1 - \alpha)). \tag{15}$$

It means that $\pi(. \mid A)$ obtains with possibility 1, and π is kept with possibility $1 - \alpha$. It can be viewed as a possibilistic conditioning generalized to the fuzzy event F. Computing the effect of

unreliable inputs as in (15) is questionable. One might wish to keep the result non-deterministic until the source is checked. Note the difference with (11): there is no conditioning on $\bar{A}$ ($\pi(\omega) = \pi(\omega \mid \Omega)$). However, contrary to (11), the equality $N(A \mid F) = \alpha$ is not warranted since, with (15), $N(A \mid F) = \min(N(A \mid A), \max(N(A), \alpha)) = N(A)$ whenever $N(A) > \alpha$. Lastly, $\pi(\omega \mid F) = \pi(\omega)$ if $\alpha = 0$ since then $F = \Omega$: no revision takes place. This behavior is very different from the case when the uncertain input is taken as a constraint. (15) can be formally generalized to any fuzzy input F, viewed as a weighted collection of nested sets (Dubois and Prade, 1991b).

4.3 - Adjustments

Williams (1994b) has defined a general form of belief change she calls "transmutations", in the setting of Spohn's functions. Given an uncertain input (A,n) taken as a constraint and a Spohn function κ describing the agent's a priori cognitive state, a transmutation of κ by (A,n) produces a Spohn function κ' such that $\kappa'(\bar{A}) = n$ and $\kappa'(A) = 0$, i.e., the degree of acceptance of A is enforced to level n. Clearly, this notion makes sense in the possibilistic setting, where a transmutation of a cognitive state π into π' using input $N'(A) = \alpha$ corresponds to enforcing $N'(A) = \alpha$ and $N'(\bar{A}) = 0$. Williams (1994a) has introduced an elementary and qualitative type of transmutation called an adjustment. An adjustment of κ by (A,n) is either a contraction κ^{-}_{A} if $n = 0$ or another belief change operation, defined as follows:

$$\begin{aligned}
\kappa^{*}_{(A,n)} &= \kappa^{-}_{A} \text{ if } n = 0 \\
&= (\kappa^{-}_{A})^{x}_{(A,n)} \text{ if } 0 < n < \kappa(\bar{A}) \\
&= \kappa^{x}_{(A,n)} \text{ otherwise}
\end{aligned}$$

where

$$\begin{aligned}
\kappa^{-}_{A}(\omega) &= 0 \text{ if } \omega \in \bar{A} \text{ and } \kappa(\omega) = \kappa(\bar{A}) \\
&= \kappa(\omega) \text{ otherwise} \\
\kappa^{x}_{(A,n)}(\omega) &= 0 \text{ if } \omega \in A \text{ and } \kappa(\omega) = \kappa(A) \\
&= \kappa(A) \text{ if either } \omega \in A \text{ and } \kappa(\omega) \neq \kappa(A) \text{ or } \omega \in \bar{A} \text{ and } \kappa(\omega) > n \\
&= n \text{ otherwise.}
\end{aligned}$$

In order to clarify the meaning of this intricate definition of adjustment, let us map it to the possibilistic setting using the transformation pointed out in section 4.1 (with $\alpha = 1 - 2^{-n}$) when $n \neq 0$. It can be shown (see Appendix) that the result of this translation is precisely

$$\pi^{*}_{(A,\alpha)}(\omega) = \max(\pi(\omega \mid A), \min(1 - \alpha, \pi(\omega \mid \bar{A})))$$

we recognize (11) for $\alpha > 0$. This result leads us to simplifying the formulation of Williams' adjustment as follows:

$$\kappa^{*}_{(A,n)}(\omega) = \min(\kappa^{*}(\omega \mid A), \max(n, \kappa^{*}(\omega \mid \bar{A})) \qquad (16)$$

where

$$\kappa^*(\omega \mid A) = +\infty \text{ if } \omega \notin A;$$
$$\kappa^*(\omega \mid A) = \kappa(\omega) \text{ if } \kappa(\omega) > \kappa(A); \kappa^*(\omega \mid A) = 0 \text{ if } \kappa(\omega) = \kappa(A).$$

$\kappa^*(\omega \mid A)$ is the Spohnian version of the qualitative form of possibilistic conditioning. Turning $\kappa^*(\omega \mid A)$ into $\kappa(\omega \mid A) = \kappa(\omega) - \kappa(A)$ when $\omega \in A$ and max into sum, one gets the Spohnian (A,n)-conditioning rule. Note that when n = 0, rule (16) gives a more drastic contraction since both A and its complement are erased from the belief state. $\kappa^*_{(A,n)}$ as per (16) does not recover the contraction κ^-_A, exactly for the same reason as its possibilistic counterpart (11). It seems somewhat artificial to enforce $\kappa^*_{(A,n)} = \kappa^-_A$ when n = 0, as done by Williams (1994b) on the basis of (16) or (11). The question is whether an input information of the form (A,0) is or is not equivalent to the pair of constraints {(A,0), ($\bar{A}$,0)}. If it is not then the result of input (A,0) should indeed be a contraction. In rule (11), the input {(A,α)} is taken as {(A,α), ($\bar{A}$,0)} when $\alpha > 0$ (and this is also true for adjustments). So this interpretation remains when $\alpha = 0$.

4.4 - Natural Revision as Revision by an Uncertain Input

In the case of Boutilier (1993)'s natural revision, the enforced input can be modelled as $N'(A) > 0$. It might be tempting to interpret it as a disjunctive input "$N'(A) = 1$ or $N'(A) = \alpha_n$ or… or $N'(A) = \alpha_2$" with scale V be made of n + 1 levels $\lambda_1 = 1 > \lambda_2 > \ldots > \lambda_n > 0$, and $\alpha_i = 1 - \lambda_i$. It suggests that the natural revision might be expressed by a series of adjustments $\pi(\omega \mid (A,\alpha_i))$ and forming their disjunction:

$$\pi^*_A(\omega) = \max_{i=2,\, n+1} \pi(\omega \mid (A,\alpha_i)) = \max(\pi(\omega \mid A), \lambda_2 * \pi(\omega \mid \bar{A}), \ldots, \lambda_n * \pi(\omega \mid \bar{A}))$$
$$= \max(\pi(\omega \mid A), (1 - \alpha_2) * \pi(\omega \mid \bar{A})).$$

For $*$ = product, this rule is found in Darwiche and Pearl (1994) under the name R-conditioning. It coincides indeed with natural revision provided that $*$ = min, N(A) = 0, and that no situation $\omega \in \bar{A}$ has a priori plausibility level $\pi(\omega) = \lambda_2$, so that when $N(\bar{A}) > 0$, i.e., $\Pi(A) < 1$ and $\Pi(\bar{A}) = 1$, we have $\Pi^*_A(\bar{A}) = \lambda_2 > \pi(\omega)$ for all ω such that $\pi(\omega) < 1$. Then natural revision comes down to raising the plausibility of the most plausible situations in A to 1 and forcing the most plausible situations in $\bar{A}$ down to λ_2 if they had plausibility 1 previously. If $\pi(\omega) = \lambda_2$ for situations $\omega \in \bar{A}$, it is always possible to let $\Pi^*_A(\bar{A}) = \lambda'$ were $1 > \lambda' > \lambda_2$ if V is infinite (= [0,1]). It gives the adjustment rule (11) for $\alpha = 1 - \lambda'$.

