Belief functions and default reasoning

S. Benferhat[1] **A. Saffiotti**[2] **P. Smets**[2]

benferha@irit.fr asaffio@ulb.ac.be psmets@ulb.ac.be

[1] IRIT (UPS)
118 Route de Narbonne
Toulouse 31062 Cedex, France

[2] IRIDIA, Université Libre de Bruxelles
50 av. F. Roosevelt
B-1050 Bruxelles, Belgium

Abstract. We present a new approach to dealing with default information based on the theory of belief functions. Our semantic structures, inspired by Adams' ε-semantics, are epsilon-belief assignments, where values committed to focal elements are either close to 0 or close to 1. We define two systems based on these structures, and relate them to other non-monotonic systems presented in the literature. We show that our second system correctly addresses the well-known problems of specificity, irrelevance, blocking of inheritance, ambiguity, and redundancy.

1. Introduction

Default reasoning is the process of drawing conclusions from i) a set of general rules which may have exceptions, and ii) a set of facts representing the available information (which is often incomplete). The conclusions so drawn are only plausible and can be revised in the light of the new information. The desirable properties for a consequence relation that capture default reasoning have been discussed at length in the AI literature. They can be summarized as follows.

- *Rarionality*: the conceqence operator used to generate plausible conclusions from a knowledge base should satisfy the rationality postulates proposed by Kraus, Lehmann and Magidor (1990).

- *Specificity*: results obtained from sub-classes should override those obtained from super-classes (Touretzky, 1984). For example, from the knowledge base Δ = {'Birds fly", "Penguins do not fly", "Penguins are birds"}, one should deduce that birds which are penguins do not fly, since penguins are a subclass of birds.

- *Property inheritance:* objects should inherit properties from super-classes unless there is contradiction on that property. For example, from the previous Δ, one should deduce that birds that are red fly, since being red is irrelevant to flying. Also, if we add the rule "Birds have legs", then one should deduce that penguins have legs too, since having legs is not a conflicting property. Failure to perform these deductions is referred to as the problem of "irrelevance" and of "inheritance blocking", respectively.

- *Ambiguity preservation*: in a situation where we have an argument in favor of a proposition, and an independent argument in favor of its negation, we should not conclude anything about that proposition.[1] The most popular example is the so-called Nixon-diamond: knowing that "Quakers are pacifists", "Republican are not pacifists", and Nixon is both a Quaker and a republican, one should not deduce that Nixon is a pacifist, nor that he is not.

- *Syntax independence*: the consequences of a knowledge base should not depend on the syntactical form used. In particular, they should not be sensitive to duplications of rules in the knowledge base ("redundancy").

In the last decade there have been several proposals for reasoning with default information. Some of them are based on the use of uncertainty models such as probability theory (Adams, 1975, Pearl, 1988), or possibility theory (Dubois and Prade, 1988, Benferhat et al., 1992). Up to now, however, no single system has been reported that correctly addresses all of the desiderata above. In this paper, we show how we can use belief functions, originally developed for modeling uncertainty (Shafer, 1976, Smets, 1988), to build a non-monotonic system that gives a satisfactory answer to all of the above issues.

There have already been a few works on representing default information with belief functions, e.g., (Hsia, 1991, Smets and Hsia, 1991). These works require the assessment of numeric values, whose origin is often an open question. Finding a solution free from such assessments would somehow avoid the problem of the origin of the numbers. In this paper, we give another interpretation of default information by using a class of "extreme" belief functions, called epsilon-belief functions, whose non-null masses are either close to 0 or close to 1. The idea of using extreme values is not new to plausible reasoning: for instance, Adams (1975) and Pearl (1990) use extreme probabilities to encode default information, and De Kleer (1990) and Poole (1993) use extreme probabilities for diagnosis.

The rest of this paper is organized as follows. In the next section,we give a short reminder on Adams' ε-semantics and Pearl' System **Z**, and recall a few notions of the theory of

[1] Note that this is different from the situation of inconsistency, where we have an argument which supports both a conclusion and its contrary, as in "$\alpha \rightarrow \beta$" and "$\alpha \rightarrow \neg\beta$".

belief functions (see (Shafer, 1976, Smets, 1988) for a complete exposition). In section 3, we introduce ε-belief functions, and show how to use them to define a non-monotonic consequence relation. This relation turns out to be too cautious, and we define two more relations in sections 4 and 5. The first one is based on the least-commitment principle, and is equivalent to system **Z**. The second one uses Dempster's rule of combination, and is incomparable with the current systems. In section 6, we study this relation in more detail and show that it correctly addresses all of the issues above.

2. Background

We are interested in *default rules* of the form "generally, if we have α then we have β", where α and β are formulae of some underlying language $\mathcal{L}$. For the goals of this paper, we assume that $\mathcal{L}$ is a classical propositional language. An *interpretation* for $\mathcal{L}$ is an assignment of a truth value in {T, F} to each formula of $\mathcal{L}$ in accordance with the classical rules of propositional calculus, we denote by Ω the set of all such interpretations (also called *worlds*). We say that an interpretation ω is a *model* of a formula α, and write $\omega \models \alpha$ iff $\omega(\alpha) = T$, and denote by [α] the set of all the models of α. We write a default rule "generally, if α then β" as $\alpha \rightarrow \beta$, where α and β are formulae of $\mathcal{L}$. Note that "→" is a *non-classical* arrow, and it should not be confused with material implication. Given a default rule $d = \alpha \rightarrow \beta$, we denote by ϕ_d the formula of $\mathcal{L}$ obtained by replacing → by material implication, namely, $\phi_d = \neg\alpha \vee \beta$. A *default base* is a multiset $\Delta = \{\alpha_i \rightarrow \beta_i, i=1,...,n\}$ of default rules. We emphasize that a base is a multiset rather than a set, i.e., $\Delta = \{\alpha \rightarrow \beta\}$ is different from $\Delta' = \{\alpha \rightarrow \beta, \alpha \rightarrow \beta\}$.

We use default bases to represent background knowledge about what normally is the case. Given a base Δ, we are interested in defining a consequence relation $\mid\sim$ between formulae of $\mathcal{L}$ that tells us which consequences we can reasonably drawn from the known facts. We would like $\mid\sim$ to fulfill the desiderata listed above. For example, given the base $\Delta = \{b \rightarrow f, p \rightarrow b, p \rightarrow \neg f\}$ (where "b" stands for "bird", "f" for "flies", and "p" for "penguin"), we would like to have $b \mid\sim f$ and $b \wedge p \mid\sim \neg f$, but not $b \wedge p \mid\sim f$.

