$\Leftrightarrow \quad pl(\alpha\wedge\delta\wedge\beta) >_\infty pl(\alpha\wedge\delta\wedge\neg\beta)$

thus proving the lemma. ■

Lemma 16. If $\alpha\wedge\text{Free}(\Delta\cup\{\alpha\})\vdash\beta$ then $\alpha \mid\!\sim_{lcd} \beta$.

Proof. First, we define $\alpha\wedge\text{Free}(\Delta\cup\{\alpha\})$ more precisely in terms of the so-called minimal inconsistent sub-bases. A sub-base A of Δ is said to be minimally inconsistent wrt α iff $\alpha\wedge\phi_A$ is inconsistent but for any $A'\subset A$ (strict inclusion) $\alpha\wedge\phi_{A'}$ is consistent. Then, we let $\text{Free}(\Delta\cup\{\alpha\})=\{d \,/\, \nexists A\subset\Delta,\ d\in A$ and A is a minimal inconsistent sub-base of $\Delta\}$. To prove our thesis, now, it is enough to prove that each lcd-bel-preferred model of α satisfies $\alpha\wedge\text{Free}(\Delta\cup\{\alpha\})$. We reason by contradiction, and assume that we have an interpretation ω which is a lcd-bel-preferred model of α but falsifies $\alpha\wedge\text{Free}(\Delta\cup\{\alpha\})$. Let $\Delta_\omega=\alpha\wedge(\wedge_{\omega\models d \text{ and } d\in\Delta}\ \phi_d)$. Since formulas of $\text{Free}(\Delta\cup\{\alpha\})$ are free from inconsistencies, then $\Delta_\omega\wedge\text{Free}(\Delta\cup\{\alpha\})$ is consistent. That is, there exists an interpretation ω' that satisfies $\Delta_\omega\wedge\text{Free}(\Delta\cup\{\alpha\})$. By using Lemma 7 we can easily check that this ω' is bel-preferred to ω, thus contradicting our assumption. ■

2) for all $d_i' \in \Delta_R$, $\varepsilon_{d_i'} >_\infty \varepsilon_{\mathbf{d}}$, and

3) for all $d \in \Delta$, $\prod_{d':\varepsilon_{d'} >_\infty \varepsilon_d} \varepsilon_{d'} >_\infty \varepsilon_d$.

The added constraints (i.e., condition 2) is consistent with the ones induced by condition 1 since none of d_i' of Δ_R satisfies $d_i' > \mathbf{d}$. From these conditions we can easily check that $pl_{\mathcal{E}}(\omega') \geq_\infty pl_{\mathcal{E}}(\omega)$. Indeed,

$$pl_{\mathcal{E}}(\omega') \approx_\infty \prod \{\varepsilon_i \mid d_i \in \Delta_R\} \,.\, \prod \{\varepsilon_i \mid d_i \in \Delta_C\}$$
$$>_\infty \varepsilon_{\mathbf{d}} \,.\, \prod \{\varepsilon_i \mid d_i \in \Delta_C\} \text{ (using condition 2) and 3) above and Property A8(c))}$$
$$>_\infty \prod \{\varepsilon_i \mid d_i \in \Delta_R\} \,.\, \prod \{\varepsilon_i \mid d_i \in \Delta_C\}$$ ■

Lemma 15. Let $\delta \in \mathcal{L}_{out(\Delta)}$ and $\alpha,\beta \in \mathcal{L}_{in(\Delta)}$. Then $\alpha \mid\!\sim_{lcd} \beta$ implies $\alpha\wedge\delta \mid\!\sim_{lcd} \beta$.

<u>Proof.</u> We recall that In(Δ) denotes the set of propositional symbols which appear in Δ and Out(Δ) = V - In(Δ) denotes the set of propositional symbols which do not appear in Δ, where V denotes the set of all propositional symbols of the language. Hence, each interpretation ω in Ω based on V can be seen as a pair of conjuncts $\omega = x \wedge y$, where x is an interpretation based only on In(Δ), and y is an interpretation based only on Out(Δ). Let $\Omega_{in(\Delta)}$ be the set of all interpretations constructed from In(Δ) and $\Omega_{Out(\Delta)}$ be the set of all interpretations constructed from Out(Δ).

It is easy to check that for any $x \in \Omega_{in(\Delta)}$ and $y \in \Omega_{Out(\Delta)}$, x and $x\wedge y$ falsify the same set of defaults in Δ. Indeed, if x falsifies d, then $x\wedge\phi_d$ is inconsistent hence $x\wedge y\wedge\phi_d$ is also inconsistent. The converse is also true since $x\wedge\phi_d$ and y are formulas built on two disjoint sets of propositional symbols (i.e., In(Δ), Out(Δ)).