Proposition 2: When $*$ = min, natural revision (8) is equivalent to the adjustment rule (11) with input (A, α) where $0 < \alpha < \min\{N(B) \mid B \text{ such that } N(B) > 0\}$ if and only if N(A) = 0.

Proof: Note that $0 < \alpha < \min\{N(B) \mid B \text{ such that } N(B) > 0\}$ implies $\max\{\pi(\omega') \mid \omega' \notin C(\pi)\} < 1 - \alpha < 1$; moreover, N(A) = 0 means that A does not contain $C(\pi)$.

Assume $\omega \notin C(\pi) \cup \{\omega \in A, \pi(\omega) = \prod(A)\}$; then $\pi(\omega \mid A) = \pi(\omega)$ if $\omega \in A$, and $\pi(\omega \mid A) = 0$ otherwise, and $\pi(\omega \mid \bar{A}) = \pi(\omega)$ if $\omega \notin A$, $\pi(\omega) \neq \prod(\bar{A})$, $\pi(w \mid \bar{A}) = 1$ if $\omega \notin A$, $\pi(\omega) = \prod(\bar{A})$ and $\pi(\omega \mid \bar{A}) = 0$ otherwise. So

if $\omega \in A$, $\pi(\omega \mid (A,\alpha)) = \max(\pi(\omega), \min(1 - \alpha, 0)) = \pi(\omega)$;
if $\omega \notin A$, $\pi(\omega) \neq \prod(\bar{A})$, $\pi(\omega \mid (A,\alpha)) = \max(0, \min(1 - \alpha, \pi(\omega))) = \pi(\omega)$ since
$\pi(\omega) \leq \max\{\pi(\omega') \mid \omega' \notin C(\pi)\} < 1 - \alpha$;
if $\omega \notin A$, $\pi(\omega) = \prod(\bar{A})$, $\pi(\omega \mid (A,\alpha)) = \max(0, \min(1 - \alpha, 1)) = 1 - \alpha$.

Assume $\omega \in \bar{A} \cap C(\pi) \neq \emptyset$. Then $\pi(\omega \mid A) = 0$ and $\pi(\omega \mid \bar{A}) = 1$. Hence $\pi(\omega \mid (A,\alpha)) = 1 - \alpha$.

Assume $\omega \in A$, and $\pi(\omega) = \prod(A)$. Then $\pi(\omega \mid A) = 1$ and $\pi(\omega \mid \bar{A}) = 0$. So, $\pi(\omega \mid (A,\alpha)) = \max(1, \min(1 - \alpha, 0)) = 1$.

Note that there is only one case when (8) and (11) give distinct possibility distributions: this is when A contains $C(\pi)$. Then $\pi(\omega \mid (A,\alpha)) = 1 - \alpha > \pi(\omega)$ for all $\omega \notin A$, with $\pi(\omega) = \prod(\bar{A})$. But this case is ruled out if and only if $N(A) = 0$. ■

When $N(A) > 0$, this is the case of trivial expansion for which $\pi^*_A = \pi$, and $\pi(\omega \mid (A,\alpha)) = \pi(\omega)$ except for $\omega \notin A$, and $\pi(\omega) = \prod(\bar{A})$. Those elements are just moved up without reaching 1. Hence the ordering of situations is generally modified since there may exist an element $\omega_1 \in A$ and $\omega_2 \notin A$ such that $\pi(\omega_1) = \pi(\omega_2) = \prod(\bar{A}) < 1$ and then $\pi(\omega_1 \mid (A,\alpha)) = \pi(\omega_1) < \pi(\omega_2 \mid (A,\alpha)) = 1 - \alpha$. So some elements of $\bar{A}$ become more plausible for the sake of letting $N'(A)$ be positive but as small as possible. This is a kind of contraction by A, while revising π by $N'(A) > 0$ when $N(A) > 0$ should leave the cognitive state as it was. This is because the disjunctive input "$N'(A) = 1$ or $N'(A) = \alpha_n$ or… or $N'(A) = \alpha_2$" (with $\alpha_2 = \alpha$) that expresses $N'(A) > 0$ is interpreted as $N'(A) = \alpha$ by the adjustment rule as being the weakest statement among them. But $N'(A) = \alpha$ is informationally stronger than $N'(A) > 0$ since the set of possibility distributions which satisfy the former is a proper subset of the set of possibility distributions that satisfy the latter. The adjusment rule is a simple means of enforcing the plausibility of some situations to a prescribed level, while natural revision no longer enforces anything when the input information is already accepted to any level.

5 - Syntactic Revisions of Possibilistic Belief Bases

Revision tools presented at the semantic level in the previous sections cannot be applied to knowledge bases right away. They should first be expressed at the syntactic level. In this section we consider layered knowledge bases made of weighted propositional formulas in possibilistic logic (Dubois, Lang and Prade, 1994). This question has also be addressed by Williams (1995, 1996) in an integer-valued setting, and the terminology of epistemic entrenchments (Gärdenfors, 1988).

5.1 - Possibilistic Logic

The syntax of possibilistic logic consists of sentences in the first order calculus to which are attached lower bounds on the degree of necessity (or possibility) of these sentences. In this section, degrees of uncertainty belong to a totally ordered set V with bottom 0 and top 1. Here we consider only the fragment of possibilistic logic with propositional sentences to which lower bounds of degrees of necessity are attached. If p is a propositional sentence, $(p\ \alpha)$ is the syntactic counterpart of the semantic constraint $N([p]) \geq \alpha$, where [p] denotes the set of models of p.

