Argumentative inference in uncertain and inconsistent knowledge bases

Salem BENFERHAT – Didier DUBOIS – Henri PRADE

Institut de Recherche en Informatique de Toulouse (I.R.I.T.)
Université Paul Sabatier – C.N.R.S.
118 route de Narbonne – 31062 TOULOUSE – FRANCE
email: {Benferha, Dubois, Prade}@irit.fr

Abstract: This paper presents and discusses several methods for reasoning from inconsistent knowledge bases. A so-called argumentative-consequence relation, taking into account the existence of consistent arguments in favor of a conclusion and the absence of consistent arguments in favor of its contrary, is particularly investigated. Flat knowledge bases, i.e. without any priority between their elements, as well as prioritized ones where some elements are considered as more strongly entrenched than others are studied under the different consequence relations which are considered. Lastly a paraconsistent-like treatment of prioritized knowledge bases is proposed, where both the level of entrenchment and the level of paraconsistency attached to a formula are propagated. The priority levels are handled in the framework of possibility theory.

Keywords: Inconsistency; consequence relation; prioritized knowledge base; uncertainty; possibilistic logic; possibility theory.

Submitted to the Ninth Annual Conference on Uncertainty in Artificial Intelligence
Washington, D.C., July 9-11, 1993

1. Introduction

One of the emerging important problems pertaining to the management of knowledge-based systems is inconsistency handling. Inconsistency may be present for several reasons: the presence of general rules with exceptions, the existence of several possibly disagreeing sources feeding the knowledge base are among the most common ones. There are two attitudes in front of inconsistent knowledge. One is to revise the knowledge base and restore consistency. The other is to cope with inconsistency. The first approach meets two difficulties: there are several ways of restoring inconsistency yielding different results, and the problem is that part of the information is thrown away and we no longer have access to it. Coping with inconsistency bypass these difficulties. However we must take a step beyond classical logic, since the presence of inconsistency enables anything to be entailed from a set of formulas.

This paper investigates several methods for coping with inconsistency by suitably defining notions of consequence capable of inferring non trivial conclusions from an inconsistent knowledge base. These consequence relationships coincide with the classical definition when the knowledge base is inconsistent. When the knowledge base is flat, i.e. made of equally reliable propositional formulas, the proposal made by Resher and Manor (1970) is very commonly used nowadays: compute the set of maximal consistent subsets of the knowledge base first, then a formula is accepted as a consequence when it can be classically inferred from all maximal consistent subsets of propositions (this is the universal consequence) or from at least one maximal consistent subset (this is the existential consequence).

However the first consequence relation is very conservative hence rather unproductive while the latter is too permissive and leads to conclusions which are themselves inconsistent with each other. A mild inference approach is proposed in this paper, that is more productive than the universal consequence but do not lead to outright contradictions. It is based on the idea of arguments that goes back to Toulmin (1956), and is related to previous proposals by Poole (1988), Pollock (1987), and Simari and Loui (1992) that were suggested in the framework of defeasible reasoning for handling exceptions. We suggest that a conclusion can be inferred from an inconsistent knowledge base if the latter contains an argument that supports this conclusion, but no argument that supports its negation.

The paper is organised as follows. Section 2 deals with flat knowledge bases and compares several notions of consequence relations that are inconsistency tolerant, including several ones that comes from the non-monotonic logic literature. Section 3 contains a thorough analysis of our argumentative inference process. Section 4 extends the argumentative inference to layered knowledge bases where layers express degrees of certainty as in possibilistic logic (Dubois et al., 1993); it refines the flat case by allowing for pieces of information of various levels. Section 5 deals with a paraconsistent-like treatment of layered inconsistent knowledge bases, whereby a formula carries two weights: its degree of certainty and the degree of certainty of its negation. Lastly a new way of combining knowledge bases issued from several sources is suggested, inspired by the argumentative inference.

2. Arguments in Flat Knowledge bases

2.1. Definition of an Argumentative Consequence Relation

For the sake of simplicity, we consider in this paper only a finite propositional language denoted by $\mathcal{L}$. We denote the set of classical interpretations by Ω, by $\vdash$ the classical consequence relation, and by Cn(S) the deductive closure of S, i.e. $Cn(S) = \{\phi \in \mathcal{L}, S \vdash \phi\}$. Let Σ be a set of propositional formulas, possibly inconsistent but not closed under Cn. We also assume that the knowledge bases manipulated in this section are flat, which means that all formulas in Σ have the same reliability.

Definition: A sub-base Σ_i of Σ is said to be consistent if it is not possible to deduce a contradiction from Σ_i, and is said to to be maximally consistent if adding any formula ϕ from Σ-Σ_i to Σ_i produces the inconsistency of $\Sigma_i \cup \{\phi\}$.

We introduce now the notion of argument:

Definition: A sub-base Σ_i of Σ is said to be an argument for a formula ϕ, if it satisfies the following conditions:

(i) $\Sigma_i \nvdash \bot$ (Σ_i is consistent),

(ii) $\Sigma_i \vdash \phi$, and

(iii) $\forall\psi\in \Sigma_i$, $\Sigma_i - \{\psi\} \nvdash \phi$

Notice that this notion of argument is identical to the one proposed by Simari and Loui (1992) and is also very similar to the notion of environment used in the terminology of the ATMS (De Kleer, 1986).

Definition: A formula ϕ is said to be an argumentative consequence of Σ, denoted by $\Sigma \vdash_{\mathcal{A}} \phi$, if and only if:

(i) there exists an argument for ϕ in Σ, and

(ii) there is no argument for $\neg\phi$ in Σ.

As a consequence of this definition, if our knowledge base contains only the two contradictory statements $\{\phi,\neg\phi\}$ then the inference $\phi \wedge \neg\phi \vdash_{\mathcal{A}} \psi$ does not hold. In other words, our approach is in agreement with the idea of paraconsistent logics (e.g. Da Costa, 1963), where they reject the principle "ex absurdo quodlibet" which allows the deduction of any formula from an inconsistent base.

The next propositions give some properties of the argumentative inference $\vdash_{\mathcal{A}}$:

Proposition 1: $\vdash_{\mathcal{A}}$ is non-monotonic

Proof:

Indeed, let us consider the following example where our knowledge base Σ contains only the formula ϕ. It is obvious that the formula ϕ is an argumentative consequence of

Σ. Let us add to Σ the information that ϕ is false, then ϕ will not be an argumentative consequence of $\Sigma' = \{\phi, \neg\phi\}$

Proposition 2: if Σ is consistent, then $\Sigma \vdash \phi$ iff $\Sigma \vdash_{\mathcal{A}} \phi$

Proof:

If a formula ϕ is a logical consequence of Σ, then there obviously exists in Σ an argument for ϕ. Since Σ is consistent, then $\neg\phi$ cannot be deduced from Σ, which means that there is no argument for $\neg\phi$ in Σ, and therefore by definition ϕ is also an argumentative consequence of Σ. The second part of the proof goes in a similar way.

Proposition 2 means that the non-monotonicity only appears in the presence of inconsistency, and the argumentative consequence resorts to what Satoh (1990) calls "lazy non-monotonic reasoning", an idea also proposed by Lin (1987).

2.2. Comparative Study of Inconsistency-Tolerant Consequence Relations

In this sub-section we compare our approach to reasoning in the presence of inconsistency to the ones reviewed by Benferhat et al. (1993). We start this comparative study by presenting the different approaches from the most conservative ones to the more adventurous ones. But first we need some further definitions:

Definition: A sub-base Σ_i of Σ is said to be minimal inconsistent if and only if it satisfies the two following requirements:

- $\Sigma_i \vdash \bot$, and
- $\forall \phi \in \Sigma_i, \Sigma_i - \{\phi\} \nvdash \bot$.

From now on, we denote by Inc(Σ) the set of propositions belonging to at least one minimal inconsistent sub-base of Σ, namely:

$$\mathrm{Inc}(\Sigma) = \{\phi, \exists\, \Sigma_i \subseteq \Sigma, \text{ such that } \phi \in \Sigma_i \text{ and } \Sigma_i \text{ is minimal inconsistent}\}$$

The set Inc(Σ) is somewhat related to the "base of nogoods" used in the terminology of the ATMS (De Kleer, 1986). Once Inc(Σ) is computed, we remove from Σ all elements of Inc(Σ), the result base is called the free base of Σ, denoted by Free(Σ) (Benferhat et al., 1992). In other words, the set Free(Σ) contains all formulae which are not involved in any inconsistency of the knowledge base Σ. Now, let us introduce the notion of the Free consequence, denoted by $\vdash_{\mathrm{Free}}$:

Definition: Let Σ be a possibly inconsistent knowledge base Σ, and let Free(Σ) be its free base. A formula ϕ is said to be a free consequence (or a sound consequence) of Σ, denoted by $\Sigma \vdash_{\mathrm{Free}} \phi$, if and only if ϕ is logically entailed from Free(Σ), namely:

$$\Sigma \vdash_{Free} \phi \quad \text{iff} \quad Free(\Sigma) \vdash \phi$$

The Free-inference relation is very conservative as we will see later. Let us now recall the approach first proposed by Resher and Manor (1970). Let Σ be a possibly inconsistent base, $MC(\Sigma)$ the set of all maximal consistent sub-bases of Σ, Resher and Manor define the universal (called also the inevitable) consequence relation in this way:

Definition: A formula ϕ is said to be an universal consequence or MC-consequence of Σ, denoted by $\Sigma \vdash_\forall \phi$, if and only if ϕ is entailed from each element of $MC(\Sigma)$, namely:

$$\Sigma \vdash_\forall \phi \quad \text{iff} \quad \forall\, \Sigma_i \in MC(\Sigma),\ \Sigma_i \vdash \phi$$

As it has been mentioned above, the Free consequence relation is more conservative then MC-consequence:

Proposition 3: Each Free-consequence is also a MC-consequence. The converse is false.