Therefore, for each bel an element of Bel_{lcd}, we have

$$pl(x) = pl(x\wedge y) \text{ for any } x \in \Omega_{in(\Delta)} \text{ and } y \in \Omega_{Out(\Delta)}$$

As a corollary, for any $x \in \Omega_{in(\Delta)}$ and for any $A \subseteq \Omega_{Out(\Delta)}$ we have:

$$pl(x) = \max \{ pl(x\wedge y_i) \ / \text{ for } y_i \in A\}. \qquad (\S)$$

The proof now is trivial using the previous equalities. Indeed, for each bel an element of Bel_{lcd}, we have:

$pl(\alpha\wedge\beta) >_\infty pl(\alpha\wedge\neg\beta)$

$\Leftrightarrow \max_{\omega\in\Omega}\{pl(\omega) \mid \omega\vDash\alpha\wedge\beta\} >_\infty \max_{\omega'\in\Omega} \{pl(\omega') \mid \omega'\vDash\alpha\wedge\neg\beta\}$

$\Leftrightarrow \max_{\omega=x\wedge y}\{pl(x\wedge y) \mid x\vDash\alpha\wedge\beta\} >_\infty \max_{\omega=x'\wedge y'} \{pl(x'\wedge y') \mid x'\vDash\alpha\wedge\neg\beta\}$

$\Leftrightarrow \max_{\omega=x\wedge y}\{pl(x) \mid x\vDash\alpha\wedge\beta\} >_\infty \max_{\omega=x'\wedge y'} \{pl(x') \mid x'\vDash\alpha\wedge\neg\beta\}$

(using § with $A = \Omega_{Out(\Delta)}$ since any $y \in \Omega_{Out(\Delta)}$, $y\wedge\alpha\wedge\beta$ is consistent)

$\Leftrightarrow \max_{\omega=x\wedge y}\{pl(x\wedge y) \mid x\vDash\alpha\wedge\beta, y\vDash\delta\} >_\infty \max_{\omega=x'\wedge y'} \{pl(x'\wedge y') \mid x'\vDash\alpha\wedge\neg\beta, y'\vDash\delta\}$

(using § with A = models of δ)

$\Leftrightarrow \max_{\omega\in\Omega}\{pl(\omega) \mid \omega\vDash\alpha\wedge\delta\wedge\beta\} >_\infty \max_{\omega'\in\Omega} \{pl(\omega') \mid \omega'\vDash\alpha\wedge\delta\wedge\neg\beta\}$

ii) $\prod_{k=1}^{i-1} \prod_{h=1}^{|\Delta_k|} \varepsilon_{kh} \;.\; \prod_{d_{ji} \in \Delta(\omega', i)} \varepsilon_{ji} \;>_\infty \varepsilon_{si}$ for all $d_{si} \notin \Delta(\omega', i)$

Then using exactly the proof of Lemma 12 (2b), we get $pl_{\mathcal{E}}(\omega') >_\infty pl_{\mathcal{E}}(\omega)$. ■

Lemma 13. Let $>_\Delta$ be an admissible preference relation on Δ, and $(\Omega, >_\Omega, \Delta, >_\Delta)$ be an admissible prioritized structure. Then, for any ω and ω' in Ω, $\omega >_\Omega \omega'$ if and only if ω is $bel_{\mathcal{E}}$-preferred to ω' for any $bel_{\mathcal{E}}$ in $Bel_\oplus(\Delta, >_\Delta)$.

Proof. Recall that, for any ω in Ω, we have $pl_{\mathcal{E}}(\omega) \approx_\infty t_\omega$, where t_ω is given by (8, section 5.1). We denote by Δ_ω the set of defaults in Δ such that ω falsifies d. For any ω and ω' in Ω, let Δ_L, Δ_C and Δ_R denote $\Delta_\omega - \Delta_{\omega'}$, $\Delta_\omega \cap \Delta_{\omega'}$ and $\Delta_{\omega'} - \Delta_\omega$, respectively.

We have:

$pl_{\mathcal{E}}(\omega) \approx_\infty \prod \{\varepsilon_i \mid d_i \in \Delta_L\} \;.\; \prod \{\varepsilon_i \mid d_i \in \Delta_C\}$

$pl_{\mathcal{E}}(\omega') \approx_\infty \prod \{\varepsilon_j \mid \varepsilon_j \in \Delta_R\} \;.\; \prod \{\varepsilon_j \mid d_j \in \Delta_C\}$.