2.1. Probabilistic semantics for default rules

Adams (1975), and later Pearl (1988), have suggested a probabilistic interpretation where a default rule $\alpha \rightarrow \beta$ is read as a constraint $P(\beta|\alpha) > 1-\varepsilon$, with P a probability distribution and ε an infinitesimal positive number. Given a set of defaults Δ, they construct a class of probability distributions A_ε such that, for each distribution P in A_ε and each default $\alpha \rightarrow \beta$ in Δ, $P(\beta|\alpha) > 1-\varepsilon$. A formula β is said to be an *ε–consequence* of α with respect to Δ, denoted by $\alpha \mid\sim_\varepsilon \beta$, if for each $P \in A_\varepsilon$ there exists a real function O such that $\lim_{\varepsilon \rightarrow 0} O(\varepsilon) = 0$ and $P(\beta|\alpha) > 1-O(\varepsilon)$. Said differently, β is a consequence of α with respect to Δ if the conditional probability $P(\beta|\alpha)$ is very high provided that the conditional probability of each default in Δ is very high. Lehmann and Magidor (1992) have shown that ε-consequence is equivalent to the system **P** of Kraus et al (1990), which is commonly regarded as the minimal core of any "reasonable" non-monotonic system.

Adams' ε-consequence is not entirely satisfactory. For example, it suffers from the problem of irrelevance mentioned above: from the default "Generally, birds fly", ε-consequence does not deduce that red birds fly also. To overcome this limitation, Pearl (1990) has proposed a default reasoning system, called **Z**, based on a ranking of default rules that respects the notion of specificity. Given a default base $\Delta = \{\alpha_i \rightarrow \beta_i \mid i = 1,...,m\}$, Pearl gives a method to rank-order the rules in Δ such that the least specific rules (i.e. with most general antecedents) get the least priority. To do this, he defines the notion of *tolerance*: a rule $\alpha \rightarrow \beta$ is said to be tolerated by a base $\{\alpha_i \rightarrow \beta_i, i = 1,...,m\}$ iff $\{\alpha \wedge \beta, \neg\alpha_1 \vee \beta_1, \ldots, \neg\alpha_m \vee \beta_m\}$ is consistent. Then, he partitions Δ into an ordered set $\{\Delta_1, \Delta_2, \ldots, \Delta_k\}$ such that rules in Δ_i are *tolerated* by all rules in $\Delta_i \cup \ldots \cup \Delta_k$. From this partition, Pearl induces a ranking K on worlds and, from this, a ranking z on formulae. Rouhly speaking, $K(\omega)$ corresponds to the index of the lowest sub-base that contains a rule violated by ω, and $z(\alpha)$ is the minimum rank of a model of α —so, $z(\alpha)$ can be read as a degree of "abnormality" of α with respect to the rules in Δ. Finally, Pearl defines a non-monotonic inference relation, denoted here by $\mid\!\sim_Z$, as follows

$$\alpha \mid\!\sim_Z \beta \iff z(\alpha \wedge \beta) < z(\alpha \wedge \neg\beta).$$

An equivalent treatment of default information has been done in the framework of possibility theory (Benferhat et al., 1992).

2.2. A reminder on belief functions

Let Ω be a finite set of worlds, one of them being the actual world. A *basic belief assignment* on Ω is a function $m: 2^\Omega \rightarrow [0, 1]$ that satisfies:[2]

$$m(\emptyset) = 0$$

$$\sum_{A \subseteq \Omega} m(A) = 1$$

The term m(A), called the *basic belief mass* given to A, represents the part of a total and finite amount of belief that supports the fact that the actual world belongs to A and does not

---

[2] In the transferable belief model (Smets and Kennes, 1994), belief functions are not necessarily normalized, i.e., we can have $m(\emptyset) > 0$. Normalization is assumed here as we only study ratios between $bel(B|A)$ and $bel(\Omega|A)$, which corresponds to studying the normalized belief functions.

support the fact that the actual world belongs to a strict subset of A. Any subset A of Ω for which m(A) > 0 is called a *focal element*.

An agent's belief can equivalently be represented by the function bel: $2^{\Omega} \rightarrow [0, 1]$, called a *belief function*, defined by

$$\text{bel}(A) = \sum_{B:\, B \subseteq A} m(B).$$

The relation between m and bel is one-to-one. The term bel(A) represents the degree of belief, of necessary support, that the actual world belongs to A. Related to bel is another function pl: $2^{\Omega} \rightarrow [0, 1]$, called a *plausibility function*, given by

$$\text{pl}(A) = \sum_{B:\, B \cap A \neq \emptyset} m(B).$$

The term pl(A) quantifies the degree of plausibility, of potential support, that the agent could give to the fact that the actual world belongs to A. When the focal elements of a basic belief assignment are singletons, then bel = pl is a probability measure. When the focal elements $A_1,\ldots,A_n$ are nested (that is, $A_1 \subseteq \ldots \subseteq A_n$), bel is called a *consonant belief function*, and bel is a necessity measure and pl is a possibility measure (Zadeh, 1978, Dubois and Prade, 1988). When m has at most one focal element A ≠ Ω, it is called a *simple support function*.

When a new piece of evidence telling that the actual world belongs to A becomes available to the agent, his/her belief is revised by the application of the so-called Dempster's rule of conditioning. The basic belief mass m(X) that was supporting the subset X of Ω, now supports X∩A. This transfer of belief masses is described by the following relation, where bel(.|A) denotes the *conditional belief function*.

$$\text{bel}(X|A) = \frac{\text{bel}(X \cup A^c) - \text{bel}(A^c)}{\text{bel}(\Omega) - \text{bel}(A^c)}$$

The impact of a piece of evidence E that bears on Ω is represented by a belief function bel that describes the agent's beliefs on Ω given E (and nothing else). Suppose the agent receives two distinct pieces of evidence E_1 and E_2, and let bel_1 and bel_2 be the belief functions induced by each evidence individually. The combined effect of E_1 and E_2 is represented by the belief function $\text{bel}_1 \oplus \text{bel}_2$ obtained by Dempster's rule of combination. The corresponding basic belief assignment, denoted by $m_1 \oplus m_2$, is given by

$$m_1 \oplus m_2(A) = \frac{\sum_{B \cap C = A} m_1(B) \cdot m_2(C)}{1 - \sum_{B \cap C = \emptyset} m_1(B) \cdot m_2(C)}.$$

3. Epsilon-belief functions

In this section, we extend the definition of ε-semantics in a belief function framework. First, we introduce the notion of epsilon-belief functions, whose values are either close to 0 or close to 1. Next, we interpret a default rule $\alpha \rightarrow \beta$ as meaning that the conditional belief $bel([\beta]|[\alpha])$ is close to 1. Finally, we define a consequence relation in a natural way: $\alpha \rightarrow \beta$ follows from a base of defaults Δ if, for each epsilon-belief function which satisfies all the default rules in Δ (i.e., the conditional belief of each rule is close to 1), we have $bel([\beta]|[\alpha])$ close to 1. It turns out that this definition gives us the same results as Adams' ε-consequence relation.