A possibilistic belief base is a finite set $\mathcal{K} = \{(p_i\ \alpha_i),\ i = 1,m\}$ of weighted (propositional) formulae that contain beliefs explicitly held by an agent. The weight indicates the agent's confidence in the corresponding formula. Note that any belief base $\mathcal{B}$ (i.e., set of propositional sentences) equipped with a complete partial ordering $\geq$ can be mapped to a possibilistic belief base, changing p and q in $\mathcal{B}$ into $(p\ \alpha)$ and $(q\ \beta)$ such that $\alpha \geq \beta$ if and only if $p \geq q$. As already pointed out the unit interval could be changed into any bounded, totally ordered set; the possibility/necessity duality is then expressed by reversing the ordering.

Reasoning in possibilistic logic is done by means of an extension of the resolution principle to weighted clauses:

$$(c\ \alpha);\ (c'\ \beta) \vdash (\mathrm{Res}(c,c')\ \min(\alpha,\beta))$$

where c and c' are propositional clauses, and Res(c,c') is their resolvent. For instance, $(\neg p \vee q\ \alpha);\ (p \vee r\ \beta) \vdash (q \vee r\ \min(\alpha,\beta))$. This inference rule presupposes that when, in a possibilistic formula $(p\ \alpha)$, p is not in clausal form, it can be turned into a set $\{(c_i\ \alpha),\ i = 1,n\}$ of weighted clauses such that p is equivalent to $c_1 \wedge c_2 \wedge \ldots \wedge c_n$. This is justified by the semantics of propositional logic, and by the fact that $N(p) \geq \alpha$ is equivalent to $N(c_i) \geq \alpha$, for i = 1,n (from now on, we write N(p) instead of N([p]) for short). Inference from a possibilistic belief base is denoted $\mathcal{K} \vdash (p\ \alpha)$, and is short for $\mathcal{K} \cup \{(\neg p\ 1)\} \vdash (\bot\ \alpha)$ (refutation method). $\mathcal{K} \vdash (\bot\ \alpha)$ can be checked by means of repeated uses of the resolution principle until the empty clause is attained, with some positive weight. The degree of inconsistency $\mathrm{inc}(\mathcal{K})$ is then defined by $\max\{\alpha \mid \mathcal{K} \vdash (\bot\ \alpha)\}$.

The set of situations in which a possibilistic logic sentence $(p\ \alpha)$ is true is a fuzzy set $[p\ \alpha]$ on Ω defined by

$$\begin{aligned} \mu_{[p\ \alpha]}(\omega) &= 1 \text{ if } \omega \in [p] \\ &= 1 - \alpha \text{ otherwise,} \end{aligned}$$

$\mu_{[p\ \alpha]}$ is the least specific (i.e., the greatest) possibility distribution π such that $N(p) = \inf_{\omega \notin [p]} 1 - \pi(\omega) \geq \alpha$. A possibilistic belief base $\mathcal{K} = \{(p_i\ \alpha_i),\ i = 1,m\}$ (with the semantics $N(p_i) \geq \alpha_i$) is represented by the possibility distribution

$$\pi(\omega) = \min_{i=1,m} \max(\mu_{[p_i]}(\omega), 1 - \alpha_i) \quad (17)$$

which extends $[\mathcal{K}] = [p_1] \cap [p_2] \cap \ldots \cap [p_m]$ from a set of sentences to a set of weighted sentences. π is the least specific possibility distribution such that $\forall\ (p\ \alpha) \in \mathcal{K}$, $N(p) \geq \alpha$, where N is computed with π. The specificity ordering is defined by $\pi \leq \pi'$ which means that π is at least as informative (= specific) as π'. Semantic entailment is defined in terms of this specificity ordering. Namely, $\mathcal{K} \models (p\ \alpha)$ if and only if $\pi \leq \max(\mu_{[p]}, 1 - \alpha)$. This notion of semantic entailment is exactly the one of Zadeh (1979).

Note that the above framework can be equivalently expressed in terms of Spohnian functions. Instead of using weights in the unit interval, we can use integers from 0 to n where n is the number of layers in the ordered belief base. An obvious understanding of a formula p being in layer i is that $\kappa(\neg p) \geq i$, which can be made equivalent to $N(p) \geq \alpha$, where $i = -\text{Log}_2(1 - \alpha)$. The most entrenched sentences are then in layer n (see for instance Williams, 1995). The minimally specific possibility distribution π that is induced by an ordered belief base is then changed into a ranking of the possible words, namely, letting $\ell(p_i)$ the layer number of p_i

$$\begin{aligned} \kappa(\omega) &= 0 \text{ if } \omega \text{ satisfies all } p_i\text{'s in } \mathcal{K} \\ &= \max_{i:\omega \models \neg p_i} \ell(p_i) \text{ otherwise.} \end{aligned}$$

This is called "minimal ranking function" by Pearl (1990) and it corresponds to the minimally specific possibility distribution π.

Possibilistic logic is sound and complete with respect to refutation based on resolution (Dubois et al., 1994). Namely, it can be checked that the degree of inconsistency of a knowledge base is reflected by the extent to which the induced possibility distribution is subnormalized:

$$\text{inc}(\mathcal{K}) = 1 - \max_{\omega \in \Omega} \pi(\omega)$$

This property is crucial in proving the equivalence between the semantic and the syntactic entailments:

$$\mathcal{K} \vdash (p\ \alpha) \text{ if and only if } \mathcal{K} \models (p\ \alpha).$$

When $\text{inc}(\mathcal{K}) > 0$, $\mathcal{K}$ is said to be partially inconsistent; when $\text{inc}(\mathcal{K}) = 1$ it is completely inconsistent. Consistent possibilistic belief bases are such that $\pi(\omega) = 1$ for some $\omega \in \Omega$. Consistency of $\mathcal{K}$ is equivalent to the consistency of the classical knowledge base obtained by removing the weights.

When $\mathcal{K}$ is consistent, we can define an ordered belief set generated by $\mathcal{K}$ as $\text{Cons}(\mathcal{K}) = \{(p\ \alpha),\ \mathcal{K} \vdash (p\ \alpha)\}$. The set function N such that $N(p) = \alpha$ for all $(p\ \alpha)$ in $\text{Cons}(\mathcal{K})$ (and 0 otherwise) is a necessity measure. The ordering $\geq_N$ generated on $\text{Cons}(\mathcal{K})$ is (up to some limit conditions) an epistemic entrenchment ordering in the sense of Gärdenfors

(1988) and an expectation ordering in the sense of Gärdenfors and Makinson (1994); see Dubois and Prade (1991a).