Proof:

(i) Let us consider our knowledge base Σ as a pair $(Inc(\Sigma), Free(\Sigma))$, and let $\Sigma_1, \ldots, \Sigma_n$ be the maximal consistent sub-bases of $Inc(\Sigma)$. It is obvious that $\Sigma_1 \cup Free(\Sigma), \ldots, \Sigma_n \cup Free(\Sigma)$ are the maximal consistent sub-bases of Σ (since $Free(\Sigma)$ are outside any conflict). Then each element of $MC(\Sigma)$ contains $Free(\Sigma)$, therefore if a formula is a Free-consequence then it is also a MC-consequence.

(ii) To show that the converse is false, let us consider the following counterexample where our base contains the five formulae:

$$\Sigma = \{A, \neg A \vee \neg B, B, \neg A \vee E, \neg B \vee E\}$$

The base Σ is inconsistent, and the inconsistency is caused by the three first formulas, which means that the free base of Σ is $Free(\Sigma)= \{\neg A \vee E, \neg B \vee E\}$. It is clear that E cannot be entailed from $Free(\Sigma)$.

In contrast with the MC-consequence, the base contains three maximal sub-bases: $\Sigma_1 = \{\neg A \vee \neg B, B, \neg A \vee E, \neg B \vee E\}$, $\Sigma_2 = \{A, B, \neg A \vee E, \neg B \vee E\}$ and $\Sigma_3 = \{A, \neg A \vee \neg B, \neg A \vee E, \neg B \vee E\}$ corresponding to the case where we remove from Σ each element of $Inc(\Sigma)$. We see that each sub-base entails E, therefore E is a MC-consequence. In the above example, it is clear that the MC-consequence involves an idea of parsimony with respect to the removal of inconsistency; each maximal consistent sub-base is obtained by removing the least number of formulas sufficient to restore consistency. This is not so when considering $Free(\Sigma)$.

There is another way to find the proof of the previous proposition noticing that:

$$Free(\Sigma) = \bigcap_{\Sigma_i \in MC(\Sigma)} \Sigma_i$$

since if a formula ϕ does not belong to $Free(\Sigma)$ then there exists a minimal inconsistent sub-base Σ_k containing ϕ, and therefore there exists at least one maximally consistent sub-base which contains Σ_k but not ϕ, which means that there exists at least one element of $MC(\Sigma)$

which does not contain ϕ, and consequently ϕ does not belong to the intersection of the elements of $MC(\Sigma)$. Then from the properties of Cn, we find:

$$Cn(Free(\Sigma))=Cn(\bigcap_{\Sigma_i \in MC(\Sigma)} \Sigma_i) \subseteq \bigcap_{\Sigma_i \in MC(\Sigma)} Cn(\Sigma_i).$$

The next proposition compares the MC-consequence to the argumentative consequence:

Proposition 4: A formula ϕ is an argumentative consequence of Σ iff
– $\exists\ \Sigma_i \in MC(\Sigma)$, such that $\Sigma_i \vdash \phi$, and
– $\nexists \Sigma_j \in MC(\Sigma)$, such that $\Sigma_j \vdash \neg\phi$.

Proof (Indications)

The idea behind the proof is to remark that if there exists an argument Σ_i for a formula ϕ in the knowledge base Σ, then we can build a maximal sub-base from Σ_i by adding as much as possible (wrt to consistency criterion) formulas of Σ-Σ_i. The converse is true, if a maximally consistent sub-base entails some formula ϕ, then we can extract an argument from Σ_i (and therefore from Σ) which entails ϕ.

Proposition 5: Each MC-consequence of Σ is also an argumentative consequence of Σ. The converse is false

Proof:

-- If ϕ is MC-consequence of Σ, then each element of $MC(\Sigma)$ enables us to infer ϕ. By consequence, i.e. $\exists\ \Sigma_i \in MC(\Sigma)$, $\Sigma_i \vdash \phi$. As, each element of $MC(\Sigma)$ is consistent and entails ϕ, then it does not exist an element of $MC(\Sigma)$ which enables us to deduce $\neg\phi$. In other words, $\nexists \Sigma_i \in MC(\Sigma)$, $\Sigma_i \vdash \neg\phi$.

Then, using the previous proposition, we conclude that ϕ is an argumentative consequence of Σ.

-- Let $\Sigma = \{A, \neg A, A \rightarrow B\}$. We have $MC(\Sigma) = \{\{A\ \ A \rightarrow B\}, \{\neg A, A \rightarrow B\}\}$. In this example, we have B which is an argumentative consequence of Σ, while it is not a MC-consequence.

One of the main drawbacks of MC-consequence is the number of $MC(\Sigma)$ which increases exponentially with the size[1] of the base and in general, it is not possible to take into account all the elements of $MC(\Sigma)$. Dubois et al. (1992) have proposed to select a non-empty subset of $MC(\Sigma)$, denoted by $Lex(\Sigma)$, and computed in the following manner:

$$\Sigma_i \in Lex(\Sigma) \text{ iff } \forall\ \Sigma_j \in MC(\Sigma),\ |\Sigma_i| \geq |\Sigma_j|$$

where $|\Sigma|$ is the cardinality of Σ[2]. A probabilistic justification of $Lex(\Sigma)$ can be found in Benferhat et al. (1993).

1 In fact the cardinality of $MC(\Sigma)$ increases exponentially with the number of the conflicts in Σ.

2 The ordering defined on elements of $MC(\Sigma)$ is called the lexicographical ordering, and corresponds to the property of parsimony advocated in diagnostic problems (Reggia et al., 1985).

In order to generate the set of plausible inferences based on Lex(Σ) from an inconsistent knowledge base, we use a definition similar to the MC-consequence:

Definition: A formula ϕ is said to be a Lex-consequence of Σ, denoted by $\Sigma \vdash_{Lex} \phi$, if and only if it is entailed from each element of Lex(Σ), namely:

$$\Sigma \vdash_{Lex} \phi \quad \text{iff} \quad \forall \Sigma_i \in Lex(\Sigma), \Sigma_i \vdash \phi$$

Proposition 6: Each MC-consequence of Σ is also a Lex-consequence of Σ. The converse is false.

This is obvious since the Lex-consequence uses a subset of MC(Σ).

The Lex-consequence and argumentative consequence are not comparable as we see in the following example:

<u>Example</u>

Let $\Sigma = \{A, \neg B \vee \neg A, B, \neg C \vee \neg A, C, \neg A \vee D\}$

We have $Lex(\Sigma) = \{\{\neg B \vee \neg A, B, \neg C \vee \neg A, C, \neg A \vee D\}\}$. Then ¬A is a Lex-consequence of Σ while it is not an argumentative consequence, since A is also present in Σ.

In contrast, D is an argumentative consequence (it derives from $\{A, \neg A \vee D\}$ while it is not a Lex-consequence.

The Lex-consequence may appear as an arbitrary selection from MC(Σ) if we consider a semantic point of view. Namely, the following situation may happen: $\Sigma_i \in Lex(\Sigma)$, $\Sigma_j \in MC(\Sigma) - Lex(\Sigma)$ and one may define Σ_k logically equivalent to Σ_j but $|\Sigma_k| > |\Sigma_i|$. However all introduced consequence relations are syntax-sensitive since Σ is not closed. Yet, the counterexample demonstrates that the Lex-consequence may implicitly delete some useful pieces of knowledge (here A). It may result in destroying some arguments, as well as some rebuttals (i.e. formulas whose presence ensure an argument for $\neg\phi$ that inhibits arguments for ϕ).

Resher and Manor (1970) have also proposed another definition of the consequence relation, called existential relation that can be described in the following way:

Definition: Let MC(Σ) be the set of all maximally consistent sub-bases of Σ. A formula ϕ is said to be an existential consequence of Σ, denoted by $\Sigma \vdash_{\exists} \phi$, if and only if there exists at least one element of MC(Σ) which entails ϕ, namely:

$$\Sigma \vdash_{\exists} \phi \quad \text{iff} \quad \exists \Sigma_i \in MC(\Sigma), \ \Sigma_i \vdash_{\exists} \phi$$

It is not hard to see that this approach is the more adventurous one, but unfortunately it has an important drawback, since this approach may lead to inconsistent set of results. Indeed, there may exists $\Sigma_i \vdash \phi$ and $\Sigma_j \vdash \neg\phi$, in which case both ϕ and $\neg\phi$ will be inhibited.

The following hierarchy summarizes the links existing between the different consequence relations studied here, the edge means the inclusion-set relation between the set of results generated by each inference relation. The top of the diagram thus corresponds to the most conservative inferences. All inferences reduce to the classical one when Σ is consistent.