Definition 1. An *epsilon-mass assignment* on Ω is a function $m_\varepsilon: 2^\Omega \rightarrow [0,1]$ such that, for each $A \subseteq \Omega$, either $m_\varepsilon(A) = 0$, or $m_\varepsilon(A) = \varepsilon_A$, or $m_\varepsilon(A) = 1-\varepsilon_A$, where ε_A is an infinitesimal. The vector $\varepsilon = (\varepsilon_1,...,\varepsilon_k)$ of the infinitesimals appearing in m_ε is called the *parameter* of m. The belief function induced from m_ε is called an *epsilon-belief function* (ε-bf).

Throughout this paper, we denote by bel_ε and pl_ε the belief and plausibility functions corresponding to a given epsilon-mass assignment m_ε.

Definition 2. Let bel_ε be an ε-bf with parameter ε. We say that bel_ε is an *εbf-model* of a default rule $\alpha \rightarrow \beta$, and write $bel_\varepsilon \models \alpha \rightarrow \beta$, iff $\lim_{\varepsilon \rightarrow 0} bel_\varepsilon([\beta]|[\alpha]) = 1$, where the limit is taken with respect to all the elements in ε going to 0.

When working with values that depend on ε, we will often use the notion of one value being infinitely larger than another, written $a >_\infty b$. We say that $a >_\infty b$ if $\lim_{\varepsilon \rightarrow 0} b/a = 0$. We also say that a and b are of the same order, written $a \approx b$, if $\lim_{\varepsilon \rightarrow 0} b/a = c$, with $c \neq 0$ and finite. The following properties will be useful for working with ε-bf's.[3]

Lemma 1. Let bel_ε be an ε-bf. For any α, β formulae of $\mathfrak{L}$,

a) $\lim_{\varepsilon \rightarrow 0} bel_\varepsilon([\beta]|[\alpha]) = 1$ iff $pl_\varepsilon([\alpha \wedge \beta]) >_\infty pl_\varepsilon([\alpha \wedge \neg\beta])$.

b) $pl_\varepsilon([\alpha]) \approx \max\{pl_\varepsilon(\omega) | \omega \models \alpha\}$.

It is interesting to note that the satisfaction relation $\models$ for ε-bf can also be defined in terms of preferential semantics (Shoham, 1988, Kraus et al., 1990). This will turn out to be useful to relate our systems to other existing ones by making use of known results. To create the link, we associate each ε-bf bel_ε with a preferential order among worlds in Ω as follows.

[3] The proofs of all the results can be found in the extended version of this paper (Benferhat et al., 1995), available from the authors.

Definition 3. Let bel_ε be an ε-bf on Ω, and α a formula of $\mathcal{L}$. The *bel-preference* induced by bel_ε is the partial order $\prec_\varepsilon$ given by: $\omega \prec_\varepsilon \omega'$ iff $pl_\varepsilon(\{\omega'\}) >_\infty pl_\varepsilon(\{\omega\})$. A model ω of α is called a *bel-preferred model* of α if there is no other world ω' that satisfies α such that $\omega' \prec_\varepsilon \omega$.

Lemma 2. Let bel_ε be an ε-bf on Ω. For any α, β formulae of $\mathcal{L}$, $bel_\varepsilon \models \alpha \rightarrow \beta$ if, and only if, each bel-preferred model of α satisfies β.

Our next step is to use ε-bf models to define the notion of *entailment* for default bases, i.e., to define which conditional assertions α→β are entailed by a default base Δ. Our first solution is a direct adaptation of the usual definition of logical entailment. We say that an ε-bf bel_ε is an εbf-model of Δ, written $bel_\varepsilon \models \Delta$, iff bel_ε is an *εbf-model* of all the rules in Δ. We denote by $Bel_\varepsilon(\Delta)$ the set of all the εbf-models of Δ. Entailment is then given by

(BF) $\quad \Delta \models_{bf} \alpha \rightarrow \beta$ iff for any bel_ε in $Bel_\varepsilon(\Delta)$, $bel_\varepsilon \models \alpha \rightarrow \beta$.

Suppose that the default base Δ represents our background knowledge. Then, for any given formula α, **(BF)** tells us which formulae are expected to be consequences of α — namely, all the formulae β for which α→β is entailed by Δ. We call *bf-consequence* this consequence relation, denoted by $\mid\!\sim_{bf}$, and let

$$\alpha \mid\!\sim_{bf} \beta \quad \text{iff} \quad \Delta \models_{bf} \alpha \rightarrow \beta.$$

Bf-consequence turns out to be equivalent to the system **P** of Kraus, Lehmann and Magidor (1990), which in turn is equivalent to Adams' ε-system.

Theorem 1. For a given Δ, $\alpha \mid\!\sim_{bf} \beta$ if, and only if, $\alpha \mid\!\sim_{\mathbf{P}} \beta$.

<u>Sketch of the proof.</u> For the only-if part, we note that Adams' infinitesimal probability distributions are a special case of our ε-bf, and then $\alpha \mid\!\sim_{bf} \beta$ only if $\alpha \mid\!\sim_{\varepsilon} \beta$ only if $\alpha \mid\!\sim_{\mathbf{P}} \beta$. To prove the if part, we use lemma 2 to show that each inference relation induced by any bel_ε in $Bel_\varepsilon(\Delta)$ is preferential. So, $\mid\!\sim_{bf}$ satisfies the rules of **P**, and then it contains all preferential consequences of Δ. ■

This result shows that we can use (infinitesimal) belief functions to give an alternative formalization of the systems ε and **P**. It also shows that $\mid\!\sim_{bf}$ suffers from the same limitations of these systems, in particular, it does not solve the problems of irrelevance and blocking of inheritance. In the next two sections, we propose two ways to define a more bold consequence relation by restricting our attention in **(BF)** to just *some* of the models in $Bel_\varepsilon(\Delta)$.

4. Entailment based on the least-commitment principle

One way to select some of the εbf-models of Δ is by using the notion of being minimally informative: intuitively, we want to look at the consequences of “only knowing” Δ (and nothing more). A similar approach has been taken, for the case of possibility measures, by Benferhat et al (1992). We recall the following

Definition 4. Let bel_1 and bel_2 two (epsilon-) belief functions over Ω. Then, bel_1 is *less committed* that bel_2 iff, for any $A \subseteq \Omega$, $pl_1(A) \geq pl_2(A)$.