Note that the restriction to $\mathcal{K}$ of the ordering on Cons($\mathcal{K}$) may fail to respect the original ordering on $\mathcal{K}$. Indeed the restriction of a generated ordered belief set to a belief set $\mathcal{K}$ satisfies a coherence property with respect to necessity functions:

For any $(p\ \alpha) \in \mathcal{K}$ and any subset $\mathcal{B} = \{(p_i\ \beta_i), i = 1,m\}$ of $\mathcal{K}$,
if $\mathcal{B} \vdash (p\ \alpha)$ then it does not hold that $\alpha < \min_{i=1,m} \beta_i$ (N-coherence)

The above condition is requested by Rescher (1976) because he is the first to propose layered belief bases and to exploit an extended modus ponens rule whereby the weight of a consequence is the minimum between the weights of the premises. Rott (1991) calls this condition EE-coherence (for empstemic entrenchment). The restriction of an epistemic entrenchment ordering to a belief base (or equivalently a N-coherent ordering) is called "ensconcement" by Williams (1994a). Starting with any ordering on a belief base $\mathcal{K}$, possibilistic inference restores the above condition on it and builds an ensconcement.

When $\mathcal{K}$ is partially inconsistent, i.e., inc($\mathcal{K}$) > 0, non-trivial deductions can still be made from $\mathcal{K}$, namely all $(p\ \alpha)$ such that $\mathcal{K} \vdash (p\ \alpha)$, and for which $\alpha >$ inc($\mathcal{K}$). Indeed $(p\ \alpha)$ is then the consequence of a consistent subpart of $\mathcal{K}$, say $\mathcal{K}^* = \{(p\ \alpha), \alpha >$ inc($\mathcal{K}$)$\}$. Non-trivial inference of p from $\mathcal{K}$ is denoted $\mathcal{K} \vdash_{\text{pref}} p$ and means $\mathcal{K}^* \vdash (p\ \alpha)$ for some positive α. When $\mathcal{K}$ is partially (but not totally) inconsistent, the associated consistent ordered belief set is $\text{Cons}_{\text{pref}}(\mathcal{K}) = \{(p\ \alpha),\ \mathcal{K} \vdash_{\text{pref}} p\}$. In that case $\max_{\omega \in \Omega} \pi(\omega) = 1 - \text{inc}(\mathcal{K}) < 1$ and $\forall p$, $\min(N(p),N(\neg p)) = \text{inc}(\mathcal{K})$. Let $\tilde{\pi}$ be a possibility distribution on Ω defined by

$$\begin{aligned} \tilde{\pi}(\omega) &= \pi(\omega) \text{ if } \pi(\omega) < 1 - \text{inc}(\mathcal{K}) \\ &= 1 \text{ otherwise,} \end{aligned}$$

then it is easy to verify (Dubois, Lang and Prade, 1994) that the necessity measures N and $\tilde{N}$ based on π (defined by (17)) and $\tilde{\pi}$ respectively are related by the following relation

$$N(p) > N(\neg p) \Rightarrow \tilde{N}(p) = N(p) > N(\neg p) = \text{inc}(\mathcal{K}) > \tilde{N}(\neg p) = 0.$$

In fact, we have that $\tilde{\pi}$ is the possibility distribution induced by $\mathcal{K}^*$.

5.2 - Revising Ordered Belief Bases: Implementing the AGM Revision

Belief change in possibilistic logic can then be envisaged as the syntactic implementation on a belief base, of change operations for possibility distributions. Let $\mathcal{K}$ be a consistent set of possibilistic formulae. Expansion of $\mathcal{K}$ by p consists of forming $\mathcal{K} \cup \{(p\ 1)$. It supplies non

trivial results provided that inc($\mathcal{K} \cup \{(p\ 1)\}$) = 0. Clearly, the possibility distribution π' that restricts the fuzzy set of situations that satisfy $\mathcal{K} \cup \{(p\ 1)\}$ is $\pi' = \min(\pi, \mu_{[p]})$.

Let us consider the case when $\mathcal{K}$ is consistent, but $\mathcal{K}' = \mathcal{K} \cup \{(p\ 1)\}$ is not, and let α = inc($\mathcal{K} \cup \{(p\ 1)\}$) > 0. The following identity is easy to prove (e.g., Dubois et al., 1994):

$$\exists\ \beta > 0,\ \mathcal{K} \cup \{(p\ 1)\} \vdash_{\text{Pref}} (q\ \beta) \text{ if and only if } N(q \mid p) > 0$$

where $N(q \mid p)$ is the necessity measure induced from $\pi(\cdot \mid [p])$, i.e., the possibility distribution expressing the contents of $\mathcal{K}$, revised with respect to the set of models of p. Indeed let π' be the possibility distribution on Ω induced by $\mathcal{K}'$, then

$$\pi' = \min(\pi, \mu_{[p]})$$
$$0 < \max_{\omega \in \Omega} \pi'(\omega) = 1 - \alpha < 1$$

and the possibility distribution $\tilde{\pi}'$ induced from the consistent part $\mathcal{K}'^*$ of $\mathcal{K}'$ (made of sentences whose weight is higher than α), is defined as

$$\begin{aligned} \tilde{\pi}'(\omega) &= \pi(\omega) \text{ if } \omega \in [p] \text{ and } \pi'(\omega) < 1 - \alpha \\ &= 1 \text{ if } \omega \in [p] \text{ and } \pi(\omega) = 1 - \alpha \\ &= \pi'(\omega) = 0 \text{ otherwise.} \end{aligned}$$

Hence the possibility distribution induced by $\mathcal{K}'^*$ is $\tilde{\pi}' = \pi(\cdot \mid [p])$, the result of revising π by [p] using the ordinal conditioning method of Section 3.2.

The possibilistic revision rule based on ordinal conditioning can thus be expressed directly on the belief base $\mathcal{K}$ by the following method (Dubois and Prade, 1992), called "brutal theory base operator" by Williams(1994a) :

i) adding p above the top layer of $\mathcal{K}$
ii) deleting all sentences whose level is below the inconsistency level α = inc($\mathcal{K} \cup \{(p\ 1)\}$).

and then a revised belief base $\mathcal{K}*_p = (\mathcal{K} \cup \{(p\ 1)\})^*$ is obtained such that Cons($\mathcal{K}*_p$) is the revision of the belief set Cons($\mathcal{K}$) by p. This belief base revision is rather drastic since all sentences $(p_i\ \alpha_i)$ with weights $\alpha_i \leq \alpha$ are thrown away, and replaced by (p 1). However it is syntax-independent. Note that this revision method works even if the weights attached to formulae are not N-coherent. Suppose $(q\ \beta) \in \mathcal{K}$ and $\mathcal{K} \vdash (q\ \gamma)$ with $\gamma > \beta$. It means that $(q\ \beta)$ can be deleted from $\mathcal{K}$ without altering its fuzzy set of models. The revision being syntax-independent, the presence or the absence of $(q\ \beta)$ in $\mathcal{K}$ will not affect the fuzzy set of models of $\mathcal{K}*_p$.