Figure 1: A comparative study of inference relations

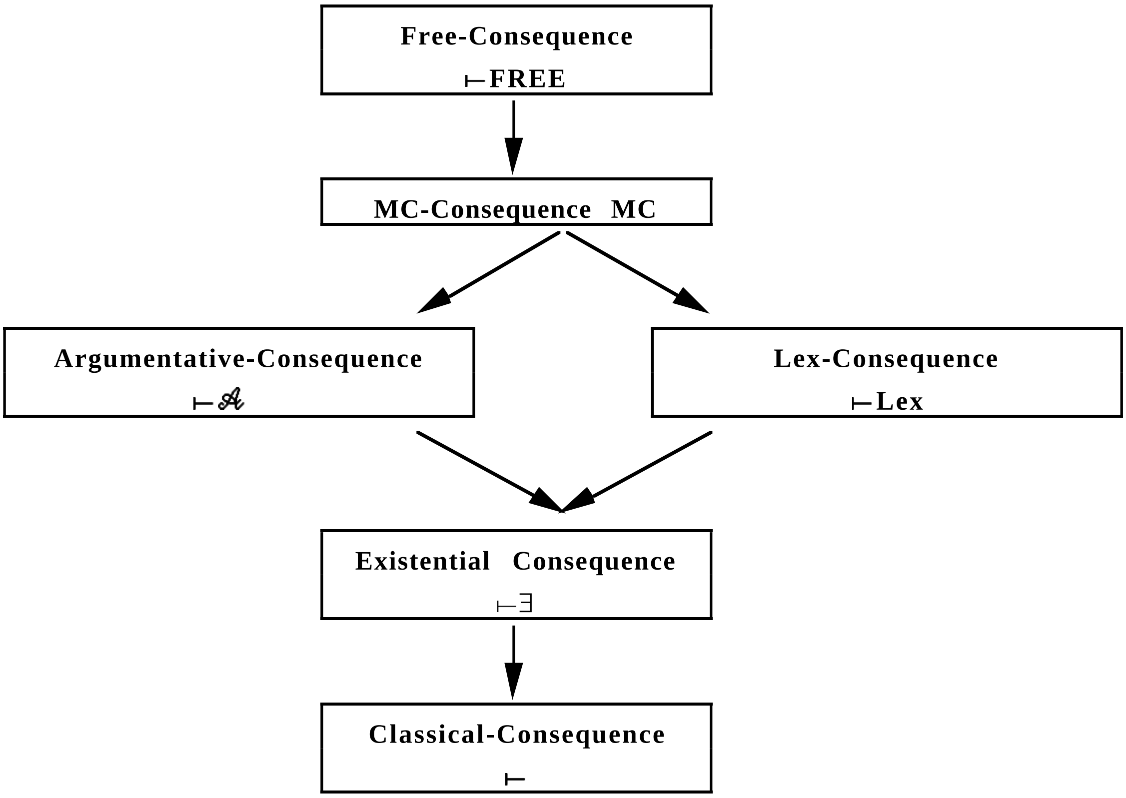

3. Properties of $\vdash_{\mathcal{A}}$:

Proposition 7 (failure of AND): We may have $\Sigma \vdash_{\mathcal{A}} \phi$, $\Sigma \vdash_{\mathcal{A}} \psi$, and *not* $\Sigma \vdash_{\mathcal{A}} \phi \wedge \psi$.

Proof

> Indeed, let us consider the following counterexample where our knowledge base is $\Sigma = \{\neg\text{bird} \vee \text{fly}, \text{bird} \vee \text{swim}, \text{bird}, \neg\text{bird}\}$. It is clear that fly and swim are both argumentative consequences of Σ, while there is no argument which supports fly $\wedge$ swim.

Proposition 7 must not be seen as a major drawback of $\vdash_{\mathcal{A}}$ since in some cases we do not want to have the AND property (as in the previous example). The $\vdash_{\mathcal{A}}$ consequence relation captures the cases when we believe in two mutually consistent properties of some objects for conflicting reasons.

Proposition 8: $\vdash_{\mathcal{A}}$ is closed under the classical consequence relation, i.e.

$$\text{If } \phi \vdash \psi \text{ then } \Sigma \vdash_{\mathcal{A}} \phi \text{ implies } \Sigma \vdash_{\mathcal{A}} \psi$$

Proof

Indeed, $\Sigma \vdash_{\mathscr{A}} \phi$ means that there exists an argument for ϕ in Σ, since $\phi \vdash \psi$. We conclude that there is also an argument for ψ in Σ. Assume now that there exists also an argument for $\neg\psi$ in Σ, then since $\phi \vdash \psi$ we conclude that there is also an argument for $\neg\phi$ in Σ which contradicts the fact that $\Sigma \vdash_{\mathscr{A}} \phi$.

An important issue when reasoning with an inconsistent knowledge base Σ is whether it is possible to construct some equivalent consistent base $Cn_{\mathscr{A}}(\Sigma)$ such that plausible inferences from Σ are the logical consequences of $Cn_{\mathscr{A}}(\Sigma)$. In this section we try to construct $Cn_{\mathscr{A}}(\Sigma)$ using the argumentative inference relation.

Propositions 7 and 8 are very important to characterise the set of argumentative consequences of a possibly inconsistent base Σ, denoted by $Cn_{\mathscr{A}}(\Sigma)$, i.e.

$$Cn_{\mathscr{A}}(\Sigma) = \{\phi, \Sigma \vdash_{\mathscr{A}} \phi\}.$$

The fact that the argumentative consequence is not closed under conjunction means that $Cn_{\mathscr{A}}(\Sigma)$ is generally not equal to its closure under Cn , namely:

$$Cn_{\mathscr{A}}(\Sigma) \neq Cn_{\mathscr{A}}(Cn_{\mathscr{A}}(\Sigma))$$

Definition: A formula ϕ is said to be a prime implicate of Σ with respect to the argumentative inference relation if and only if:

(i) $\Sigma \vdash_{\mathscr{A}} \phi$

(ii) $\nexists\phi'$, such that $\phi' \vdash \phi$ and $\Sigma \vdash_{\mathscr{A}} \phi'$

If ϕ is a prime implicate then it can be inferred from a maximal consistent subset of Σ. However, if $\Sigma_i \in MC(\Sigma)$, then the conjunction of formulas in Σ_i is not a prime implicate since it can be defeated by other maximal consistent subsets of Σ. Indeed, $\forall$ i $\neq$ j, $\Sigma_i \vdash \neg\Sigma_j$. Moreover there is at most one prime implicate inferred by a given maximal consistent subset Σ_i. Indeed, if $\Sigma \vdash_{\mathscr{A}} \phi$, $\Sigma \vdash_{\mathscr{A}} \psi$, and $\Sigma_i \vdash \phi$, $\Sigma_i \vdash \psi$ then $\Sigma \vdash_{\mathscr{A}} \phi \wedge \psi$ and neither ϕ nor ψ are prime implicates. Hence the number of prime implicates of Σ is at most the number of maximal consistent subsets. It can be strictly less: $\Sigma=\{A, \neg A\}$ has only one prime implicate, the tautology.

Let $R_1, \ldots, R_n$ be the set of prime implicates of Σ, then $Cn_{\mathscr{A}}(\Sigma)$ can be seen as the union of the deductive closure of each R_i under Cn, namely:

$$Cn_{\mathscr{A}}(\Sigma) = Cn(R_1) \cup \ldots \cup Cn(R_n) \text{ such that:}$$

And it is easy to check that $\forall$ i, j = 1,n $\Sigma \nvdash_{\mathscr{A}} R_i \wedge R_j$

Examples

(1) let Σ be the following base:

$$\Sigma = \{\neg A \vee B, A \vee C, A, \neg A\}.$$

Then:

$$Cn_{\mathscr{A}}(\Sigma) = Cn(\{C\}) \cup Cn(\{B\})$$

(2) Consider now Σ as:

$$\Sigma = \{\neg A \vee B, A \vee B, A, \neg A, C\}.$$

Then:

$$Cn_{\mathscr{A}}(\Sigma) = Cn(\{B \wedge C\})$$

Let us consider this last property. The fact that $\Sigma \vdash_{\mathscr{A}} R_i$, $\Sigma \vdash_{\mathscr{A}} R_j$ and $\Sigma \nvdash_{\mathscr{A}} R_i \wedge R_j$ can correspond to two cases:

(i) • No argument supporting $R_i \wedge R_j$ can be found. In that case the arguments Σ_i and Σ_j supporting R_i and R_j respectively are inconsistent. Indeed, if $\Sigma_i \cup \Sigma_j$ is consistent then $\Sigma_i \cup \Sigma_j \vdash_{\mathscr{A}} R_i \wedge R_j$ and there would exist an argument supporting $R_i \wedge R_j$.

(ii) • There is an argument for $\neg R_i \vee \neg R_j$. In that case it is not clear that if Σ_i and Σ_j are arguments for R_i and R_j respectively, if $\Sigma_i \cup \Sigma_j$ is inconsistent. Indeed the fact that $\Sigma_i \cup \Sigma_j$ is consistent, hence supports $R_i \wedge R_j$ can no longer used to contradict the assumption that R_i and R_j are prime implicates. Proving that $\Sigma_i \cup \Sigma_j$ should be inconsistent even in that case remains an open problem.

Anyway the arguments supporting the prime implicates can be viewed as a set of scenarii extracted from Σ, that express different points of views on what is the actual information contained in Σ. These points of view are incompatible in the sense that the subsets Σ_i and Σ_j supporting two prime implicates R_i and R_j should not be mixed (even if not inconsistent). The fact that $Cn_{\mathscr{A}}(\Sigma)$ still reflects conflicts lying in Σ can be seen as follows: the argumentative inference forbids that two prime implicates R_i and R_j be inconsistent. However the set $\{R_1, \ldots, R_n\}$ can be globally inconsistent, namely one argumentative consequence of Σ can be defeated by other consequences grouped together.

Example

Consider the set $\Sigma = \{\neg A, \neg B, A, B, \neg C \vee \neg D, \neg A \vee B\}$

The maximal consistent subsets of Σ are:

$$\Sigma_1 = \{\neg A, \neg B, \neg C \vee \neg D, \neg A \vee B\}$$
$$\Sigma_2 = \{\neg A, B, \neg C \vee \neg D, \neg A \vee B\}$$
$$\Sigma_3 = \{A, B, \neg C \vee \neg D, \neg A \vee B\}$$
$$\Sigma_4 = \{\neg B, A, \neg C \vee \neg D\}$$

Consider the three formulas:

$$\phi_1 = (\neg A \wedge \neg B \wedge (\neg C \vee \neg D)) \vee (\neg C \wedge D \wedge (A \vee B))$$
$$\phi_2 = (\neg A \wedge B \wedge (\neg C \vee \neg D)) \vee (C \wedge \neg D \wedge (A \vee \neg B))$$

$$\phi_3 = (A \wedge B \wedge (\neg C \vee \neg D)) \vee (\neg C \wedge \neg D \wedge (\neg A \vee \neg B))$$

It is easy to see that $\Sigma_1 \vdash \phi_1$, $\Sigma_2 \vdash \phi_2$ and $\Sigma_3 \vdash \phi_3$, but we never have $\Sigma_i \vdash \phi_j$ for $i \neq j$.