The *least-commitment principle* (Smets, 1988) states that, in order to model an item of information by a belief function, we should use the least committed belief function that is compatible with the information. Note that the least committed belief function representing a formula α is given by the simple support function that gives mass 1 to [α] and 0 anywhere else. We show how to build an εbf-model of Δ based on this principle. We start by allocating a quasi-unitary mass to the set $[\phi_\Delta]$ of the worlds where (the propositional equivalent of) all the defaults in Δ are satisfied, and the remaining mass ε to Ω. If there were no conflict in the defaults, this allocation would be an εbf-model of Δ. When there are conflicts, however, this ε-bf will not satisfy some of the defaults in Δ — namely, those that inherit a conflicting property from a more general class. Then, we put aside the defaults that are already satisfied, and put almost all of the free mass ε on the set $[\phi_{\Delta'}]$ corresponding to the still unsatisfied defaults, leaving a small ε’ on Ω. This new ε-bf is an εbf-model of Δ’ (and of Δ) if we have no conflicts in Δ’. Otherwise, we iterate the procedure until the ε-bf will satisfy all the defaults in Δ.

More precisely, let $\varepsilon = (\varepsilon_1, ..., \varepsilon_n)$ be a vector of infinitesimals such that

(ε_{LC}) $\varepsilon_i >_\infty \varepsilon_{i+1}$ for any i = 1, ..., n–1.

where n is the cardinality of Δ. We build an ε-bf in the following way.

Step 0. Let i = 0, $\Delta_0 = \Delta$, $sat_0 = \emptyset$, m_0 s.t. $m_0(\Omega) = 1$ and $m_0(A) = 0$ otherwise.

Step 1. Repeat until $\Delta_i = \emptyset$

1a. Let i = i+1

1b. Let $\Delta_i = \Delta_{i-1} - sat_{i-1}$

1c. Let bel_i be the ε-bf given by:
$m_i(\Omega) = \varepsilon_i$, $m_i([\phi_{\Delta_i}]) = m_{i-1}(\Omega) - \varepsilon_i$, $m_i(A) = m_{i-1}(A)$ otherwise.

1d. Let $sat_i = \{d \in \Delta_i \mid bel_i \models d\}$

1e. If $sat_i = \emptyset$ and $\Delta_i \neq \emptyset$ then Fail.

Step 2. Return bel_{i-1}.

Note that all the focal elements are nested —the inner one being $[\phi_\Delta]$— and then the final ε-bf returned by Step 2 is a consonant belief function. The procedure fails to find an ε-bf if Δ is inconsistent, in this case, we have $sat_i = \emptyset$ and $\Delta_i \neq \emptyset$.

Example 1. Let $\Delta = \{ b \to f, p \to b, p \to \neg f \}$ (where "b" stands for "bird", "f" for "flies", and "p" for "penguin"), and let us apply the previous algorithm. We have $\Delta_1 = \Delta$, and bel_1 given by: $m_1[\phi_\Delta] = 1-\varepsilon_1$, $m_1(\Omega) = \varepsilon_1$, and $m_1(.) = 0$ elsewhere. We compute the set sat_1 of defaults which are satisfied by bel_1. We have:

$pl_1([b\wedge f]) = 1-\varepsilon_1+\varepsilon_1 = 1$, $pl_1([b\wedge\neg f]) = \varepsilon_1$,
$pl_1([p\wedge b]) = \varepsilon_1$, $pl_1([p\wedge\neg b]) = \varepsilon_1$,
$pl_1([p\wedge\neg f]) = \varepsilon_1$, $pl_1([p\wedge f]) = \varepsilon_1$.

Therefore, $sat_1 = \{b\to f\}$. We iterate, and get Δ_2 by removing $b\to f$ from Δ. Then:

$$m_2(\Omega) = \varepsilon_2,\ m_2([\phi_{\Delta_2}]) = \varepsilon_1-\varepsilon_2,\ m_2([\phi_{\Delta_2}]) = 1-\varepsilon_1,\ m_2(A) = 0 \text{ otherwise}$$

with ε_1 infinitely larger than ε_2. We have:

$$pl_2([p\wedge b]) = \varepsilon_1,\ pl_2([p\wedge\neg b]) = \varepsilon_2,\ pl_2([p\wedge\neg f]) = \varepsilon_1, \text{ and } pl_2([p\wedge f]) = \varepsilon_2.$$

All the defaults in Δ_2 are satisfied and the algorithm ends returning bel_2. ■

We denote by $Bel_{LC}(\Delta)$ the family of ε-bf's built by the procedure above — the elements of this family differ in the choice of ε, provided that (ε_{LC}) is satisfied. This family "behaves well" for our goals: it is a subset of $Bel_\varepsilon(\Delta)$, and it induces a unique ordering $\prec_\varepsilon$ on the worlds in **Ω**. The latter property means that we can decide entailment by just looking at one element of $Bel_{LC}(\Delta)$.

Lemma 3. Let Δ be a default base. Then:

a) Any element of $Bel_{LC}(\Delta)$ is an εbf-model of Δ.
b) Let bel_1 and bel_2 be two elements of $Bel_{LC}(\Delta)$, and $\prec_1$ and $\prec_2$ the corresponding orderings induced on **Ω**. Then, $\prec_1 \equiv \prec_2$.

It is interesting to compare $Bel_{LC}(\Delta)$ with the result of the stratification proposed by Pearl (1988). We can show that the focal elements of $Bel_{LC}(\Delta)$ correspond exactly to the elements of the partition of Δ obtained by **Z**.

Lemma 4. Let $\Delta = \Delta_1\cup\ldots\cup\Delta_n$ be the stratification given by system **Z**, and let bel_i the ε-bf built by our algorithm at step i. Then, for any default $\alpha\to\beta$ in Δ,

a) $\alpha\to\beta$ is tolerated by Δ iff $pl_1([\alpha_i]) = 1$
b) $\alpha\to\beta \in \Delta_i$ iff $bel_i \models \alpha\to\beta$.

So, our algorithm produces the same ranking over Δ than **Z**. However, our approach does not require an a priori definition of the notion of "tolerance", but relies on the notion of being "less committed". Which of these two notions provides a more natural starting point is a matter of opinion.

We now use the set $Bel_{LC}(\Delta)$ to give our second definition of entailment. It is similar to **(BF)**, but we restrict the attention to the εbf-models that are in $Bel_{LC}(\Delta)$.

(LC) $\Delta \models_{lc} \alpha \to \beta$ iff for any bel_ε in $Bel_{LC}(\Delta)$, $bel_\varepsilon \models \alpha \to \beta$.