Note that when $\Pi(p) > 0$, $N(q \mid p) > 0$ is equivalent to $N(\neg p \vee q) > N(\neg p \vee \neg q)$, i.e., in terms of epistemic entrenchment (Gärdenfors, 1988), $\neg p \vee q$ is more entrenched than $\neg p \vee \neg q$, and corresponds to a characteristic condition for having q in the (ordered) belief set obtained by revising Cons($\mathcal{K}$) with respect to p, in Gärdenfors (1988). We just showed that this revision is easily implemented in the possibilistic belief base itself, without making the underlying ordered belief set explicit. This result goes against the often encountered claim that working with epistemic entrenchment orderings would be intractable. Note that the revision produces a new epistemic entrenchment ordering which is in accordance with the min-based conditioning, but not the natural revision.

5.3 - Parsimonious Revisions of Ordered Belief Bases

A more parsimonious revision scheme for possibilistic belief bases $\mathcal{K}$ receiving an input p is to consider all subsets of $\mathcal{K}$ that fail to infer $(\neg p\ \alpha)$ for $\alpha > 0$. If $\mathcal{H}$ is such a subset, then the result of the revision could be $\mathcal{H} \cup \{(p,1)\}$. We may take advantage of the ordering in $\mathcal{K}$ to make the selection. Namely we may restrict ourselves to $\mathcal{H}$ such that $\forall\ (q\ \alpha) \notin \mathcal{H}$, $\mathcal{H} \cup \{(q\ \alpha), (p\ 1)\} \vdash (\bot\ \alpha)$ in the possibilistic logic sense, i.e., $(q\ \alpha)$ is involved in the contradiction. This proposal, made independently in Dubois et al. (1992) corresponds to selecting a preferred subbase in the sense of Brewka (1989). This selection process leads to a unique solution if $\mathcal{K}$ is totally ordered. In case of ties, further refinement can be made using a lexicographic ordering of the weights of the sentences not in $\mathcal{H}$, as proposed in Dubois et al. (1992). These revision processes, that also relate to Nebel (1992)'s syntax-based revision schemes, are systematically studied in Benferhat et al. (1993, 1995).

This kind of revision process cannot be expressed at the semantic level, where all sentences in the knowledge base $\mathcal{K} \cup \{(p\ 1)\}$ have been combined into a possibility distribution on Ω, and revision is performed on the aggregated possibility distribution. Especially if $(q\ \beta) \in \mathcal{K}$ and $\beta < \text{inc}(\mathcal{K} \cup \{(p\ 1)\})$ then $\min(\pi, \mu_{[p]}) \leq \max(\mu_{[q]}, 1 - \beta)$, i.e., everything happens as if $(q\ \beta)$ had never been in $\mathcal{K}$. The alternative syntactic revision rule, explained above, breaks the minimal inconsistent subsets of $\mathcal{K} \cup \{(p\ 1)\}$ in a parsimonious way, enabling pieces of evidence like $(q\ \beta)$ to be spared when they are not involved in the inconsistency of $\mathcal{K} \cup \{(p\ 1)\}$.

Example: Consider the belief base $\mathcal{K} = \{(\neg p\ \alpha), (q\ \beta)\}$ with $\beta < \alpha$. Then

$$\begin{aligned}\pi(\omega) &= \min(\max(1 - \mu_{[p]}(\omega), 1 - \alpha), \max(\mu_{[q]}(\omega), 1 - \beta)) \\ &= 1 \text{ if } \omega \models \neg p \wedge q \\ &= 1 - \alpha \text{ if } \omega \models p \\ &= 1 - \beta \text{ if } \omega \models \neg p \wedge \neg q.\end{aligned}$$

Revising by input p at the semantic level leads to consider

$$\pi'(\omega) = \min(\pi(\omega), \mu_{[p]}(\omega)) = \begin{cases} 1 - \alpha \text{ if } \omega \models p \\ 0 \text{ otherwise.} \end{cases}$$

Hence $\pi(\omega \mid [p]) = 1$ if $\omega \models p$ and 0 otherwise. Hence $\mathcal{K}*_p = \{(p\ 1)\}$. Acting at the syntactic level, the preferred sub-base of $\{(p\ 1), (\neg p\ \alpha), (q\ \beta)\}$ that contains p is $\{(p\ 1), (q\ \beta)\}$. Note that although $\min(\pi, \mu_{[p]}) \leq \max(\mu_{[q]}, 1 - \beta)$ we no longer have $\pi(\omega \mid [p]) \leq \min(\mu_{[q]}(\omega), 1 - \beta)$, i.e., adding the low certainty formulas consistent with $\mathcal{K}*_p$ leads to a non-trivial expansion of $\pi(\cdot \mid [p])$. It points out a weakness of the semantic views of revision, which is particularly true with numerical approaches: the representation of the cognitive state is lumped, i.e., the pieces of belief are no longer available and the semantic revision process cannot account for the structure of the cognitive state that is made explicit in the ordered belief base.

5.4 - Adjustments of Ordered Belief Bases

Possibilistic base revision can be extended to the case of uncertain inputs. This has been done by Williams (1995) for her so-called adjustments. We have pointed out that if the input information is of the form $(p\ \alpha)$, an adjustment is the possibilistic counterpart of Jeffrey's rule as given by equation (11) and $* = \min$ when the weight is positive. Williams (1995) gives a rather complex recipe to achieve the adjustment on the belief base itself. Stated in terms of necessity measures, the adjustment comes down to computing, for all formulae q the degree (see Lemma 1)

$$N(q \mid (p\ \alpha)) = \min\ (N(q \mid p), \max(\alpha, N(q \mid \neg p)).$$

Let us use this equation in proposing a method for adjusting possibilistic knowledge bases. Let $\mathcal{K} = \{(p_i,\alpha_i) \mid i = 1,n\}$ be a possibilistic knowledge base and consider an input information of the form $N'(p) = \alpha$. An adjustment algorithm can proceed as follows in order to compute the adjusted knowledge base $\mathcal{K}\alpha_p$.