[⇒] Assume that $\omega >_\Omega \omega'$. As $(\Omega, >_\Omega, \Delta, >_\Delta)$ is an admissible prioritized structure, then, $\omega >_\Omega \omega'$ iff for each d_i in Δ_L there exits d_j in Δ_R such that $d_j >_\Delta d_i$, that is, $\varepsilon_i >_\infty \varepsilon_j$ or equivalently, $\kappa(\varepsilon_i) < \kappa(\varepsilon_j)$.

Let $\{\Delta_{L1},\ldots,\Delta_{Lk}\}$ be a partition of Δ_L such that Δ_{Li} contains all the ε_{ij} having the same ε_i' in Δ_R, namely $\varepsilon_i' <_\infty \varepsilon_{ij}$ for $j=1,|\Delta_{Li}|$. Using the constraints on the infinitesimals used to recover Geffner's system, we also have for any i:

$\prod \{\varepsilon_{ij} \mid \varepsilon_{ij} \in \Delta_i\} >_\infty \varepsilon_i'$.

Therefore:

$pl_{\mathcal{E}}(\omega) \approx_\infty \prod \{\varepsilon_i \mid d_i \in \Delta_L\} \;.\; \prod \{\varepsilon_i \mid d_i \in \Delta_C\}$

⇔ $pl_{\mathcal{E}}(\omega) \approx_\infty \prod \{\varepsilon_{1j} \mid \varepsilon_{1j} \in \Delta_{L1}\} \;.\; \ldots \;.\; \prod \{\varepsilon_{kj} \mid \varepsilon_{kj} \in \Delta_{Lk}\} \;.\; \prod \{\varepsilon_i \mid d_i \in \Delta_C\}$

⇒ $pl_{\mathcal{E}}(\omega) >_\infty \varepsilon_1' \;.\; \ldots \;.\; \varepsilon_k' \;.\; \prod \{\varepsilon_i \mid d_i \in \Delta_C\}$ (Using A9(f))

⇒ $pl_{\mathcal{E}}(\omega) >_\infty \prod \{\varepsilon_i' \mid \varepsilon_i' \in \Delta_R\} \;.\; \prod \{\varepsilon_i \mid d_i \in \Delta_C\}$

⇒ $pl_{\mathcal{E}}(\omega) >_\infty pl_{\mathcal{E}}(\omega')$.

[⇐] Assume now that $pl_{\mathcal{E}}(\omega) >_\infty pl_{\mathcal{E}}(\omega')$ holds for any $bel_{\mathcal{E}}$ in $Bel_\oplus(\Delta, >_\Delta)$, and suppose that ω is not preferred (in the sense of Geffner) to ω'. This means that there exists **d** of Δ_L such that none of d_i' of Δ_R satisfies $d_i' >$ **d**. We proceed as in Lemma 12 by constructing a $bel_{\mathcal{E}}$ which is in $Bel_\oplus(\Delta, >_\Delta)$ and such that $pl_{\mathcal{E}}(\omega') \geq_\infty pl_{\mathcal{E}}(\omega)$. $bel_{\mathcal{E}}$ is constructed such that it satisfies the three following conditions:

1) for all $d, d' \in \Delta$, if $d >_\Delta d'$ then $\varepsilon_{d'} >_\infty \varepsilon_d$, and

$$pl_{\mathcal{E}}(\omega') \;>_\infty\; \prod_{d_{ji}\in \Delta(\omega,i)-\Delta(\omega',i)} \varepsilon_{ji} \cdot \prod_{d_{ji}\in \Delta(\omega',i)\cap\Delta(\omega,i)} \varepsilon_{ji} \cdot \prod_{k=i+1}^{n} \prod_{d_{jk}\in \Delta(\omega,k)} \varepsilon_{jk}$$

$$>_\infty\; \prod_{d_{ji}\in \Delta(\omega,i)} \varepsilon_{ji} \cdot \prod_{k=i+1}^{n} \prod_{d_{jk}\in \Delta(\omega,k)} \varepsilon_{jk}$$

$$>_\infty\; \prod_{k=1}^{i-1} \prod_{d_{jk}\in \Delta(\omega',k)} \varepsilon_{jk} \cdot \prod_{d_{ji}\in \Delta(\omega,i)} \varepsilon_{ji} \cdot \prod_{k=i+1}^{n} \prod_{d_{jk}\in \Delta(\omega,k)} \varepsilon_{jk}$$

$$\Rightarrow\; pl_{\mathcal{E}}(\omega') \;>_\infty\; pl_{\mathcal{E}}(\omega)$$

and hence the thesis. ■

Theorem 5. Let $\mid\!\sim_B$ be the inference relation of Brewka's preferred sub-theories system. Then, for any given Δ, $\alpha \mid\!\sim_B \beta$ if, and only if, $\alpha \mid\!\sim_{\oplus 3} \beta$.

Proof.