Moreover $\phi_1 \wedge \phi_2 \wedge \phi_3 \vdash \bot$.

This result can be viewed as a weakness of the argumentative inference which avoids obvious direct contradictions, but does not escape hidden ones. It confirms the fact that $Cn_{\mathscr{A}}(\Sigma)$ is a heterogeneous set of properties that pertain to distinct views of the world. This means that a question-answering system whereby a question "is it true that ϕ" is answered by yes or no after computing $\Sigma \vdash_{\mathscr{A}} \phi$ is not really informative enough. The system must also supply the argument for ϕ. This way of coping with inconsistency looks natural, and the arguments for ϕ and ψ should enable the user to decide whether these two plausible conclusions can be accepted together or not.

4. Arguments in prioritized knowledge bases

The use of priorities among formulas is very important to appropriately revise inconsistent knowledge bases. For instance Gärdenfors (1988) has proved that any revision process that satisfies natural requirements is implicitly based on such a set of priorities. Similarly a proper treatment of default rules also leads to prescribe priority levels (e.g. Pearl, 1990). In these two cases, the handling of priorities has been shown to be completely in agreement with possibilistic logic (Dubois, Lang and Prade, 1992b; Benferhat et al., 1992). Fox et al. (1992) also manipulate arguments of different levels in a way completely consistent with possibilistic logic.

In the prioritized case, a knowledge base can be viewed as a layered knowledge base $\Sigma = B_1 \cup \ldots \cup B_n$, such that formulas in B_i have the same level of priority or certainty and are more reliable than the ones in B_j where $j > i$. This stratification is modelled by attaching a weight $\alpha \in [0,1]$ to each formula with the convention that $(\phi\ \alpha_i) \in B_i$, $\forall i$ and $\alpha_1 = 1 > \alpha_2 > \ldots > \alpha_n > 0$.

A sub-base $\Sigma_i = E_1 \cup \ldots \cup E_n$ of $\Sigma = B_1 \cup \ldots \cup B_n$ where $\forall\ j = 1,n$, $E_j \subseteq B_j$ is said to be consistent if: $\Sigma_i \nvdash \bot$ and is said to be maximal consistent if adding any formula from $(\Sigma - \Sigma_i)$ to Σ_i produces an inconsistent knowledge base.

Before introducing the notion of argument in prioritized knowledge base, let us define the notion of entailment in a layered base, named π-entailment:

Definition: Let $\Sigma = B_1 \cup \ldots \cup B_n$ be a layered knowledge base. A formula ϕ is said to be a π-consequence of Σ with weight α_i, denoted by $\Sigma \vdash_\pi (\phi\ \alpha_i)$[1], if and only if:

(i) $B_1 \cup \ldots \cup B_i$ is consistent, and
(ii) $B_1 \cup \ldots \cup B_i \vdash \phi$
(iii) $\forall\, j < i,\ B_1 \cup \ldots \cup B_j \nvdash \phi$

It is clear that in the presence of inconsistency the π-entailment and the classical entailment have not the same behaviour. Indeed in classical logic if our base Σ is inconsistent then any formula can be deduced from Σ and the base becomes useless. In a stratified base, the situation is better since it is possible to use only a consistent subbase of Σ (in general not maximal), denoted by $\pi(\Sigma)$, induced by the levels of priority and defined in this way:

$$\pi(\Sigma) = B_1 \cup \ldots \cup B_i, \text{ such that } \pi(\Sigma) \text{ is consistent and } B_1 \cup \ldots \cup B_{i+1} \text{ is inconsistent}$$

The remaining sub-base $\Sigma - \pi(\Sigma)$ is simply inhibited. It is not hard to check that the following result holds:

$$\Sigma \vdash_\pi \phi \quad \text{iff} \quad \pi(\Sigma) \vdash \phi$$

However, this way of dealing with inconsistency is not entirely satisfactory, since it suffers from a principal drawback named "drowning problem" in (Benferhat et al., 1993), as we can see in the following examples:

<u>Examples:</u>

• Let Σ be the following stratified knowledge base:

$$\Sigma = \{\{\neg A \vee \neg B\}, \{A\}, \{B\}, \{C\}\}$$

This notation of the form $\{B_1, B_2, \ldots, B_n\}$, where the weights are omitted is used for the sake of simplicity. This base is of course inconsistent, and only the subset $\Sigma_i = \{\{\neg A \vee \neg B\}, \{A\}\}$ is valid, and therefore C can not be deduced despite the fact that C is outside the conflict.

• A particular case of the drowning effect is called "blocking property inheritance" (Pearl, 1990; Benferhat et al., 1992). This is can be illustrated from the following set of stratified defaults:

$$\Sigma = \{\{p\}, \{\neg p \vee b, \neg p \vee \neg f\}, \{\neg b \vee f, \neg b \vee w\}\}$$

where p, b, f and w means respectively penguin, bird, fly and wings. From this base it is not possible for a penguin to inherit properties of birds (in our example to inherit property of having wings), while the only undesirable property for a penguin is "flying".

One way of solving the drowning problem is to recover the inhibited free defaults, denoted by IFree(Σ), and defined in this way:

---

[1] The definition of $\vdash_\pi$ is identical to the one proposed in possibilistic logic (Dubois, Lang and Prade, 1989, 1991, 1992b).

$$IFree(\Sigma) = Free(\Sigma) \cap (\Sigma - \pi(\Sigma))$$

Then once the inhibited free set has been computed, we define the new inference relation in this way:

Definition: A formula ϕ is said to be a π+Free-consequence of Σ, if and only if it is entailed logically from $\pi(\Sigma)$ and IFree(Σ), namely:

$$\Sigma \vdash_{\pi+free} \phi \quad \text{iff} \qquad \pi(\Sigma) \cup IFree(\Sigma) \vdash \phi$$

Proposition: Each π-consequence of Σ is also a π+Free-consequence of Σ.

This is obvious since the π-consequence uses a subset of the base used by $\vdash_{\pi+Free}$.

Brewka (1989) (see also Roos (1992)) has proposed a more adventurous approach to reason with inconsistent and layered knowledge bases, the idea is to take advantage of the stratification of the base to rank-order the maximal consistent sub-bases of Σ and keep only the best ones, namely the "so-called preferred sub-bases".

Let $\Sigma = B_1 \cup \ldots \cup B_n$ be a layered knowledge base. A preferred sub-base Σ_i is constructed by starting with a maximal consistent sub-base of B_1, then we add to Σ_i as many formulas of B_2 as possible (wrt to consistency criterion), and so on. Formally, Σ_i is a preferred sub-base of Σ if it can be constructed as follows:

$$\Sigma_i = E_1 \cup E_2 \cup \ldots \cup E_n$$

where $\forall$ j = 1,n, $E_1 \cup E_2 \cup \ldots \cup E_j$ is a maximal consistent sub-base of $B_1 \cup B_2 \cup \ldots \cup B_j$.

Preferred subbases have also been independently introduced by (Dubois, Lang and Prade, 1992a) in the setting of possibilistic logic under the name of strongly maximal consistent subbases. They are such that $\Sigma_i \cup \{(\phi\ \alpha)\} \vdash_\pi (\bot\ \alpha)$, $\forall (\phi\ \alpha) \in \Sigma$.

Definition: Let Pref(Σ) be the set of preferred sub-bases of Σ. A formula ϕ is said to be a preferred consequence of Σ, denoted by $\Sigma \vdash_{pref} \phi$, if and only if it is entailed from each element of Pref(Σ), namely:

$$\Sigma \vdash_{pref} \phi \qquad \text{iff} \qquad \forall \Sigma_i \in Pref(\Sigma),\ \Sigma_i \vdash \phi$$

Proposition: Each π+Free-consequence of Σ is also a preferred consequence of Σ. The converse is false

Proof:

• Let us prove that each preferred sub-base of Σ contains $\pi(\Sigma) \cup IFree(\Sigma)$. Let $\Sigma = B_1 \cup \ldots \cup B_n$. First it is easy to check that Ifree(Σ) is included in each sub-base since it is outside any conflict. We have also $\pi(\Sigma)$ included in any preferred sub-base. Indeed, let E a preferred sub-base, then we have by recurrence:

- if B_1 is inconsistent then $\pi(\Sigma) = \emptyset$ then $\pi(\Sigma) \subseteq E$,
- Now assume that $B_1 \cup \ldots \cup B_i$ is consistent and $B_1 \cup \ldots \cup B_{i+1}$ is inconsistent, then it is clear that $B_1 \cup \ldots \cup B_i \subseteq E$ and therefore $\pi(\Sigma) \subseteq E$.

• To see that the converse is false, let us consider the following counterexample:

$$\Sigma = \{\{A, B, \neg A \vee \neg B\}, \{\neg A \vee C, \neg B \vee C\}\}$$

It is easy to see that C is a preferred consequence of Σ, while it is not a π+Free-consequence of Σ.

The Lex-consequence relation described in the case of flat knowledge bases has also been proposed in the case of stratified knowledge bases (Dubois et al., 1992). The objective is to reduce the number of elements of Pref(Σ), by selecting the elements which satisfy the following requirement:

$$\Sigma_i = E_1 \cup \ldots \cup E_n \in \text{Lex}(\Sigma) \quad \text{iff} \quad \forall\, \Sigma_j = E'_1 \cup \ldots \cup E'_n \in \text{Pref}(\Sigma),\ |E_k| \geq |E'_k| \text{ for } k = 1,n$$

The definition of Lex-consequence is identical to the one presented in the case of a flat knowledge base, namely a formula ϕ is a Lex-consequence of Σ if and only if it is entailed from each element of Lex-consequence.