Adjustment of $\mathcal{K}$ by $N'(p) = \alpha$

i) compute $\mathcal{K}*_p$ and $\mathcal{K}*_{\neg p}$ using the above brutal syntactic revision method

ii) the adjusted belief base $\mathcal{K}\alpha_p$ is then defined by $\{(q\ \beta), q \in \mathcal{K}*_p \cup \mathcal{K}*_{\neg p}, \beta = \min(\beta^+, \max(\beta^-,\alpha)\}$ where β^+ is the weight of q in $\mathcal{K}*_p$ and β^- the weight of q in $\mathcal{K}*_{\neg p}$, with the understanding that the weight of a formula is 0 when it does not appear in the corresponding knowledge base.

Let us check that indeed $\mathcal{K}\alpha_p \vdash (p,\alpha)$ and is consistent. Several cases must be considered

a) $\mathcal{K} \cup \{(p,1)\}$ and $\mathcal{K} \cup \{(\neg p,1)\}$ are consistent

Then $\mathcal{K}*_p = \mathcal{K} \cup \{(p,1)\}$, and $\mathcal{K}*_{\neg p} = \mathcal{K} \cup \{(\neg p,1)\}$. Hence $(p,\alpha) \in \mathcal{K}\alpha_p$, but since $\mathcal{K}$ does not contain $\neg p$ at any non-zero level, $\neg p$ does not appear in $\mathcal{K}\alpha_p$ either. If p_i appears in $\mathcal{K}$, it appears in both $\mathcal{K}_p$ and $\mathcal{K}_{\neg p}$, and so the weight β_i of p_i in $\mathcal{K}\alpha_p$ is $\beta = \min(\alpha_i, \max(\alpha,\alpha_i)) = \alpha_i$. Moreover $\mathcal{K} \vdash (p,\alpha)$ does not hold since $\mathcal{K}$ is consistent with $(\neg p,1)$. Then we have $\mathcal{K}\alpha_p = \mathcal{K} \cup \{(p,\alpha)\}$ and $\mathcal{K}\alpha_p \vdash (p,\alpha)$.

b) $\mathcal{K} \vdash (p,\beta)$ ***for*** $\alpha < \beta$

In that case $\mathcal{K}*_p = \mathcal{K} \cup \{(p,1)\}$ and $inc(\mathcal{K} \cup \{(\neg p,1)\}) = \beta$ so $\mathcal{K}_{\neg p} = T(\mathcal{K},\beta) \cup \{(\neg p,1)\}$ where $T(\mathcal{K},\beta)$ is the result of truncating $\mathcal{K}$ by deletion of formulas whose weight is at most β. Then $\mathcal{K}^{\alpha}_p = \{(p,\alpha)\} \cup \{(p_i,\alpha_i),\ \alpha_i > \beta \text{ or } \alpha_i \leq \alpha\} \cup \{(p_i,\alpha),\ \alpha < \alpha_i \leq \beta\}$ where the formulas (p_i,α_i) with $\alpha < \alpha_i \leq \beta$ are discounted to level (p_i,α). Indeed if p_i appears in $\mathcal{K}*_p$ and $\mathcal{K}*_{\neg p}$ (because $\alpha_i > \beta$) then again α_i obtains. If p_i does not appear in $\mathcal{K}*_{\neg p}$, then its weight in $\mathcal{K}^{\alpha}_p$ is $\min(\alpha_i,\alpha)$. Clearly, $\mathcal{K}*_{(p,\alpha)} \vdash (p,\alpha)$ since all formulas of weight β that are instrumental in proving p from $\mathcal{K}$ now have weight α and all formulas of weight $\gamma \in (\alpha,\beta)$ that could be used to prove p have weight α as well. Note that even formulas p_i not involved in proving p get weight α if $\alpha < \alpha_i \leq \beta$, which indicates the brutal behavior of this method.

c) $\mathcal{K} \vdash (p,\beta)$ ***for*** $\alpha \geq \beta$

This case if similar to the previous one from a formal point of view, but is in fact of the same nature as (a) since the input corresponds to a reinforcement of the certainty of p. It is easy to check that $\mathcal{K}^{\alpha}_p = \mathcal{K} \cup \{(p,\alpha)\}$ is a valid way of adjusting $\mathcal{K}^{\alpha}_p$. If (p,β) is already explicitly in $\mathcal{K}$, then one may delete it of course.

d) $\mathcal{K} \vdash (\neg p,\beta)$

Then, since $\mathcal{K}*_p = T(\mathcal{K},\beta)$, $(p_i,\alpha_i) \in \mathcal{K}^{\alpha}_p$ for $\alpha_i > \beta$. But since $(p_i,\alpha_i) \notin \mathcal{K}*_p$ when $\alpha_i \leq \beta$, $\mathcal{K}^{\alpha}_p = \{(p,\alpha)\} \cup T(\mathcal{K},\beta)$. The syntactic change is then drastic again.

Two comments may shed some light on the interest of the above method. First natural revision is easy to capture under the form of an input of the form $N(p) \geq \alpha$ where $0 < \alpha < \min\{\alpha_i,\ (p,\alpha_i) \in \mathcal{K}\}$. The procedure for natural revision is as for adjustment in cases (a) and (d). Case (c) never occurs. In case (b), natural revision just does not affect $\mathcal{K}$. Next, the type of adjustment carried out in cases (b) and (d) is very drastic as explained above. If a more parsimonious adjustment must be achieved in case (b), it means that all minimal subsets of $\mathcal{K}$ that imply p with a level higher than α must be computed. Let $\mathcal{S}$ be such a minimal subset and $\mathcal{S} \vdash (p,\gamma)$ with $\beta \geq \gamma > \alpha$. Then for every such $\mathcal{S}$, some $(p_i, \gamma) \in \mathcal{S}$ must be turned into (p_i,α), but other formulas in $\mathcal{S}$ may remain as they were. Clearly the computational task in then heavier than with the brutal syntactic adjustment, which in terms requires only two satisfiability tests be carried out at most. The same remark applies to case (d) where $T(\mathcal{K},\beta)$ can be changed into another more parsimoniously revised $\mathcal{K}$ as mention in 5.3.

6 - Conclusion

Possibility theory can capture many forms of belief change encountered in the recent literature, while keeping the connection with probabilistic forms of belief change. Some results pertaining to the revision of syntactic belief bases completely ordered in terms of necessity degrees have been provided. They should be considered in the framework of more general tools

for merging belief bases, including symmetric combination methods for several belief bases issued from distinct sources, as recently studied by Benferhat et al. (1996).