• ($\Rightarrow$) $\alpha \mid\!\sim_B \beta$

$\Leftrightarrow$ $\forall\omega\in[\alpha\wedge\neg\beta]$, $\exists\omega'\in[\alpha\wedge\beta]$ such that ω' is B-preferred to ω

$\Rightarrow$ $\forall\omega\in[\alpha\wedge\neg\beta]$, $\exists\omega'\in[\alpha\wedge\neg\beta]$ such that $pl_{\mathcal{E}}(\omega') >_\infty pl_{\mathcal{E}}(\omega)$ in each $bel_{\mathcal{E}}$ of $Bel_{\oplus 3}(\Delta)$ (due to Lemma 12)

$\Rightarrow$ $\alpha \mid\!\sim_{\oplus 3} \beta$.

• ($\Leftarrow$) Assume that $\alpha \mid\!\sim_{\oplus 3} \beta$ holds but $\alpha \mid\!\sim_B \beta$ does not hold, this means that $\exists\omega'\in[\alpha\wedge\neg\beta]$ such that $\forall\omega\in[\alpha\wedge\beta]$, ω is not B-preferred to ω'. Again, we can easily construct a $bel_{\mathcal{E}}$ of $Bel_{\oplus 3}(\Delta)$ such that $pl_{\mathcal{E}}(\omega') \geq_\infty pl_{\mathcal{E}}(\omega)$ holds for any $\omega\in[\alpha\wedge\beta]$ (and hence contradicts the hypothesis). Indeed, let $\omega\in[\alpha\wedge\beta]$, we proceeds as in Lemma 12 by distinguishing two cases:

1) ω' is B-preferred to ω, then, by the if part of the proof of Lemma 12, $pl_{\mathcal{E}}(\omega') >_\infty pl_{\mathcal{E}}(\omega)$, or
2) ω' is not *B-preferred* than ω', then, either

2a) for any j, $[\omega]_j = [\omega']_j$, but this means that $pl_{\mathcal{E}}(\omega') = pl_{\mathcal{E}}(\omega)$ since they exactly falsify the same set of defaults, or

2b) there exists an index i such that, $\forall j>i$, $[\omega]_j=[\omega']_j$ and neither $[\omega]_i \supset[\omega']_i$ nor $[\omega']_i\supset[\omega]_i$. In this case, we define $bel_{\mathcal{E}}$ as one satisfying the two following conditions (given by Lemma 12 (2b)): for i=2,n:

i) $\displaystyle\prod_{k=1}^{i-1} \prod_{h=1}^{|\Delta_k|} \varepsilon_{kh} >_\infty \varepsilon_{si}$ for all s = 1, ..., $|\Delta_i|$,

ii) $\prod_{k=1}^{i-1} \prod_{h=1}^{|\Delta_k|} \varepsilon_{kh} \cdot \prod_{d_{ji} \in \Delta(\omega', i)} \varepsilon_{ji} >_\infty \varepsilon_{si}$ for all $d_{si} \notin \Delta(\omega', i)$

These constraints are obtained by simply assuming a refinement of the stratification of Δ where each Δ_i is in fact split in two parts where defaults falsified by ω' are less prioritary than those satisfied by ω'. Clearly, condition i guarantees that $bel_{\mathcal{E}}$ belongs to $Bel_{\oplus 3}(\Delta)$, and condition ii) makes explicit the splitting of each Δ_i. Hence, we get:

$$pl_{\mathcal{E}}(\omega') \approx_\infty \prod_{k=1}^{n} \prod_{d_{jk} \in \Delta(\omega',k)} \varepsilon_{jk}$$

$$= \prod_{k=1}^{i-1} \prod_{d_{jk} \in \Delta(\omega',k)} \varepsilon_{jk} \cdot \prod_{d_{ji} \in \Delta(\omega',i)} \varepsilon_{ji} \cdot \prod_{k=i+1}^{n} \prod_{d_{jk} \in \Delta(\omega',k)} \varepsilon_{jk}$$

$$= \prod_{k=1}^{i-1} \prod_{d_{jk} \in \Delta(\omega',k)} \varepsilon_{jk} \cdot \prod_{d_{ji} \in \Delta(\omega',i)} \varepsilon_{ji} \cdot \prod_{k=i+1}^{n} \prod_{d_{jk} \in \Delta(\omega,k)} \varepsilon_{jk}$$

(since for k=i+1,n we have Δ(ω',k)=Δ(ω,k))

$$= \prod_{k=1}^{i-1} \prod_{d_{jk} \in \Delta(\omega',k)} \varepsilon_{jk} \cdot \prod_{d_{ji} \in \Delta(\omega',i)-\Delta(\omega,i)} \varepsilon_{ji} \cdot \prod_{d_{ji} \in \Delta(\omega',i) \cap \Delta(\omega,i)} \varepsilon_{ji} \cdot \prod_{k=i+1}^{n} \prod_{d_{jk} \in \Delta(\omega,k)} \varepsilon_{jk} \qquad (\S)$$