Proposition: Each preferred-consequence of Σ is also a Lex-consequence of Σ. The converse is false.

The proof is obvious since the Lex-consequence generally uses a strict subset of Pref(Σ).

Now, we propose to extend the argumentative inference to layered knowledge bases, and to compare it with the inferences proposed above.

Definition: A sub-base Σ_i of Σ is said to be an argument for a formula ϕ with a weight α if it satisfies the following conditions:

(i) $\Sigma_i \nvdash \bot$

(ii) $\Sigma_i \vdash_\pi (\phi\ \alpha)$

(iii) $\forall\, (\psi\ \beta) \in \Sigma_i,\ \Sigma_i - \{(\psi\ \beta)\} \nvdash_\pi (\phi\ \alpha)$

Definition: A formula ϕ is said to be an argumentative consequence of Σ, denoted by $\Sigma \vdash_{\mathcal{A}} (\phi\ \alpha)$, if and only if:

(i) there exists an argument for $(\phi\ \alpha)$ in Σ , and

(ii) for each argument of $(\neg\phi\ \beta)$ in Σ, we have $\beta < \alpha$.

We now give the procedure which determines if ϕ is an argumentative consequence of a stratified knowledge base $\Sigma = B_1 \cup \ldots \cup B_n$. The procedure presupposes the existence of an algorithm which checks if there exists an argument for a given formula in some base. This can be achieved by using the variant of a refutation method proposed for example in (Lin, 1987).

The procedure is based on a construction of the maximal argument of ϕ and its contradiction. First we start with the sub-base B_1, and we check if there is a consistent sub-base of B_1 which entails ϕ or $\neg\phi$. If the response is respectively Yes-No then ϕ is an argumentative consequence of Σ with a weight $\alpha_1 = 1$, by symmetry if the response is No-Yes then $\neg\phi$ is in this case the argumentative consequence of Σ. Now if the response is Yes-Yes then we have a conflict. If the answer corresponds to one of the answers given above then the algorithm stops.

In the last case, if the response is No-No we repeat the same cycle described above with $B_1 \cup B_2$. The algorithm stops when we use all the knowledge base Σ.
Notice that the algorithm involves at most n steps, and the cost of the step depends on the complexity of the refutation procedure used in (2).

Procedure Argumentative_consequence (Input: Σ, ϕ; Output: Result)
Begin

(* The variable Result has four possible values:
- t: if $\Sigma \vdash_{\mathcal{A}} \phi$
- f: if $\Sigma \vdash_{\mathcal{A}} \neg\phi$
- T: if $\Sigma \nvdash_{\mathcal{A}} \phi$ and $\Sigma \nvdash_{\mathcal{A}} \neg\phi$
- $\perp$: conflict (case where the maximal arguments for ϕ and $\neg\phi$ have the same weights *)

1. i := 1; Result1 = Result2 = false; S = B_1;

2. (a) If there exist an argument for ϕ in S then Result1 = true;
(b) If there exist an argument for $\neg\phi$ in S then Result2 = true;

3. (a) if Result1 = true and Result2 = false then Result = t;
(b) if Result1 = false and Result2 = true then Result = f;
(c) if Result1 = true and Result2 = true the Result = $\perp$;
(d) if i = n then Result = T;
else i := i + 1; S = S $\cup$ B_i; go to 2;

End

As discussed in the case of a flat knowledge base, the inference relation $\vdash_{\mathcal{A}}$ is non-monotonic, and if our knowledge base is consistent then the set of formulas generated by $\vdash_{\mathcal{A}}$ is identical to the one generated by the "classical" inference rule $\vdash_{\pi}$.

The next proposition shows that $\vdash_{\mathcal{A}}$ is a faithful extension of the inference π-entailment.

Proposition 4: If $\Sigma \vdash_{\pi} (\phi\ \alpha)$ then $\Sigma \vdash_{\mathcal{A}} (\phi\ \alpha)$. The converse is false.

Proof

-- Indeed, if $\pi(\Sigma)$ infers a formula ϕ with weight α, then we have an argument for ϕ with weight α, since $\pi(\Sigma)$ is consistent we cannot infer $\neg\phi$ with weight greater than α, therefore ϕ is also an argumentative consequence of Σ.

-- To show that the converse is false, let us consider the following example where our knowledge contains the following set of formulas $\Sigma = \{\{\neg B\}, \{B\}, \{C\}\}$, the formula C is an argumentative consequence of Σ while it is not a possibilistic consequence of Σ.

Proposition 5: Each π+Free-consequence of Σ is also an argumentative consequence of Σ. The converse is false

Proof:

• Indeed if a formula is π+Free-entailed from Σ, then we have two possibilities:

- $\pi(\Sigma)$ entails alone the formula ϕ, then ϕ is a π-consequence of Σ and therefore using the previous proposition ϕ is also an argumentative consequence.

- $\pi(\Sigma) \cup \text{IFree}(\Sigma)$ entails ϕ but $\pi(\Sigma)$ does not entail it, there is an argument for ϕ. Assume there is also an argument Σ' for $\neg\phi$ in Σ, which means that $\Sigma' \vdash \neg\phi$ and therefore $\pi(\Sigma) \cup \text{IFree}(\Sigma) \cup \Sigma'$ is inconsistent, by consequence this contradicts the fact that the rules in IFree are outside any contradiction.

• Let us consider the following example where our knowledge base is:

$$\Sigma = \{\{A\}, \{\neg A \vee C, \neg C\}, \{\neg C \vee D\}\}$$

This knowledge is inconsistent, and $\prod(\Sigma) \cup \text{IFree}(\Sigma)$ is $\{\{A\}, \{\neg C \vee D\}\}$ and therefore it is not possible to deduce D, while D is an argumentative consequence of Σ.

The argumentative consequence is not comparable to the Pref-consequence nor the Lex-consequence, as we see in the following example:

Example

• Let $\Sigma = \{\{A, \neg B \vee \neg A, B, C\}, \{\neg C \vee \neg A\}, \{\neg A \vee D\}\}$

We have:

- $\text{Pref}(\Sigma) = \{\{\{A, \neg B \vee \neg A, C\}, \{\neg A \vee D\}\}, \{\{A, B, C\}, \{\neg A \vee D\}\},$
$\{\{\neg B \vee \neg A, B, C\}, \{\neg C \vee \neg A\}, \{\neg A \vee D\}\}\}$

- $\text{Lex}(\Sigma) = \{\{\{\neg B \vee \neg A, B, C\}, \{\neg C \vee \neg A\}, \{\neg A \vee D\}\}\}$.

Then $\neg A$ is a Lex-consequence of Σ while it is not an argumentative consequence, since A is also present in Σ. Note that one may object to the deletion of A from $\text{Lex}(\Sigma)$, given its high priority. Hence the Lex-consequence looks debatable. In contrast, D is an argumentative consequence (it derives from $\{A, \neg A \vee D\}$ while it is not a Pref-consequence and a lex-consequence. Again the Pref-consequence forgets the argument, because A and $\neg A \vee D$ never belong to the same preferred subbase.

• Let $\Sigma = \{\{A\}, \{\neg A\}, \{\neg A \vee \neg D, A \vee D\}\}$, we have $\text{Pref}(\Sigma) = \{\{A\}, \{\neg A \vee \neg D, A \vee D\}\}$. In this case $\neg D$ is a Pref-consequence, while it is not an argumentative consequence of Σ. Again, the argument for D is killed by $\text{Pref}(\Sigma)$.

As we have done in the non-stratified case, we summarize the relationships between the different consequence relations:

Figure 2: A comparative study of inference relations in stratified knowledge

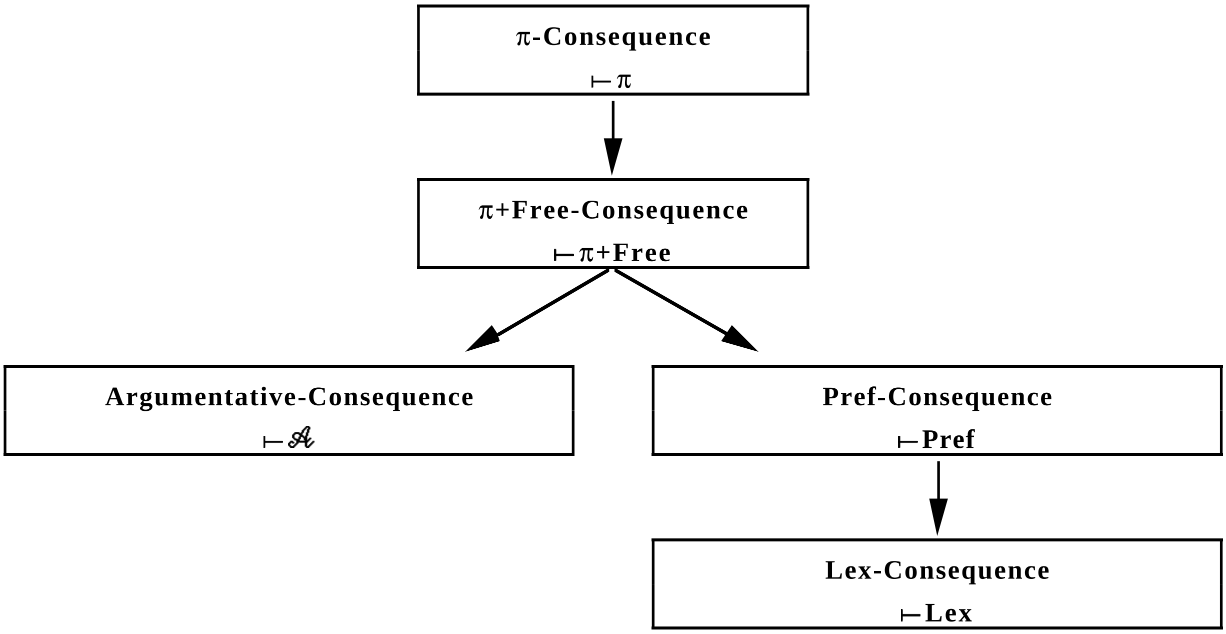

5. Paraconsistent-Like Reasoning in Layered Knowledge Bases

In the preceding sections we have seen how in the case of flat and prioritized knowledge bases it is possible to use consistent subparts of it in order to define different types of consequences which are still meaningful. Levels of priority or of certainty attached to formulas have also been used to distinguish between strong and less strong arguments in favor of a proposition or of its contrary. However it is possible to go one step further in the use of the certainty or priority levels by i) attaching to a proposition ϕ not only the (greatest) weight α attached to a possibilistic logical proof of ϕ (in the sense of $\vdash_{\pi}$) from a consistent subbase, but also the weight β attached to the strongest argument in favor of $\neg\phi$ if any, and ii) by continuing to infer from premises such (ϕ, α, β) propagating the weights α and β. It will enable us distinguishing between consequences obtained only from "free" propositions in the knowledge base Σ for which $\beta = 0$ (i.e. propositions for which there is no argument in Σ in favor of their negation), and consequences obtained using also propositions which are not free (for which there exist both a weighted argument in their favor and a weighted argument in favor of their negation).

More formally, the idea is first to attach to any proposition in the considered possibilistic knowledge base Σ two numbers reflecting the extent to which we have some certainty that the proposition is true and to what extent we have some certainty that the proposition is false, and then to provide some extended resolution rule enabling us to infer from such propositions. Namely, a prioritized knowledge base Σ as introduced in Section 4 is now viewed as a possibilistic knowledge base where each priority level α attached to a formula ϕ is interpreted as a lower bound of a necessity measure N, i.e. $N(\phi) \geq \alpha$ (Dubois, Lang and Prade, 1989, 1991, 1992b). See the Appendix for a refresher on possibilistic logic. For each ϕ, such that $(\phi\ \alpha)$ is in Σ, we compute the largest weight α' associated with an argument for ϕ and the largest weight β' associated with an argument for $\neg\phi$ in the sense of Section 4. If there exists no argument in favor of $\neg\phi$, we will take $\beta' = 0$; it means in this case that $(\phi\ \alpha)$ is among the free elements of Σ since ϕ is not involved in the inconsistency of Σ (otherwise there would exist an argument in favor of $\neg\phi$).

In the general case, we shall say that ϕ has a level of "paraconsistency" equal to $\min(\alpha',\beta')$. Classically and roughly speaking, the idea of paraconsistency, first introduced by Da Costa (1963), is to say that we have a paraconsistent knowledge about ϕ if we both want to state ϕ and to state $\neg\phi$. It corresponds to the situation where we have conflicting information about ϕ. In a paraconsistent logic we do no want to have every formula ψ deducible as soon as the knowledge base contains ϕ and $\neg\phi$ (as it is the case in classical logic). The idea of paraconsistency is "local" by constrast with the usual view of inconsistency which considers the knowledge base in a global way. It is why we speak here of paraconsistent information when $\min(\alpha',\beta') > 0$. Note that in this process we may improve the lower bound α into a larger one α' if $\exists\ \Sigma_i \subseteq \Sigma$, Σ_i consistent and $\Sigma_i \vdash_\pi (\phi\ \alpha')$ (similarly for β' if $(\neg\phi\ \beta)$ is already present in Σ). Then Σ is changed into a new knowledge base Σ' where each formula $(\phi\ \alpha)$ of Σ is replaced by $(\phi\ \alpha'\ \beta')$. Moreover if $\alpha' < \beta'$, i.e. the certainty in favor of $\neg\phi$ is greater than the one in favor of ϕ, we replace $(\phi\ \alpha'\ \beta')$ by $(\neg\phi\ \beta'\ \alpha')$. If ϕ is under a clausal form, $\neg\phi$ is a conjunction $\varphi_1 \wedge \ldots \wedge \varphi_n$; in this case we will replace $(\neg\phi\ \beta'\ \alpha')$ by the clauses $(\varphi_i\ \beta'\ \alpha')$, i = 1,n in order to keep Σ' under a clausal form if Σ was under a clausal form. See the Appendix for a justification. Let us consider an example

$\Sigma = \{(\neg A \vee B\ \ \alpha), (A\ \beta), (\neg B\ \gamma)\ (B\ \delta), (\neg B \vee C\ \ \varepsilon)1\ (\neg C\ \rho)\}$.

Then

$$\begin{aligned}\Sigma' = \{\ & (\neg A \vee B\ \ \alpha\ \ \min(\beta,\gamma)), (A\ \beta\ \min(\alpha,\gamma)),\\ & (\neg B\ \max(\gamma, \min(\rho,\varepsilon)), \max(\delta, \min(\alpha,\beta))),\\ & (B\ \max(\delta, \min(\alpha,\beta)), \max(\gamma, \min(\rho,\varepsilon))),\\ & (\neg B \vee C\ \ \varepsilon\ \ \min(\rho,\delta)), (\neg C\ \rho\ \min(\varepsilon,\delta))\ \}.\end{aligned}$$

Depending on the ordering between the weights we will keep either $(\neg B\ x\ y)$ or $(B\ y\ x)$ depending if $x > y$ or $y > x$. If $x = y$ we will keep both of them in Σ'.

In a second step an extended resolution rule can be proposed in order to infer from propositions in Σ'. This rule, which can be justified in a possibilistic framework (see the Appendix for a proof) is the following, expressed in clausal form

$$\frac{\begin{array}{c}(A \vee B \;\; \alpha' \; \beta') \\ (\neg B \vee C \;\; \gamma' \; \delta')\end{array}}{(A \vee C \;\; \varepsilon' \; \rho')}$$

with $\varepsilon' = \min(\max(\gamma',\beta'),\ \max(\alpha',\delta'))$
$\rho' = \max(\beta',\delta')$.

When $\beta' = \delta' = 0$, i.e. the premises are not paraconsistent, we obtain $\varepsilon' = \min(\alpha',\gamma')$, $\rho' = 0$. In other words we recover the possibilistic resolution rule in the consistent case. Clearly we have $\varepsilon' \geq \rho'$, i.e. the inference preserves the inequality between the weight. We also observe that the degree of paraconsistency of the conclusion namely $\min(\varepsilon',\rho') = \max(\beta',\delta')$ is equal to the maximum of the degrees of paraconsistency of the two premises namely $\min(\alpha',\beta') = \beta'$ and $\min(\gamma',\delta') = \delta'$. Thus the inference rule extends the standard possibilistic resolution and in case of paraconsistent premise(s), propagates this paraconsistency to the conclusion. In the case where one of the premises is not paraconsistent, i.e. $\beta' = 0$ for instance, the degree of certainty $\varepsilon' = \min(\gamma', \max(\alpha',\delta'))$ of the conclusion is greater than its degree of paraconsistency $\rho' = \delta'$ only if the degree of certainty α' of the non-paraconsistent premise is greater than δ' and $\gamma' > \delta'$ (i.e. $\gamma' \neq \delta'$). Otherwise the conclusion which is obtained is such that $\varepsilon' = \rho' = \delta'$, i.e. nothing emerges from inconsistency.

Let us consider an example

Σ = { (A 1), (¬A ∨ B 0.8), (¬B 0.6), (¬A ∨ C 0.5), (¬D 0.3), (¬A ∨ D 0.4), (¬A ∨ E 0.7), (¬F ∨ G 0.5), (F 1), (¬F 0.2), (¬H ∨ I 0.3), (I 0.4) }.

Observe that $\Sigma \vdash_\pi (\bot\ 0.6)$, i.e. the global level of inconsistency of the base is 0.6. Then we have

Σ' = { (A 1 0.6), (¬A ∨ B 0,8 0,6), (¬A ∨ C 0.5 0), (D 0.4 0.3), (¬A ∨ D 0.4 0.3), (¬A ∨ E 0.7 0), (¬F ∨ G 0.5 0), (F 1 0.2), (¬H ∨ I 0.3 0), (I 0.4 0) }.

Applying the "paraconsistent" resolution rule yields

(C 0.6 0.6), (E 0.7 0.6), (G 0.5 0.2), (I 0.3 0).

This shows that

– non-paraconsistent premises such as (¬A ∨ C 0.5 0) with a rather low degree of certainty resolved with another premise (here (A 1 0.6)) whose level of paraconsistency is larger than this degree of certainty, lead to fully blurred paraconsistent conclusions, here (C 0.6 0.6). By contrast $\vdash_{Free+\pi}$ would enable to get (C 0.5), while the refutation procedure used in $\vdash_\pi$ yields (C 0.6), reflecting the global inconsistency of the base

- if the non-paraconsistent premise is sufficiently certain with respect to the paraconsistency of the other premise, e.g. $(\neg A \vee E\ \ 0.7)$, and $(A\ \ 1\ \ 0.6)$, the conclusion, here $(E\ \ 0.7\ \ 0.6)$ would not be completely blurred, even if this certainty is less than the global level of inconsistency of the base (e.g. $(G\ \ 0.5\ \ 0.2)$ obtained from $(\neg F \vee G\ \ 0.5\ \ 0)$, $(F\ \ 1\ \ 0.2)$)
- if the premises are not paraconsistent, (e.g. $(\neg H \vee I\ \ 0.3\ \ 0)$, $(H\ \ 0.4\ \ 0)$), we obtain a non-paraconsistent conclusion, as with $\vdash_{\text{Free}+\pi}$.