Acknowledgements: This work has been partially supported by the European ESPRIT Basic Research Action No. 6156 entitled "Defeasible Reasoning and Uncertainty Management Systems (DRUMS-II)".

References

Alchourrón C.E.P., Gärdenfors P., Makinson D. (1985) On the logic of theory change: Partial meet functions for contraction and revision. J. of Symbolic Logic, 50, 510-530.

Benferhat S., Cayrol C., Dubois D., Lang J., Prade H. (1993) Inconsistency management and prioritized syntax-based entailment. Proc. of the 13th Inter. Joint Conf. on Artificial Intelligence (IJCAI'93), Chambéry, 640-645.

Benferhat S., Dubois D., Prade H. (1995) How to infer from inconsistent beliefs without revising?. Proc. of the 14th Inter. Joint Conf. on Artificial Intelligence (IJCAI'95), Montréal, Canada, Aug. 20-25, 1449-1455.

Benferhat S., Dubois D., Prade H. (1996) Syntactic and semantic approaches to information combination in possibilistic logic. In: Aggregation of Evidence under Fuzziness (B. Bouchon-Meunier, ed.), Physica-Verlag, Berlin, to appear.

Boutilier C. (1993) Revision sequences and nested conditionals. Proc. of the 13th Inter. Joint Conf. on Artificial Intelligence (IJCAI'93), Chambéry, 519-525.

Boutilier C., Goldszmidt M. (1993) Revision by conditional beliefs. Proc. of the 11th National Conf. on Artificial Intelligence (AAAI'93), Washington, 649-654.

Brewka G. (1989) Preferred subtheories: An extended logical framework for default reasoning. Proc. of the 11th Inter. Joint Conf. on Artificial Intelligence (IJCAI'89), Detroit, 1043-1048.

Cox R.T. (1946) Probability, frequency and reasonable expectation. Amer. J. of Physics, 14, 1-13.

Darwiche A., Pearl J. (1994) On the logic of iterated belief revision. Proc. of the 5th Conf. on Theoretical Aspects of Reasoning About Knowledge (TARK'94) (R. Fagin, ed.), Pacific Grove, CA, Morgan Kaufmann, 5-23.

Dubois D., Fodor J.C., Prade H., Roubens M. (1996) Aggregation of decomposable measures with application to utility theory. Theory and Decision, 41, 59-95.

Dubois D., Lang J., Prade H. (1992) Inconsistency in possibilistic knowledge bases: To live with it or not live with it. In: Fuzzy Logic for the Management of Uncertainty (L.A. Zadeh, J. Kacprzyk, eds.), Wiley, New York, 335-351

Dubois D., Lang J., Prade H. (1994) Possibilistic logic. In: Handbook of Logic in Artificial Intelligence and Logic Programming, Vol. 3 (D.M. Gabbay et al., eds.), Oxford University Press, 439-513.

Dubois D., Prade H. (1986) Possibilistic logic under matrix form. In: Fuzzy Logic in Knowledge Engineering (C.V. Negoita, H. Prade, eds.), Verlag TÜV Rheinland, Köln, 112-126.

Dubois D., Prade H. (1990) Aggregation of possibility masures. In: Multiperson Decision Making using Fuzzy Sets and Possibility Theory (J. Kacprzyk, M. Fedrizzi, eds.), Kluwer, Dordrecht, 55-63.

Dubois D., Prade H. (1991a) Epistemic entrenchment and possibilistic logic. Artificial Intelligence, 50, 223-239.

Dubois D., Prade H. (1991b) Updating with belief functions, ordinal conditional functions and possibility measures. In: Uncertainty in Artificial Intelligence, Vol. 6 (P. Bonissone et al. eds.), Elsevier Science Publ., New York, 311-329.

Dubois D., Prade H. (1992) Belief change and possibility theory. In: Belief Revision (P. Gärdenfors, ed.), Cambridge University Press, Cambridge, 142-182.

Dubois D., Prade H. (1993) Belief revision and updates in numerical formalisms —An overview, with new results for the possibilistic framework—. Proc. of the 13th Inter. Joint Conf. on Artificial Intelligence (IJCAI'93), Chambéry, 620-625.

Dubois D., Prade H. (1995) Possibility theory as a basis for qualitative decision theory. Proc. of the Inter. Joint Conf. on Artificial Intelligence (IJCAI'95), Montréal, 1924-1930.

Gärdenfors P. (1988) Knowledge in Flux: Modeling the Dynamics of Epistemic States. The MIT Press.

Gärdenfors P., Makinson D. (1994) Nonmonotonic inference based on expectation ordering. Artificial Intelligence, 65, 197-245.

Goldszmidt M., Pearl J. (1992) Rank-based systems: A simple approach to belief revision, belief update, and reasoning about evidence and actions. Proc. of the 3rd Inter. Conf. on Principles of Knowledge Representation and Reasoning (KR'92) (B. Nebel et al., eds.), Cambridge, 661-672.

Grove A. (1988) Two modellings for theory change. J. of Philosophical Logic, 15, 157-170.

Jeffrey R. (1965) The Logic of Decision. McGraw-Hill, New York.

Lehrer K., Wagner C.G. (1981) Rational Consensus in Science and Society. D. Reidel, Boston.

Nebel B. (1992) Syntax-based approaches to belief revision. In: Belief Revision (P. Gärdenfors, ed.), Cambridge University Press, Cambridge, UK, 52-88.

Pearl J. (1990) System Z: A natural ordering of defaults with tractable applications to default reasoning. Proc. of the 3rd Conf. on Theoretical Aspects of Reasoning about Knowledge (TARK'90) (R. Parikh, ed.), Morgan Kaufmann, 121-135.

Ramer A. (1989) Conditional possibility measures. Int. J. of Cybernetics and Systems, 20, 233-247.

Rescher N. (1976) Plausible Reasoning. Van Gorcum, Amsterdam

Rott H. (1991) A nonmonotonic conditional logic for belief revision I. In: The Logic of Theory Change (A. Furmann, M. Morreau, eds.), LNAI 465, Springer Verlag, Berlin, 135-183.

Spohn W. (1988) Ordinal conditional functions: A dynamic theory of epistemic states. In: Causation in Decision, Belief Change, and Statistics, Vol. 2 (W.L. Harper, B. Skyrms, eds.), D. Reidel, Dordrecht, 105-134.

Williams M.A. (1994a) On the logic of theory base change. Proc. of the Europ. Workshop on Logics in Artificial Intelligence (JELIA'95), LNCS 838, Springer Verlag, 86-105.