As we did for **(BF)**, we define the lc-consequence relation for a given base Δ by:

$$\alpha \mid\!\sim_{lc} \beta \quad \text{iff} \quad \Delta \models_{lc} \alpha \to \beta.$$

Example 2. We can use the ε-bf bel_2 built in Example 1 to check that we have $b \wedge p \mid\!\sim_{lc} \neg f$. In fact, we have $pl_2([b \wedge p \wedge \neg f]) = \varepsilon_1$, $pl_2([b \wedge p \wedge f]) = \varepsilon_2$, and $\varepsilon_1 >_\infty \varepsilon_2$. Hence, by lemma 1, $\lim_{\varepsilon \to 0} bel_2([\neg f]|[b \wedge p]) = 1$, and then $b \wedge p \mid\!\sim_{lc} \neg f$. ■

As $Bel_{LC}(\Delta)$ is a subset of $Bel_\varepsilon(\Delta)$, lc-consequences include bf-consequences, in particular, they include the consequences of system **P**. As it turns out, the LC consequence relation is strictly larger than $\mid\!\sim_{\mathbf{P}}$, and precisely coincides with Pearl's system **Z**, as it appears from lemma 4 above.

Theorem 2. For a given Δ, $\alpha \mid\!\sim_{lc} \beta$ if, and only if, $\alpha \mid\!\sim_{\mathbf{Z}} \beta$.

5. Entailment based on Dempster's rule

The second way that we propose to strengthen the entailment relation **(BF)** is by considering only the εbf-models of Δ that can be built by using Dempster's rule of combination. The intuitive argument goes as follows. Suppose we regard each default in Δ as being one item of evidence provided by one of several *distinct* sources of information.[4] Then, it makes sense to represent each default individually by one belief function, and combine these belief functions by Dempster's rule to obtain a representation of the aggregate effect of all the defaults. We can then define entailment by looking at the conditionals that are satisfied by the combined belief function.

---

[4] The construction given here should extend to the case where each source S_i of information provides an entire base Δ_i of defaults: we could use the least-commitment principle to build a representative belief function for each Δ_i as above, and then combine these representatives by Dempster's rule.

There are two technical choices to be made. The first one is which belief function to use to represent each default. Sticking to the arguments used in the last section, we propose to use the least committed εbf that satisfy that default: for a default $d = \alpha \rightarrow \beta$, this is the simple support function that allocates a mass of $1-\varepsilon_d$ to $[\phi_d]$, and the remaining mass ε_d to Ω. We note by ssf_d this function. Given a default base $\Delta = \{d_1,...,d_n\}$, then, we consider the ε-bf obtained by combining all the ssf_d's by Dempster's rule:

$$bel_{\oplus} = \oplus\{ssf_d | d \in \Delta\}.$$

The ε-bf's obtained in this way may be complex: the focal elements are not nested, and the mass values may include products of several ε_d (or $1-\varepsilon_d$) for different d. Luckily, the plausibility of each world can be approximated in a simple way.

Lemma 5. Let $bel_{\oplus}$ be an ε-bf built from Δ as above. Then, for any world ω in Ω,

$$pl_{\oplus}(\{\omega\}) \approx \prod\{\varepsilon_d \mid d \in \Delta \text{ s.t. } \omega \not\models \phi_d\}.$$

The second question, now, is how to choose the infinitesimals ε_d's. We impose constraints on the elements of ε based on the two following principles:

- *Auto-deduction principle.* We want to have $bel_{\oplus} \models \alpha \rightarrow \beta$ for each default $\alpha \rightarrow \beta$ in Δ (i.e., $bel_{\oplus}$ must be an εbf-model of Δ).
- *Least commitment principle.* We want each ssf_d to be least committed (hence, each ε_d must be as large as possible).

For each $d = \alpha \rightarrow \beta$ in Δ, the first principle generates the constraint (cf. lemma 1)

$$\max\{pl_{\oplus}(\omega) \text{ s.t. } \omega \models \alpha \wedge \beta\} >_{\infty} \max\{pl_{\oplus}(\omega) \text{ s.t. } \omega \models \alpha \wedge \neg\beta\}$$

which is equivalent to (cf. Lemma 5)

$$\textbf{(C}_\textbf{d}\textbf{)} \qquad \max_{\omega \models \alpha \wedge \beta} \left\{ \prod_{d:\ \omega \not\models \phi_d} \varepsilon_d \right\} >_{\infty} \max_{\omega \models \alpha \wedge \neg\beta} \left\{ \prod_{d:\ \omega \not\models \phi_d} \varepsilon_d \right\}.$$

By solving the system given by all the C_d's, we can get constraints of the form $\varepsilon_i >_{\infty} \varepsilon_j$ between some of the elements of ε (and between products of elements). The second principle can be used to sanction equivalencies between unconstrained elements of ε by the following argument. Suppose that, for some ε_i and ε_j, neither $\varepsilon_i >_{\infty} \varepsilon_j$ nor $\varepsilon_j >_{\infty} \varepsilon_i$ is the case, as both ε_i and ε_j should be made as large as possible, no one can be larger than the other, and then $\varepsilon_i = \varepsilon_j$.

We now describe an algorithm to solve a set *C* of constraints of the form **(C_d)**. We call *term* any product of elements of ε, and *complex term* any term containing at least two elements of ε. The algorithm returns a set $\xi = \{\xi_0,\dots,\xi_n\}$ of equivalence classes of terms such that: 1) all the terms in a class are of the same order, 2) any term in ξ_i is infinitely larger than any term in ξ_j when $i < j$, and 3) any element of ε is in some class.

Step 0. Let $i = 0$, $A_t = \{t_i$ s.t. t_i is a term and $1 >_\infty t_i$ is in $C\}$.

- 0.a Let ξ_0 be a set of ε_i such that there exists a complex term t in A_t which contains ε_i
- 0.b. If $\xi_0 = \emptyset$ then $\xi_0 = A_t$ else $\xi_1 = A_t - \xi_0$ and $i = 1$

Step 1. Repeat until $C = \emptyset$

- 1a. Let $i = i+1$. Remove from C any satisfied constraint
- 1b. Let A_t be a set of terms in C which does not appear in the right side of any constraint of C
- 1c. Let ξ_i be a set of ε_i which does not appear in any of set $\xi_{j<i}$ and where there exists a complex term t in A_t which contains ε_i
- 1.d. If $\xi_i \neq \emptyset$ then $A_t := A_t - \xi_i$ and $i = i+1$
- 1.e. Let A_ε be the set of ε_i which does not appear neither in any constraint of C nor in any $\xi_{j<i}$.
- 1.f. Let $\xi_i = A_t + A_\varepsilon$

Step 2. Return the sets $\xi_{j=1,i}$.