Generally speaking, if a clause $A \vee B$ is more paraconsistent than the clause $\neg B \vee C$ is certain, then $A \vee C$ will be completely blurred by paraconsistency, since from a logical point of view, being more certain that $\neg A \wedge \neg B$ is true than we are certain that $\neg B \vee C$ is true, $A \vee B, \neg A \wedge \neg B \vdash A \vee C$ applies with a greater level of certainty than $A \vee B, \neg B \vee C \vdash A \vee C$.

We can observe that using the "paraconsistent" resolution rule locally in knowledge base Σ', may yield a the same proposition with different weights, namely $(\phi\ \alpha'\ \beta')$, $(\phi\ \alpha''\ \beta'')$. There only exists a partial ordering between such pairs, namely a more certain and less paraconsistent conclusion should be preferred; this is less obvious when we have to choose between a highly certain but highly paraconsistent conclusion and a conclusion with low certainty and low paraconsistency.

Lastly observe that, *for the formulae explicitly in Σ*, the paraconsistent approach gives the same results as $\vdash_{\mathcal{A}}$. They differ for other conclusions since the paraconsistent approach propagates the effects of local inconsistency.

6. Combining knowledge bases

Baral et al. (1992) propose several approaches to combine knowledge bases, and one of them is very similar to what (Brewka, 1989) calls "preferred sub-theories". The idea is to assume a total ordering between different bases $\Sigma_1 < \Sigma_2 < \ldots < \Sigma_n$, such that Σ_i is more reliable than Σ_j for j<i. A resulting base is constructed from Σ_1 by adding as many formulas as possible from Σ_2 (wrt consistency criterion), then as many formulas as possible from Σ_3, and so on. The principal problem is that the resulting base is not unique.

Cholvy (1992) has proposed two approaches to merge bases according to suspicious attitude or to trusting attitude. The suspicious attitude is very conservative since for example the result of merging two knowledge bases $\Sigma_1 < \Sigma_2$ is equal to the union of the two bases if they are not conflictual, and is equal to Σ_1 in other case. In the trusting attitude, the approach is very similar to (Baral et al., 1992) and produces always one resulting base, but unfortunately the approach is very restrictive since the knowledges bases to be merged must be sets of literals.

In the context of possibilistic logic, an approach has been proposed in (Dubois, Lang and Prade, 1992c) for the fusion of knowledge bases. To each base Σ_i is associated its possibility distribution π_i, and the question is how to construct the possibility distribution π of the resulting base Σ_{result}. One way to define π is to consider the weighted union of the possibility distributions associated with each source; namely $\pi = \max(\pi_1, \ldots, \pi_n)$, the corresponding knowledge base is the intersection of the deductively closed bases $Cn(\Sigma_i), \ldots, Cn(\Sigma_n)$. It is clear that this approach is very cautious. In the same paper, another approach has been proposed considering the union of the deductively closed bases $Cn(\Sigma_i), \ldots, Cn(\Sigma_n)$. However when the resulting base is inconsistent, then some formulas will be inhibited by the drowning effect (Benferhat et al., 1993).

We suggest a new approach to merge n knowledge bases $\Sigma_1, \ldots, \Sigma_n$. For this aim we use of a variation of the argumentative consequence relation, denoted by $\vdash_{\mathcal{AM}}$ $\mathcal{M}$ for the multi-sources, and which is defined in the following way:

$$\Sigma_1, \ldots, \Sigma_n \vdash_{\mathcal{AM}} (\phi\ \alpha) \text{ iff:}$$
$$\exists \Sigma_i, \text{ such that } \Sigma_i \vdash (\phi\ \alpha), \text{ and}$$
$$\nexists \Sigma_j, \text{ such that } \Sigma_j \vdash (\neg\phi\ \beta) \text{ such that } \beta > \alpha.$$

Then the knowledge base resulting is: $\Sigma_{result} = \{(\phi\ \alpha) / \Sigma_1, \ldots, \Sigma_n \vdash_{\mathcal{AM}} (\phi\ \alpha)\}$. Because Σ_{result} can be inconsistent, this approach should be used for question-answering purposes only, and each response should be accompanied with its argument.

8. Conclusion

The proposed notion of argumentative inference is appearing for several reasons. First it is an extension of classical inference (in the flat case) and possibilistic inference (in the layered case) that copes with inconsistency in a very "ecological" way. Namely it is very faithful to the actual contents of the knowledge base, and does not do away with information contained in it, as opposed to the approaches based on preferred and lexicographically preferred subbases. It avoids the drowning effect of standard possibilistic logic by salvaging sentences whose level of entrenchment is low but are not involved in any contradiction set. Another advantage is that it is amenable to efficient standard implementation methods based on classical resolution.

Also it avoids outright contradictory responses (such that ϕ and $\neg\phi$), although several deduced sentences can be globally inconsistent. But as pointed out earlier, the arguments supporting a set of more than two globally contradictory sentences are distinct, so that the reality of this contradiction is debatable, and only reflects the presence of different points of view. Anyway it seems that it is the price to pay in order to remain faithful to an inconsistent knowledge base. Another result of the paper is the use of local contradictions as a specific weight attached to sentences. This approach only partially avoids the drowning effect, but leads to more informative responses than possibilistic logic since not only the certainty of the formula is evaluated, but also its level of conflict.

Lastly it would be interesting to apply the above result to default reasoning and study in such a framework the argumentative inference and the one proposed by Simari and Loui (1992).

References

C. Baral, S. Kraus, J. Minker, V.S. Subrahmanian (1992) Combining knowledge bases consisting in first order theories. Computational Intelligence, 8(1), 45-71.

S. Benferhat, D. Dubois, H. Prade (1992) Representing default rules in possibilistic logic. Proc. of the 3rd Inter. Conf. on Principles of Knowledge Representation and Reasoning (KR'92) (J. Allen, R. Fikes, E Sandewall, eds.), Morgan & Kaufmann, Cambridge, Mass., Oct 26-29.

S. Benferhat, C. Cayrol, D. Dubois, J. Lang, H. Prade (1993) Inconsistency management and prioritized syntax-based entailment. Submitted for publication.

G. Brewka (1989) Preferred subtheories: an extended logical framework for default reasoning. Proc. of the 11th Inter. Joint Conf. on Artificial Intelligence (IJCAI'89), Detroit, Michigan, Aug. 20-25, 1043-1048.

L. Cholvy (1992) A logical approach to multi-sources reasoning. In: Applied Logic Conference: Logic at Work, Amesterdam.

N.C.A. Da Costa (1963) Calculus propositionnels pour les systèmes formels inconsistants. Compte Rendu Acad. des Sciences (Paris), 257, 3790-3792.

D. Dubois, J. Lang, H. Prade (1989) Automated reasoning using possibilistic logic: semantics, belief revision and variable certainty weights. Preprints of the 5th Workshop on Uncertainty in Artificial Intelligence, Windsor, Ontario, Aug. 18-20, 81-87. Revised version in IEEE Trans. on Data and Knowledge Engineering, 1993, to appear.

D. Dubois, J. Lang, H. Prade (1991) Handling uncertainty, context, vague predicates, and partial inconsistency in possibilistic logic. Preprints of the Workshop on Fuzzy Logic in Artificial Intelligence (IJCAI), Sydney, Australia, Aug. 24-30, 13-23.

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

D. Dubois, J. Lang, H. Prade (1992b) Possibilistic logic. In: Handbook of Logic in Artificial Intelligence and Logic Programming, Vol. 3 (D.M. Gabbay, ed.), Oxford University Press, to appear.

D. Dubois, J. Lang, H. Prade (1992c) Dealing with multi-source information in possibilistic logic. Proc. of the 10th Europ. Conf. on Artificial Intelligence (ECAI'92), Vienna, Austria, Aug. 3-7, 38-42.

J. De Kleer (1986a) An assumption-based TMS. Artificial Intelligence, 28, 127-162.

J. Fox, P. Krause, S. Ambler (1992) Arguments, contradictions and practical reasoning. Proc. of the 10th Europ. Conf. on Artificial Intelligence (ECAI'92), Vienna, Austria, Aug. 3-7, 623-627.

P. Gärdenfors (1988) Knowledge in Flux — Modeling the Dynamic of Epistemic States. The MIT Press, Cambridge, Mass.

F. Lin (1987) Reasoning in the presence of inconsistency. Proc. of the 6th National Conf. on Artificial Intelligence (AAAI'87), Seattle, Wash., July 13-17, 139-143.

J. Pearl (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 (M. Vardi, ed.), Morgan Kaufmann, San Mateo, CA, 121-135.

G. Pinkas, R.P. Loui (1992) Reasoning from inconsistency: a taxonomy of principles for resolving conflict. Proc. of the 3rd Inter. Conf. on Principles of Knowledge Representation and Reasoning (KR'92) (J. Allen, R. Fikes, E Sandewall, eds.), Morgan-Kaufmann, Cambridge, Mass., Oct. 26-29, 709-719.

J.L. Pollock (1987) Defeasible reasoning. Cognitive Science, 11, 481-518.

D. Poole (1988) A logical framework for default reasoning. Artificial Intelligence, 36, .

J.A. Reggia, D.S. Nau, and P.Y. Wang (1985) A formal model of diagnosis inference. Information Sciences, 37, 227-285.

R. Reiter (1987) A theory of diagnosis from first principles. Artificial Intelligence 32, 57-95.

N. Resher, R. Manor (1970) On inference from inconsistent premises. Theory and Decision, 1, 179-219.

N. Roos (1992) A logic for reasoning with inconsistent knowledge. Artificial Intelligence, 57, 69-103.

K. Satoh (1990) A probabilistic interpretation for lazy non-monotonic reasoning. Proc. of the 8th Proc. of the National Conf. on Artificial Intelligence (AAAI-90), Boston, July 29-Aug. 3, 659-664.