Williams M.A. (1994b) Transmutations of knowledge systems. Proc. of the 4th Inter. Conf. on Principles of Knowledge Representation and Reasoning (KR'94) (J. Doyle et al., eds.), Bonn, Morgan Kaufmann, 619-629.

Williams M.A. (1995) Iterated theory base change: A computational model. Proc. of the 14th Inter. Joint Conf. on Artificial Intelligence (IJCAI'95), Montréal, 1541-1547.

Williams M.A. (1996) Towards a practical approach to belief revision: Reason-based change. Proc. of the 5th Inter. Conf. on Principles of Knowledge Representation and Reasoning (KR'96), Cambridge, MA, Nov.

Williams P.M. (1980) Bayesian conditionalization and the principle of minimum information. British J. of the Philosophy of Science, 31, 131-144.

Zadeh L.A. (1978) Fuzzy sets as a basis for a theory of possibility. Fuzzy Sets and Systems, 1, 3-28.

Zadeh L.A. (1979) A theory of approximate reasoning. In: Machine Intelligence, Vol. 9 (J.E. Hayes, D. Michie, L.I. Mikulich, eds.), Elsevier, New York, 149-194.

Appendix: Mapping Adjustment of Kappa Functions to the Possibilistic Setting

An adjustment of κ by (A,n) is either a contraction κ^{-}_{A} if n = 0 or another belief change operation, defined by

$$\begin{aligned}
\kappa^{*}_{(A,n)} &= \kappa^{-}_{A} \text{ if } n = 0 \\
&= (\kappa^{-}_{A})^{x}_{(A,n)} \text{ if } 0 < n < \kappa(\bar{A}) \\
&= \kappa^{x}_{(A,n)} \text{ otherwise}
\end{aligned}$$

where

$$\begin{aligned}
\kappa^{-}_{A}(\omega) &= 0 \text{ if } \omega \in \bar{A} \text{ and } \kappa(\omega) = \kappa(\bar{A}) \\
&= \kappa(\omega) \text{ otherwise} \\
\kappa^{x}_{(A,n)}(\omega) &= 0 \text{ if } \omega \in A \text{ and } \kappa(\omega) = \kappa(A) \\
&= \kappa(A) \text{ if either } \omega \in A \text{ and } \kappa(\omega) \neq \kappa(A) \\
&\text{or } \omega \in \bar{A} \text{ and } \kappa(\omega) > n \\
&= n \text{ otherwise.}
\end{aligned}$$

It is obvious that $\kappa^{-}_{A}(\omega)$ becomes $\pi^{-}_{A}(\omega)$ using the mapping $\pi(\omega) = 2^{-\kappa(A)}$, and considering the definition of a contraction in the possibilistic setting. Hence Williams contraction exactly corresponds to possibilistic contraction. Let us consider $\kappa^{x}_{(A,n)}$. The corresponding possibilistic rule is (noticing $\alpha = 1 - 2^{-n}$)

$$\begin{aligned}
\pi^{x}_{(A,\alpha)}(\omega) &= 1 \text{ if } \omega \in A \text{ and } \pi(\omega) = \prod(A) \\
&= \pi(\omega) \text{ if } \omega \in A \text{ and } \pi(\omega) < \prod(A) \\
&= \pi(\omega) \text{ if } \omega \in \bar{A} \text{ and } \pi(\omega) < 1 - \alpha \\
&= 1 - \alpha \text{ if } \omega \in \bar{A} \text{ and } \pi(\omega) \geq 1 - \alpha.
\end{aligned}$$

This definition can then be simplified

$$\begin{aligned}
\pi^{x}_{(A,\alpha)}(\omega) &= \pi(\omega \mid A) \text{ if } \omega \in A \\
&= \min(1 - \alpha, \pi(\omega)) \text{ if } \omega \notin A.
\end{aligned}$$

This is the belief change operation (15) with $* = \min$: this operation has been given a clear meaning in Section 3.2 since (A,n) is thus an unsure input that can be rejected if too uncertain. The adjustment then can be expressed as follows in the possibilistic setting:

$$\begin{aligned}\pi^*_{(A,\alpha)} &= \pi^-_A \text{ if } \alpha = 0 \\ &= (\pi^-_A)^x_{(A,n)} && \text{if } 0 < \alpha < N(A) \\ &= \pi^x_{(A,n)} && \text{otherwise.}\end{aligned}$$

Assume $N(A) > \alpha > 0$. Then:

$$\begin{aligned}(\pi^-_A)^x_{(A,n)}(\omega) &= \pi^-_A(\omega \mid A) \text{ if } \omega \in A \\ &= \min(1 - \alpha, \pi^-_A(\omega)) \text{ if } \omega \notin A.\end{aligned}$$

It is not difficult to check that

i) $\pi^-_A(\omega \mid A) = \pi(\omega \mid A)$. Indeed if $\omega \notin A$ then both sides are zero. If $\omega \in A$ then $\pi^-_A(\omega) = \pi(\omega)$ everywhere;
ii) if $\omega \notin A$ then $\max(\pi(\omega), \pi(\omega \mid \bar{A})) = \pi(\omega \mid \bar{A})$.

Hence

$$\begin{aligned}\pi^*_{(A,\alpha)}(\omega) &= \pi(\omega \mid A) \text{ if } \omega \in A \\ &= \min(1 - \alpha, \pi(\omega \mid \bar{A})) \text{ otherwise.}\end{aligned}$$

Assume now $N(A) \leq \alpha$; then, as stated above, $\pi^*_{(A,\alpha)} = \pi^x_{(A,\alpha)}$. But since $N(A) \leq \alpha$, $\prod(A) \geq 1 - \alpha$, so when $\omega \notin A$:

$$\begin{aligned}&\text{if } \pi(\omega) < 1 - \alpha \text{ then } \pi(\omega \mid \bar{A}) = \pi(\omega); \\ &\text{if } \pi(\omega) = 1 - \alpha \text{ then } \pi(\omega \mid \bar{A}) \geq 1 - \alpha;\end{aligned}$$

hence $\min(1 - \alpha, \pi(\omega)) = \min(1 - \alpha, \pi(\omega \mid \bar{A}))$ for $\omega \notin A$. It becomes clear that $\pi^*_{(A,\alpha)}$ takes the same form whether $N(A) > \alpha$ or not and coincides with the counterpart of Jeffrey's rule in the possibilistic setting, that is, rule (11) with $* = \min$.