Example 3. Let again $\Delta = \{b\to f,\ p\to\neg f,\ p\to b\}$. The simple support functions corresponding to the three defaults in Δ are given by

$m_1([\neg b\vee f]) = 1 - \varepsilon_1,\ m_1(\Omega) = \varepsilon_1,$
$m_2([\neg p\vee\neg f]) = 1 - \varepsilon_2,\ m_2(\Omega) = \varepsilon_2,$ and
$m_3([\neg p\vee b]) = 1 - \varepsilon_3,\ m_3(\Omega) = \varepsilon_3.$

The requirement of auto-deductivity gives us the following three constraints:

$\max\{pl_\oplus(\{\omega\})|\omega\models b\wedge f\} >_\infty \max\{pl_\oplus(\{\omega\})|\omega\models b\wedge\neg f\}$ i.e., $1 >_\infty \varepsilon_1$
$\max\{pl_\oplus(\{\omega\})|\omega\models p\wedge\neg f\} >_\infty \max\{pl_\oplus(\{\omega\})|\omega\models p\wedge f\}$ i.e., $\max\{\varepsilon_1, \varepsilon_3\} >_\infty \varepsilon_2$
$\max\{pl_\oplus(\{\omega\})|\omega\models p\wedge b\} >_\infty \max\{pl_\oplus(\{\omega\})|\omega\models p\wedge\neg b\}$ i.e., $\max\{\varepsilon_1, \varepsilon_2\} >_\infty \varepsilon_3$

Let us apply the previous algorithm. The set ξ_0 contains exactly one element ε_1. When we put ε_0 in the highest value then all the constraints will be satisfied. Therefore the result is $\xi_0=\{\varepsilon_1\}$, $\xi_1=\{\varepsilon_2, \varepsilon_3\}$, and $\varepsilon_1 >_\infty \varepsilon_2 = \varepsilon_3$. ■

We denote by $Bel_\oplus(\Delta)$ the family of ε-bf's built by Dempster's rule and whose parameter ε satisfy the constraints above. As we did for the $Bel_{LC}(\Delta)$ family, we make sure that the elements of $Bel_\oplus(\Delta)$ have the right properties for our goals: they are εbf-models of Δ, and they induce a unique ordering on Ω.

Lemma 6. Let Δ be a default base. Then:

a) Any element of $Bel_{\oplus}(\Delta)$ is an εbf-model of Δ.

b) Let bel_1 and bel_2 be two elements of $Bel_{\oplus}(\Delta)$, and $\prec_1$ and $\prec_2$ the corresponding orderings induced on Ω. Then, $\prec_1 \equiv \prec_2$.

Our third and last definition of entailment, called LCD (Least-Commitment plus Dempster's rule), is obtained by focusing on the εbf-models that are in $Bel_{\oplus}(\Delta)$.

(LCD) $\Delta \models_{lcd} \alpha \rightarrow \beta$ iff for any bel_ε in $Bel_{\oplus}(\Delta)$, $bel_\varepsilon \models \alpha \rightarrow \beta$

As usual, we define a corresponding consequence relation $\mid\!\sim_{lcd}$ by:

$$\alpha \mid\!\sim_{lcd} \beta \quad \text{iff} \quad \Delta \models_{lcd} \alpha \rightarrow \beta.$$

Note that, as all the elements of $Bel_{\oplus}(\Delta)$ are εbf-model of Δ, the $\mid\!\sim_{lcd}$ relation is as least as strong as $\mid\!\sim_{bf}$. In particular, $\mid\!\sim_{lcd}$ satisfies the KLM properties for system **P** (Kraus et al., 1990). In fact, **LCD** is strictly stronger than **P**. For example, **LCD** correctly addresses the irrelevance problem, as shown by the following example.

Example 4. We first show how to use the result of the previous example to verify that $b \wedge p \mid\!\sim_{lcd} \neg f$. In fact, by applying lemmas 1(b) and 5, we have $pl_{\oplus}([b \wedge p \wedge f]) \approx \varepsilon_2$ and $pl_{\oplus}([b \wedge p \wedge \neg f]) \approx \varepsilon_1$, and we know that $\varepsilon_1 >_\infty \varepsilon_2$. Next, consider a new property "red" (r) unrelated to b, p and f. We expect that red birds fly (note that this is not the case in system **P**). For any bel in $Bel_{\oplus}$, and its corresponding pl, we have

$$pl([b \wedge r \wedge f]) \approx \max\{pl([b \wedge r \wedge f \wedge p]), pl([b \wedge r \wedge f \wedge \neg p])\} \approx \max\{\varepsilon_2, 1\} = 1$$
$$pl([b \wedge r \wedge \neg f]) \approx \max\{pl([b \wedge r \wedge \neg f \wedge p]), pl([b \wedge r \wedge \neg f \wedge \neg p])\} \approx \max\{\varepsilon_1, \varepsilon_1\} = \varepsilon_1$$

Hence $pl([b \wedge r \wedge f]) >_\infty pl([b \wedge r \wedge \neg f])$, which implies $b \wedge r \mid\!\sim_{lcd} f$ as desired. ■

The following theorem summarizes the relation between **(LCD)** and system **P**.

Theorem 3. For a given Δ, if $\alpha \mid\!\sim_{\mathbf{P}} \beta$ then $\alpha \mid\!\sim_{lcd} \beta$. The converse is not true.

6 Analysis of the LCD consequence relation

We have seen that the LCD consequence relation gives us strictly more than system **P**, in particular, it correctly answers the problem of irrelevance. In this section, we study in more detail the patterns of reasoning that are captured by LCD. To do this, we consider the desiderata listed in the introduction, and show how LCD addresses them. We also contrast the LCD solution with the one obtained by other existing systems that go beyond system **P**.

We start by property inheritance. Several systems, including Pearl's system **Z**, suffer from the problem of inheritance blocking. The canonical example is built by adding to the usual penguin problem the default $b \rightarrow l$ (generally, birds have legs). From this, system **Z** cannot deduce that penguins have legs also, i.e., $p \not\mid\!\sim_{Z} l$.[5] By contrast, LCD allows that deduction, as shown below.

Example 5. Let $\Delta = \{b \rightarrow f,\ p \rightarrow \neg f,\ p \rightarrow b,\ b \rightarrow l\}$, where l stands for legs. The simple support functions are those in Example 3, plus

$$m_4([\neg p \vee l]) = 1 - \varepsilon_4,\ \ m_4(\Omega) = \varepsilon_4.$$

The constraint that $pl_{\oplus}$ must satisfy are the same as in Example 3, plus

$$\max\{pl_{\oplus}(\{\omega\}) | \omega \models b \wedge l\} >_{\infty} \max\{pl_{\oplus}(\{\omega\}) | \omega \models b \wedge \neg l\} \qquad \text{i.e., } 1 >_{\infty} \varepsilon_4$$