G.R. Simari, R.P. Loui (1992) A mathematical treatment of defeasible reasoning and its implementation. Artificial Intelligence, 53, 125-157.

S. Toulmin (1956) The Uses of Argument. Cambridge University Press, Cambridge, UK.

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

Appendix

In the semantics of possibilistic logic, we associate with each formula $(\phi\ \alpha)$ in a knowledge base Σ a possibility distribution $\pi_{(\phi\ \alpha)}$ defined on the set of interpretations Ω (assumed to be finite here for simplicity)

$$\begin{aligned}\pi_{(\phi\ \alpha)}(\omega) &= 1 && \text{if } \omega \vDash \phi \\ &= 1 - \alpha && \text{if } \omega \vDash \neg\phi \end{aligned} \qquad \text{(A1)}$$

Then a global possibility distribution π_Σ is obtained by combining all the $\pi_{(\phi\ \alpha)}$, by

$$\pi_\Sigma = \min_{(\phi\ \alpha)\in\Sigma} \pi_{(\phi\ \alpha)}.$$

π_Σ reflects a preferential ordering among the interpretations induced by the levels of certainty attached to the formulas. It can be shown that the consistency of the classical knowledge base obtained from Σ by ignoring the weights α is equivalent to

$$\exists\omega,\ \pi_\Sigma(\omega) = 1. \qquad \text{(A2)}$$

To π_Σ is associated a possibility measure (Zadeh, 1978)

$$\Pi_\Sigma(\phi) = \max\{\pi_\Sigma(\omega) \in [0,1],\ \omega \models \phi\},\ \forall\phi, \qquad \text{(A3)}$$

and a dual necessity measure $N_\Sigma(\phi) = 1 - \Pi_\Sigma(\phi)$. In this case the possibility distribution π_Σ is said to be normalized. Then we have the following equivalence (Dubois, Lang and Prade, 1989, 1991, 1992b)

$$\alpha \text{ is the greatest element in } [0,1] \text{ such that}$$
$$\Sigma \vdash_\pi (\psi\ \alpha) \Leftrightarrow N_\Sigma(\psi) = \alpha \Leftrightarrow \Pi_\Sigma(\neg\psi) = 1 - \alpha \qquad \text{(A4)}$$

where $\vdash_\pi$ corresponds to the repeated application of the extended resolution rule

$$\frac{\begin{array}{c}(A \vee B\ \ \alpha)\\(\neg B \vee C\ \ \gamma)\end{array}}{(A \vee C\ \ \min(\alpha,\gamma))} \qquad \text{(A5)}$$

to Σ, put in clausal form, to which the clauses corresponding to $(\neg\psi\ 1)$ are added (extension of the refutation procedure). Then α is the greatest weight attached to the empty clause which can be obtained by repeated application of (A5) to Σ. Since $N(\phi \wedge \psi) = \min(N(\phi),N(\psi))$, the transformation of the formulas in Σ in a clausal form is straightforward. See (Dubois, Lang and Prade, 1989, 1991, 1992b) for details.

When Σ is not consistent, i.e. π_Σ is not normalized, we define the degree of inconsistency of Σ as

$$\text{Incons}(\Sigma) = 1 - \max_{\omega\in\Omega} \pi_\Sigma(\omega). \qquad \text{(A6)}$$

Then it can be shown that any consequence $(\psi\ \alpha)$ obtained by $\vdash_\pi$ such that $\alpha > \text{Incons}(\Sigma)$ is still valid since based on a consistent subpart of Σ. Unfortunately there may exist free elements in Σ with weights less that Incons(Σ). Their use in the inference process inhibited by the level of inconsistency of Σ. See Dubois, Lang and Prade, 1989, 1991, 1992b.

Here we propose to keep the basic features of the possibilistic inference machinery, namely i) to represent the pieces of knowledge by means of possibility distributions, ii) to combine possibility distributions, iii) to compute possibility measures associated with the resulting possibility distributions. However, the idea is not to associate with Σ a unique possibility distribution as in the consistent case, but rather to distinguish in Σ its free subpart Σ_{Free} from the other elements in Σ which are not free, and then to build one (normalized) possibility distribution π_{Free} associated with the free subpart of Σ and for each non-free element ϕ in Σ to define the subnormalized possibility distribution

$$\pi_{\phi}(\omega) = 1 - \alpha' \text{ if } \omega \models \phi$$
$$\pi_{\phi}(\omega) = 1 - \beta' \text{ if } \omega \models \neg\phi$$

where α' and β' are respectively the greatest α' and β' such that $\exists\ \Sigma_i \subset \Sigma$, Σ_i consistent and $\Sigma_i \vdash_\pi (\phi\ \alpha)$, $\exists\ \Sigma_j \subset \Sigma$, Σ_j consistent and $\Sigma_j \vdash_\pi (\neg\phi\ \beta)$. Such a proposition will be said to be paraconsistent in Σ at the degree $\min(\alpha',\beta')$ since we are both certain at the degree $\alpha' > 0$ that ϕ is true and certain at the degree $\beta' > 0$ that ϕ is false. In the consistent case all the formulas in Σ are such that $\min(\alpha',\beta') = 0$ (or more generally possibility distributions attached to free formulas, which are among the components giving π_{Free} by combination)..

In case nothing can be inferred from Σ_{Free} about a statement of interest, we can combine π_{Free} [1] with some subnormalized possibility distribution and still compute a possibility measure, but then the result of the combination will be subnormalized, and the result of the inference will be necessarily paraconsistent. However this paraconsistency will reflect only the subnormalized possibility distribution(s) which are indeed used in the inference process. In other words, their level of paraconsistency may be less than Incons(Σ).

At the syntactic level, this will correspond to an inference rule starting with the two premises

$$(A \vee B\ \ \alpha'\ \beta')$$
$$(\neg B \vee C\ \ \gamma'\ \delta')$$

These two premises translate respectively into the possibility distributions

$$\pi_1(\omega) = 1 - \beta' \text{ if } \omega \models A \vee B\ ;\ \pi_1(\omega) = 1 - \alpha' \text{ if } \omega \models \neg A \wedge \neg B$$
$$\pi_2(\omega) = 1 - \delta' \text{ if } \omega \models \neg B \vee C\ ;\ \pi_2(\omega) = 1 - \gamma' \text{ if } \omega \models B \wedge \neg C$$

By combination we obtain $\pi = \min(\pi_1,\pi_2)$ defined by

$$\begin{aligned}
\pi(\omega) &= \min(1 - \beta', 1 - \delta') \text{ if } \omega \models A \wedge B \wedge C\\
&= \min(1 - \beta', 1 - \delta') \text{ if } \omega \models A \wedge \neg B \wedge C\\
&= \min(1 - \beta', 1 - \gamma') \text{ if } \omega \models A \wedge B \wedge \neg C\\
&= \min(1 - \beta', 1 - \delta') \text{ if } \omega \models A \wedge \neg B \wedge \neg C\\
&= \min(1 - \beta', 1 - \delta') \text{ if } \omega \models \neg A \wedge B \wedge C\\
&= \min(1 - \alpha', 1 - \delta') \text{ if } \omega \models \neg A \wedge \neg B \wedge C\\
&= \min(1 - \beta', 1 - \gamma') \text{ if } \omega \models \neg A \wedge B \wedge \neg C\\
&= \min(1 - \alpha', 1 - \delta') \text{ if } \omega \models \neg A \wedge \neg B \wedge \neg C
\end{aligned}$$

[1] Or more generally possibility distributions attached to free formulas, which are among the components giving π_{Free} by combination.

Thus we get

$$\Pi(A \vee C) = \max(\min(1-\beta', 1-\delta'), \min(1-\beta', 1-\gamma'), \min(1-\alpha', 1-\delta')) \triangleq 1-\rho'$$
$$\Pi(\neg A \wedge \neg C) = \max(\min(1-\beta', 1-\gamma'), \min(1-\alpha', 1-\delta')) \triangleq 1-\varepsilon'$$

i.e. we have the following pattern

$$\frac{\begin{array}{c}(A \vee B \;\; \alpha' \; \beta') \\ (\neg B \vee C \;\; \gamma' \; \delta')\end{array}}{(A \vee C \;\; \varepsilon', \rho')}$$

with

$$\varepsilon' = \min(\max(\beta',\gamma'), \max(\alpha',\delta'))$$
$$\rho' = \min(\max(\beta',\delta'), \varepsilon').$$

N.B. : A particular case of this formula where $\beta' = 0$ corresponding to a modus ponens pattern with only one paraconsistent premise, has been suggested in Dubois, Lang and Prade (1991). It seems reasonable to restrict the use of this rule to case where $\alpha' \geq \beta'$ and $\gamma' \geq \delta'$, i.e. when the premises are considered as more certain than their contrary.

In order to always work with clauses such that $\alpha' \geq \beta'$, we have to transform the weighted formula $(A \vee B \;\; \alpha \; \beta)$ with $\alpha < \beta$ into a set of such clauses. This formula is represented by a possibility distribution

$$\pi(\omega) = 1-\beta \text{ if } \omega \vDash A \vee B \;\; ; \;\; \pi(\omega) = 1-\alpha \text{ if } \omega \vDash \neg A \wedge \neg B.$$

Thus

$$\Pi(\neg A) = \max(1-\alpha, 1-\beta) = 1-\alpha, \; \Pi(A) = 1-\beta, \; \Pi(\neg B) = \max(1-\alpha, 1-\beta) = 1-\alpha, \; \Pi(B) = 1-\beta.$$

Then this corresponds to two weighted clauses $(\neg A \; \beta \; \alpha)$ and $(\neg B \; \beta \; \alpha)$.