Now from the above condition ii), we have:

$$\prod_{k=1}^{i-1} \prod_{h=1}^{|\Delta_k|} \varepsilon_{kh} \cdot \prod_{d_{ji} \in \Delta(\omega', i)} \varepsilon_{ji} >_\infty \varepsilon_{si} \quad \text{for all } d_{si} \notin \Delta(\omega', i)$$

$$\Rightarrow \prod_{k=1}^{i-1} \prod_{h=1}^{|\Delta_k|} \varepsilon_{kh} \cdot \prod_{d_{ji} \in \Delta(\omega', i)} \varepsilon_{ji} >_\infty \varepsilon_{si} \quad \text{for all } d_{si} \in \Delta(\omega, i) - \Delta(\omega', i)$$

(Δ(ω, i) - Δ(ω', i) is not empty)

$$\Rightarrow \prod_{k=1}^{i-1} \prod_{d_{jk} \in \Delta(\omega',k)} \varepsilon_{jk} \cdot \prod_{d_{ji} \in \Delta(\omega', i)} \varepsilon_{ji} >_\infty \varepsilon_{si} \quad \text{for all } d_{si} \in \Delta(\omega, i) - \Delta(\omega', i)$$

(using the property A8(c) since $\Delta(\omega,k) \subset \Delta_k$))

$$\Rightarrow \prod_{k=1}^{i-1} \prod_{d_{jk} \in \Delta(\omega',k)} \varepsilon_{jk} \cdot \prod_{d_{ji} \in \Delta(\omega',i)-\Delta(\omega,i)} \varepsilon_{ji} >_\infty \varepsilon_{si} \quad \text{for all } d_{si} \in \Delta(\omega, i) - \Delta(\omega', i)$$

(using the property A8(c) since $\Delta(\omega', i) - \Delta(\omega, i) \subset \Delta(\omega', i)$)

$$\Rightarrow \prod_{k=1}^{i-1} \prod_{d_{jk} \in \Delta(\omega',k)} \varepsilon_{jk} \cdot \prod_{d_{ji} \in \Delta(\omega',i)-\Delta(\omega,i)} \varepsilon_{ji} >_\infty \prod_{d_{ji} \in \Delta(\omega,i)-\Delta(\omega',i)} \varepsilon_{ji}$$

(using the property A8(c))

Hence, replacing in (§) and using the property A8(c) we get:

(using the property A8(c) since $\Delta(\omega,k)\subset\Delta_k$)

$$\Rightarrow \prod_{k=1}^{i-1} \prod_{d_{jk}\in\Delta(\omega,k)} \varepsilon_{jk} >_\infty \prod_{d_{ji}\in\Delta(\omega',i)-\Delta(\omega,i)} \varepsilon_{ji}$$

(using again the property A8(c))

Hence, (*) becomes:

$$pl_{\mathcal{E}}(\omega) >_\infty \prod_{d_{ji}\in\Delta(\omega',i)} \varepsilon_{ji} \cdot \prod_{k=i+1}^{n} \prod_{d_{jk}\in\Delta(\omega',k)} \varepsilon_{jk}$$

(since $\prod_{d_{ji}\in\Delta(\omega',i)-\Delta(\omega,i)} \varepsilon_{ji}\cdot\prod_{d_{ji}\in\Delta(\omega,i)} \varepsilon_{ji} = \prod_{d_{ji}\in\Delta(\omega',i)} \varepsilon_{ji}$, and recall that $\Delta(\omega,k)\subset\Delta(\omega',k)$)

$$\Rightarrow pl_{\mathcal{E}}(\omega) >_\infty \prod_{k=1}^{i-1} \prod_{d_{jk}\in\Delta(\omega',k)} \varepsilon_{jk} \cdot \prod_{d_{ji}\in\Delta(\omega',i)} \varepsilon_{ji} \cdot \prod_{k=i+1}^{n} \prod_{d_{jk}\in\Delta(\omega',k)} \varepsilon_{jk}$$

$$\Rightarrow pl_{\mathcal{E}}(\omega) >_\infty pl_{\mathcal{E}}(\omega')$$

and thence the thesis.