We apply our algorithm to this set. The first layer, ξ_0, contains exactly two elements: ε_1 and ε_4. Once we constraint ε_1 and ε_4 to have the highest value, all the constraints are satisfied. Therefore we get: $\xi_0 = \{\varepsilon_1, \varepsilon_4\} >_{\infty} \xi_1 = \{\varepsilon_2, \varepsilon_3\}$. To see if penguins have legs, we compute

$$\begin{aligned} pl([p \wedge l]) &\approx \max\{pl([p \wedge l \wedge b \wedge f]), pl([p \wedge l \wedge b \wedge \neg f]), pl([p \wedge l \wedge \neg b \wedge f]), pl([p \wedge l \wedge \neg b \wedge \neg f])\} \\ &\approx \max\{\varepsilon_2, \varepsilon_1, \varepsilon_2\varepsilon_3, \varepsilon_2\} = \varepsilon_1 \end{aligned}$$

$$\begin{aligned} pl([p \wedge \neg l]) &\approx \max\{pl([p \wedge \neg l \wedge b \wedge f]), pl([p \wedge \neg l \wedge b \wedge \neg f]), pl([p \wedge \neg l \wedge \neg b \wedge f]), pl([p \wedge \neg l \wedge \neg b \wedge \neg f])\} \\ &\approx \max\{\varepsilon_2\varepsilon_4, \varepsilon_1\varepsilon_4, \varepsilon_2\varepsilon_3\varepsilon_4, \varepsilon_2\varepsilon_4\} = \varepsilon_1\varepsilon_4 \end{aligned}$$

Therefore, $pl([p \wedge l]) >_{\varepsilon} pl([p \wedge \neg l])$, which implies $p \mid\!\sim_{lcd} l$ as desired. ■

Another desiderata listed in the introduction was the ability to stay uncommitted in cases of ambiguity. The following example shows a case of ambiguity where system **Z** would deduce an undesired result, while LCD does not.

Example 6. Let $\Delta = \{b \rightarrow f,\ p \rightarrow \neg f,\ p \rightarrow b,\ m \rightarrow f\}$, where the last default means "Generally, objects with metal-wings fly." The simple support functions are again those in Example 3, plus

$$m_4([\neg m \vee f]) = 1 - \varepsilon_4,\ \ m_4(\Omega) = \varepsilon_4.$$

The constraints that $pl_{\oplus}$ must satisfy are the same as in Example 3, plus

[5] Goldszmidt and Pearl (1991) have suggested an extension of **Z**, called $\mathbf{Z}^+$, that correctly handles this example. Unfortunately, $\mathbf{Z}^+$ does not solve the problem of inheritance blocking in general: if we add the rules "Generally, legless birds are birds" and "Generally, legless birds do not have legs" to our base, then $\mathbf{Z}^+$ cannot deduce both of "Legless birds fly" and "Penguins have legs"—it will just deduce one of them, depending on the ranking. This problem is solved by LCD.

$\max\{pl_{\oplus}(\{\omega\})|\omega\models m\wedge f\} >_{\infty} \max\{pl_{\oplus}(\{\omega\})|\omega\models m\wedge\neg f\}$ i.e., $1 >_{\infty} \varepsilon_4$

We get the same ordering as before, with ε_4 in the top class: $\xi_0 = \{\varepsilon_1, \varepsilon_4\} >_{\infty} \xi_1 = \{\varepsilon_2, \varepsilon_3\}$. Consider now a bird that is a penguin and has metal wing. Given the base Δ, we can not say whether or not this beast will fly — we are in a case of ambiguity. And indeed we have:

$pl([b\wedge p\wedge m\wedge f]) \approx \varepsilon_2$
$pl([b\wedge p\wedge m\wedge \neg f]) \approx \varepsilon_1\varepsilon_4$.

As the ordering above says nothing of the relative magnitude of ε_2 and $\varepsilon_1\varepsilon_4$, we do not have neither $b\wedge p\wedge m \mid\!\sim_{lcd} f$ nor $b\wedge p\wedge m \mid\!\sim_{lcd} \neg f$. Notice, by contrast, that **Z** would give us the arbitrary result $b\wedge p\wedge m \mid\!\sim_{Z} \neg f$. ■

The following theorem summarizes the relation between LCD and system **Z**.

Theorem 4. The consequence relations $\mid\!\sim_{lcd}$ and $\mid\!\sim_{\mathbf{Z}}$ are incomparable.

Another solution to get non-monotonic inferences that go beyond those of system **P** is to use the so-called lexicographic approach (Dubois et al., 1991, Benferhat et al., 1993, Lehmann, 1993). One well know drawback of this approach is that it is syntax-sensitive: repetitions of the same default in Δ may change the result. We show that LCD is not sensitive to duplications.

Example 7. Consider a variant of the Quaker-Republican problem where the rule "Generally, Quaker are pacifists" has been duplicated: $\Delta = \{q\rightarrow p, q\rightarrow p, r\rightarrow\neg p\}$. By using a lexicographic approach, we would deduce $q\wedge r \rightarrow p$, while we would prefer to acknowledge the ambiguity and deduce nothing. In LCD, we have

$m_1([\neg q\vee p]) = 1- \varepsilon_1,\ m_1(\Omega) = \varepsilon_1,$
$m_2([\neg q\vee p]) = 1- \varepsilon_2,\ m_2(\Omega) = \varepsilon_2,$
$m_3([\neg r\vee\neg p]) = 1- \varepsilon_3,\ m_3(\Omega) = \varepsilon_3,$

together with the constraints

$\max\{pl_{\oplus}(\{\omega\})|\omega\models q\wedge p\} >_{\infty} \max\{pl_{\oplus}(\{\omega\})|\omega\models q\wedge\neg p\}$ i.e., $1 >_{\infty} \varepsilon_1\varepsilon_2$
$\max\{pl_{\oplus}(\{\omega\})|\omega\models q\wedge p\} >_{\infty} \max\{pl_{\oplus}(\{\omega\})|\omega\models q\wedge\neg p\}$ i.e., $1 >_{\infty} \varepsilon_1\varepsilon_2$
$\max\{pl_{\oplus}(\{\omega\})|\omega\models r\wedge\neg p\} >_{\infty} \max\{pl_{\oplus}(\{\omega\})|\omega\models r\wedge p\}$ i.e., $1 >_{\infty} \varepsilon_3$

By applying our algorithm, we get one single class, and so $\varepsilon_3 \approx \varepsilon_1\varepsilon_2$. Then

$pl([q\wedge r\wedge p]) \approx \varepsilon_3$ and $pl([q\wedge r\wedge\neg p]) \approx \varepsilon_1\varepsilon_2$,

and we have neither $q\wedge r \mid\!\sim_{lcd} p$ nor $q\wedge r \mid\!\sim_{lcd} \neg p$, as desired. ■

In general, we state the following

Theorem 5. LCD consequence and the lexicographic approach are incomparable.

We have shown that the LCD consequence relation behaves well with respect to all of the desiderata stated in the introduction. Unfortunately, there are cases where LCD gives us results whose intuitive acceptability is questionable.