($\Leftarrow$) To see that the "only if" part holds, assume that $pl_{\mathcal{E}}(\omega) >_\infty pl_{\mathcal{E}}(\omega')$ holds for each $bel_{\mathcal{E}}$ of $Bel_{\oplus 3}(\Delta)$ and suppose, by way of contradiction, that ω is *not* B-preferred to ω'. We distinguish two cases:

1) ω' is B-preferred to ω, then, by the if part of the proof, $pl_{\mathcal{E}}(\omega') >_\infty pl_{\mathcal{E}}(\omega)$, for each $bel_{\mathcal{E}}$ of $Bel_{\oplus 3}(\Delta)$ which contradicts our hypothesis, or

2) ω' is not *B-preferred* than ω', then, either

2a) for any j, $[\omega]_j = [\omega']_j$, but this means that $pl_{\mathcal{E}}(\omega') = pl_{\mathcal{E}}(\omega)$ since they exactly falsify the same set of defaults, which contradicts our hypothesis, or

2b) there exists an index i such that, $\forall j>i$, $[\omega]_j=[\omega']_j$ and neither $[\omega]_i \supset[\omega']_i$ nor $[\omega']_i\supset[\omega]_i$.

In this case, we can easily construct a $bel_{\mathcal{E}}$ which belongs to $Bel_{\oplus 3}(\Delta)$ and such that $pl_{\mathcal{E}}(\omega') >_\infty pl_{\mathcal{E}}(\omega)$ (and hence contradicts the hypothesis). $bel_{\mathcal{E}}$ should satisfy the two following conditions:

for i=2,n:

i) $\displaystyle\prod_{k=1}^{i-1} \prod_{h=1}^{|\Delta_k|} \varepsilon_{kh} >_\infty \varepsilon_{si}$ for all s = 1, ..., $|\Delta_i|$, Hence, ω must be lex-preferred to ω', and our proof is concluded. ■

Lemma 12. Let ω and ω' be elements of Ω. Then, ω is B-preferred to ω' if and only if ω is $bel_{\mathcal{E}}$-preferred to ω' in each $bel_{\mathcal{E}}$ of $Bel_{\oplus 3}(\Delta)$.

Proof. The proof is similar to the one of Lemma 11. We again start by recalling to preference ordering used in Brewka's system. Let Δ be a stratified base as above. An interpretation ω is said to be *B-preferred* to ω' (with respect to Δ) if and only if there exists a layer index i such that:

1) if $i<n$, $[\omega']_i \subset [\omega]_i$, and
2) $\forall n \geq j > i$, $[\omega]_j = [\omega']_j$.

where $[\omega]_i$ denotes the set of defaults of Δ_i satisfied by ω.

($\Rightarrow$) The proof goes as the corresponding one in Lemma 11. By our stipulations:

$$pl_{\mathcal{E}}(\omega) \approx_\infty \prod \{\varepsilon_d \mid d \in \Delta \text{ s.t. } \omega \not\models \phi_d\}$$
$$= \prod_{k=1}^{n} \prod_{d_{jk} \in \Delta(\omega,k)} \varepsilon_{jk}$$

where $\Delta(\omega, k)$ is the set of defaults in Δ_k that are falsified by ω. Now, assume that ω is *B-preferred* than ω', and let i such that $\forall j>i$, $[\omega]_j=[\omega']_j$, and $[\omega]_i \supset [\omega']_i$ (i.e., $\forall j>i$, $\Delta(\omega, j) = \Delta(\omega', j)$, and $\Delta(\omega, i) \subset \Delta(\omega', i)$). Then, each $bel_{\mathcal{E}}$ of $Bel_{\oplus 3}(\Delta)$ we have:

$$pl_{\mathcal{E}}(\omega) \approx_\infty \prod_{k=1}^{n} \prod_{d_{jk} \in \Delta(\omega,k)} \varepsilon_{jk}$$
$$= \prod_{k=1}^{i-1} \prod_{d_{jk} \in \Delta(\omega,k)} \varepsilon_{jk} \cdot \prod_{d_{ji} \in \Delta(\omega,i)} \varepsilon_{ji} \cdot \prod_{k=i+1}^{n} \prod_{d_{jk} \in \Delta(\omega,k)} \varepsilon_{jk}$$
$$= \prod_{k=1}^{i-1} \prod_{d_{jk} \in \Delta(\omega,k)} \varepsilon_{jk} \cdot \prod_{d_{ji} \in \Delta(\omega,i)} \varepsilon_{ji} \cdot \prod_{k=i+1}^{n} \prod_{d_{jk} \in \Delta(\omega',k)} \varepsilon_{jk}$$

(since for k=i+1,n we have $\Delta(\omega',k)=\Delta(\omega,k)$)

$$>_\infty \prod_{d_{ji} \in \Delta(\omega',i)-\Delta(\omega,i)} \varepsilon_{ji} \cdot \prod_{d_{ji} \in \Delta(\omega,i)} \varepsilon_{ji} \cdot \prod_{k=i+1}^{n} \prod_{d_{jk} \in \Delta(\omega',k)} \varepsilon_{jk} \qquad (*)$$

Indeed, from the constraints on ε_{ij}, we have:

$$\prod_{k=1,i-1} \prod_{j=1,|\Delta_k|} \varepsilon_{kj} >_\infty \varepsilon_{is} \text{ for any } d_{is} \in \Delta_i$$
$$\Rightarrow \prod_{k=1}^{i-1} \prod_{d_{jk} \in \Delta(\omega,k)} \varepsilon_{jk} >_\infty \varepsilon_{ji} \text{ (where } d_{ij} \in \Delta(\omega',i)-\Delta(\omega,i) \text{ which is not empty)}$$

(⇒) Assume that ω is lex-preferred to ω', and let i such that $\forall j>i$, $|[\omega]_j|=|[\omega']_j|$, and $|[\omega]_i|>|[\omega']_i|$. Then, using the convention that product taken on empty sets of indices is 1,

$$pl_{\mathcal{E}}(\omega) \approx_\infty \prod_{k=1}^{n} \delta_k^{f(\omega,k)}$$

$$= \prod_{k=1}^{i-1} \delta_k^{f(\omega,k)} \;.\; \delta_i^{f(\omega;i)} \;.\; \prod_{j=i+1}^{n} \delta_j^{f(\omega,j)}$$

$$= \prod_{k=1}^{i-1} \delta_k^{f(\omega,k)} \;.\; \delta_i^{f(\omega;i)} \;.\; \prod_{j=i+1}^{n} \delta_j^{f(\omega',j)} \quad \text{(since } f(\omega',j)=f(\omega,j)) \qquad (*)$$

From the constraints on δ_i, we have:

$$\prod_{k=1}^{i-1} \delta_k^{|\Delta_k|} \;>_\infty\; \delta_i$$

$$\Rightarrow \prod_{k=1}^{i-1} \delta_k^{f(\omega,k)} \;>_\infty \delta_i \qquad \text{(since } f(\omega,k)\le|\Delta_k| \text{ and } \delta\in(0,1))$$

$$\Rightarrow \prod_{k=1}^{i-1} \delta_k^{f(\omega,k)} \;>_\infty \delta_i^{f(\omega',i)-f(\omega,i)} \qquad \text{(where } f(\omega',i)-f(\omega,i) \text{ is positive)}$$

both ⇒ resulting from Lemma A8(c) and the transitivity of $>_\infty$. Hence, (*) becomes:

$$pl_{\mathcal{E}}(\omega) >_\infty \delta_i^{f(\omega';i)-f(\omega;i)} \;.\; \delta_i^{f(\omega;i)} \;.\; \prod_{j=i+1}^{n} \delta_j^{f(\omega';j)}$$

$$\Rightarrow \quad pl_{\mathcal{E}}(\omega) >_\infty \delta_i^{f(\omega',i)} \;.\; \prod_{j=i+1}^{n} \delta_j^{f(\omega',j)}$$

$$\Rightarrow \quad pl_{\mathcal{E}}(\omega) >_\infty \prod_{k=1}^{i-1} \delta_k^{f(\omega',k)} \;.\; \delta_i^{f(\omega',i)} \;.\; \prod_{j=i+1}^{n} \delta_j^{f(\omega',j)}$$

$$\Rightarrow \quad pl_{\mathcal{E}}(\omega) >_\infty pl_{\mathcal{E}}(\omega')$$

and thence the thesis.

(⇐) The reverse implication can be shown to hold in a similar way. Assume that $pl_{\mathcal{E}}(\omega) >_\infty pl_{\mathcal{E}}(\omega')$ and suppose, by way of contradiction, that ω is *not* lex-preferred to ω'. We distinguish two cases:

1) ω' is neither lex-preferred to ω, and hence ω and ω' falsifies exactly the same number of rules in each Δ_i; but this means that $pl_{\mathcal{E}}(\omega) = pl_{\mathcal{E}}(\omega')$, which contradicts our hypothesis; or
2) ω' is lex-preferred to ω; but then, using if part of the proof, $pl_{\mathcal{E}}(\omega') >_\infty pl_{\mathcal{E}}(\omega)$, which again contradicts the hypothesis.