Example 8. Consider the Quaker-Republican problem with the extra rule "Generally, ecologists are pacifist": $\Delta = \{q \to p, e \to p, q \to \neg p\}$. We have

$$m_1([\neg q \vee p]) = 1 - \varepsilon_1,\ \ m_1(\Omega) = \varepsilon_1,$$
$$m_2([\neg e \vee p]) = 1 - \varepsilon_2,\ \ m_2(\Omega) = \varepsilon_2,$$
$$m_3([\neg r \vee \neg p]) = 1 - \varepsilon_3,\ \ m_3(\Omega) = \varepsilon_3,$$

together with the constraints

$$\max\{pl_{\oplus}(\{\omega\}) | \omega \vDash q \wedge p\} >_{\infty} \max\{pl_{\oplus}(\{\omega\}) | \omega \vDash q \wedge \neg p\} \qquad \text{i.e., } 1 >_{\infty} \varepsilon_1$$
$$\max\{pl_{\oplus}(\{\omega\}) | \omega \vDash e \wedge p\} >_{\infty} \max\{pl_{\oplus}(\{\omega\}) | \omega \vDash e \wedge \neg p\} \qquad \text{i.e., } 1 >_{\infty} \varepsilon_2$$
$$\max\{pl_{\oplus}(\{\omega\}) | \omega \vDash r \wedge \neg p\} >_{\infty} \max\{pl_{\oplus}(\{\omega\}) | \omega \vDash r \wedge p\} \qquad \text{i.e., } 1 >_{\infty} \varepsilon_3$$

Since all the elements of ε are free, our algorithm puts all of them in the same class, i.e., $\varepsilon_1 = \varepsilon_2 = \varepsilon_3$. Then, we have

$$pl([q \wedge e \wedge r \wedge p]) \approx \varepsilon_3 >_{\infty} pl([q \wedge e \wedge r \wedge \neg p]) \approx \varepsilon_1 \varepsilon_2,$$

and so ecologists who are Quakers and republicans are pacifist: $q \wedge e \wedge r \mid\!\sim_{lcd} p$. ■

The last example may be disappointing, in that we may consider that we are in a case of ambiguity and we should stay silent. The reason for the answer given by LCD is to be found in the multi-source interpretation of the approach: having two sources to independently justify a conclusion is regarded as a stronger reason to accept that conclusion. The assumption of independence between the sources is essential to use Dempster's rule of combination, it may be interesting to study variants of LCD based on different rules of combination.

7. Conclusions

We have detailed a new approach to deal with default information based on a special class of belief functions, and have used it to define three non-monotonic consequence relations. The last one, LCD consequence, appears to be particularly attractive. We have proved that LCD is stronger than system **P**, and thus it satisfies the rationality postulates of Kraus, Lehmann and Magidor (1990). Moreover, we have given examples showing that LCD correctly addresses the well-known problems of irrelevance (example 4), of blocking of inheritance (example 5), of ambiguity (example 6), and of redundancy (example 7). In this, LCD has a distinctive advantage over all currently existing approaches.

It is interesting to notice that, despite this good behavior, LCD does not satisfy rational monotonicity. Moreover, although LCD is insensitive to the number of repetitions of the same default rule, it is sensitive to the number of different rules supporting the same conclusion. This is not surprising given that LCD is based on the interpretation of defaults as items of information provided by independent sources. We plan to explore variants of LCD that abandon the assumption of independence.

References

Adams, E. W. (1975) *The logic of conditionals*. Reidel, Dordrecht.

Benferhat, S., Dubois, D. and Prade, H. (1992) Representing default rules in possibilistic logic. Procs. of the 3rd Conf. on Principles of Knowledge Representation and Reasoning KR'92.

Benferhat, S., Cayrol, C., Dubois, D., Lang, J. and Prade, H. (1993) Inconsistency management and prioritized syntax-based entailment. Procs. of IJCAI'93.

Benferhat, S., Saffiotti, A. and Smets, Ph (1995) Belief functions and default reasoning. Technical Report TR/IRIDIA/95-5, Université Libre de Bruxelles, Belgium.

Dubois, D. and Prade, H. (1988) (with the collaboration of Farreny H., Martin- Clouaire R., Testemale C.) *Possibility Theory – An Approach to Computerized Processing of Uncertainty*. Plenum Press, New York

De Kleer, J. (1990) Using Crude Probability Estimates to Guide Diagnosis. *Artificial Intelligence*, 45, 381-391.

Gabbay, D. M. (1985) Theoretical foundations for non-monotonic reasoning in expert systems. In: K.R. Apt (Ed.) *Logics and models of Concurrent Systems*. Springer Verlag, Berlin, 439-457.

Goldszmidt, M. and Pearl, J. (1991) System Z^+: A formalism for reasoning with variable-strength defaults. Procs.. of the AAAI Conference, Anaheim, CA.

Hsia, Y-T. (1991) A belief-Function Semantics for Cautious Non-Monotonicity, Technical Report TR/IRIDIA/91-3, Université Libre de Bruxelles, Belgium.

Kraus, S., Lehmann, D. and Magidor, M. (1990) Non-monotonic reasoning, preferential models and cumulative logics. *Artificial Intelligence* **44**, 167-207.

Lehmann, D. (1993) Another perspective on Default Reasoning. Technical report, Hebrew University.

Lehmann, D. and Magidor, M. (1992) What does a conditional knowledge base entail? *Artificial Intelligence*, **55**, 1-60.

Pearl, J. (1988) *Probabilistic Reasoning in Intelligent Systems: Networks of Plausible Inference*. Morgan Kaufmann Pub. Inc., San Mateo, CA.

Pearl J. (1990) System Z: A natural ordering of defaults with tractable applications to default reasoning. Procs. of Theoretical Aspects of Reasoning about Knowledge, 121-135.

Poole, D. (1993) Average-case analysis of a search algorithm for estimating prior and posterior probabilities in Bayesian networks with extreme probabilities. Procs. of IJCAI'93, 606-612.

Shafer, G. (1976) *A Mathematical Theory of Evidence*. Princeton University Press, New Jersey.

Shoham, Y. (1988) *Reasoning About Change — Time and causation from the standpoint of Artificial Intelligence*. MIT Press, Cambridge, MA.

Smets, Ph. (1988) Belief functions. In: Ph. Smets, E. H. Mamdani, D. Dubois, H. Prade (Eds.) *Non-Standard Logics for Automated Reasoning* . Academic Press, New York, 253-286.

Smets, Ph and Hsia, Y-T. (1991) Default Reasoning and the Transferable Belief Model. In: P. Bonissone, M. Henrion, L. Kanal and J. Lemmer (Eds.) *Uncertainty in Artificial Intelligence 6*. North Holland, Amsterdam, 495-504.

Smets, Ph and Kennes, R. (1994) The transferable belief model. *Artificial Intelligence* **66**, 191-234.

Touretzky, D. (1984) Implicit ordering of defaults in inheritance systems. Procs. of AAAI'84, 322-325.

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