$$C(\omega) \;=\; \sum\{c(d) \mid d\in\Delta \text{ s.t. } \omega\not\models\phi_d\} \;=\sum_{i=1}^{n} i.k_i(\omega)$$

where $k_i(\omega)$ is the number of defaults of Δ_i which are not satisfied by the world ω and $n = |\Delta|$. We say that ω is penalty-preferred to ω' iff $C(\omega) < C(\omega')$. For a given formula α, we say that an interpretation ω is α-penalty-preferred if it satisfies α and there is no interpretation ω' satisfying α such that $C(\omega') < C(\omega)$. Finally, β is a penalty-consequence of α and Δ iff each α-penalty-preferred interpretation satisfies β. Turning now attention to our system, we have $\varepsilon_d = \delta^i$ for $d\in\Delta_i$. It is easy to see that, for each world ω,

$$\begin{aligned} pl_{\mathcal{E}}(\omega) &\approx_\infty \prod\{\varepsilon_d \mid d\in\Delta \text{ s.t. } \omega\not\models\phi_d\} \\ &= \prod_{i=1,n} \delta^{i\cdot k_i(\omega)} \\ &= \delta^{\sum_{i=1,n} i\cdot k_i(\omega)}. \end{aligned}$$

Then, we have $pl_{\mathcal{E}}(\omega) \approx_\infty \delta^{C(\omega)}$. Hence $pl_{\mathcal{E}}(\omega) >_\infty pl_{\mathcal{E}}(\omega')$ iff $C(\omega) < C(\omega')$ iff ω is penalty-preferred to ω', hence the thesis. ■

Lemma 11. Let ω and ω' be elements of Ω, and let $bel_{\mathcal{E}}$ be any element of $Bel_{\oplus 2}(\Delta)$. Then, ω is lex-preferred to ω' if and only if ω is $bel_{\mathcal{E}}$-preferred to ω'.

Proof. We start by recalling the definition of lex-preference (see (Benferhat et al., 1993) for more details). Let Δ be the given base, and $\{\Delta_1,\ldots, \Delta_n\}$ be the given stratification for it. Then, an interpretation ω is said to be *lex-preferred* than ω' if and only if there exist an index $1\leq i\leq n$ such that:

1) $|[\omega]_i|>|[\omega']_i|$, and
2) if $i<n$, $\forall n\geq j>i$, $|[\omega]_j|=|[\omega']_j|$,

where $|[\omega]_i|$ is the number of rules in Δ_i satisfied by ω. A non-monotonic consequence relation, called a lexicographical entailment and denoted by $\mid\!\sim_{lex}$, is defined from this ordering in the same way as in penalty logic.

Note that:

$$pl_{\mathcal{E}}(\omega) \approx_\infty \prod\{\varepsilon_d \mid d\in\Delta \text{ s.t. } \omega\not\models\phi_d\},$$

and by our stipulations

$$\prod\{\varepsilon_d \mid d\in\Delta \text{ s.t. } \omega\not\models\phi_d\} = \prod_{k=1}^{n} \delta_k^{f(\omega;k)}$$

where $f(\omega, k)$ is the number of defaults in Δ_k that are falsified by ω.

where each F is a focal element of $m_\oplus$. Let's analyze these focal elements. By construction, each F is the intersection of the focal elements of some of the m_d's, i.e., $F = [\phi_{\Delta'}]$ for some subset Δ' of Δ. It is easy to check that

$$m_\oplus(F) = \prod_{d\in\Delta'} (1-\varepsilon_d) \cdot \prod_{d\notin\Delta'} \varepsilon_d \quad . \qquad (\S)$$

Clearly, each factor in the first product is in $\mathbb{E}^1$; then, by repeatedly applying Lemma A8(d),

$$m_\oplus(F) \approx_\infty \prod_{d\notin\Delta'} \varepsilon_d \quad .$$

With κ denoting the order (see definition A1), we have:

$$\begin{aligned}
\kappa(pl_\oplus(\omega)) &= \min_{F:\ \omega\in F} \kappa(m_\oplus(F)) && \text{by Lemma 3 and A2(b),}\\
&= \min_{F:\ \omega\in F} \kappa\Big(\prod_{d\notin\Delta'} \varepsilon_d\Big) \\
&= \min_{F:\ \omega\in F} \sum_{d\notin\Delta'} \kappa(\varepsilon_d)\ , && \text{by Lemma A2(c).}
\end{aligned}$$

As each $\kappa(\varepsilon_d)$ is positive and finite, the minimum is achieved when Δ' is the 'largest' possible subbase of Δ, hence when $\Delta' = \Delta_\omega = \{d: \omega\models\phi_d, d\in\Delta\}$, any other Δ'' that satisfies $\omega\models\phi_{\Delta''}$ being a subset of Δ_ω, in which case the sum would be larger.

We distinguish two cases.

(i) $\Delta_\omega = \Delta$, that is, ω satisfies all the defaults in Δ. (Note that $[\phi_\Delta]$ is the only focal element for which this may be the case.) Then, the second product in (§) is empty, and so $m_\oplus(F_\omega) \approx_\infty 1$, as stated in our thesis.

(ii) $\Delta_\omega \neq \Delta$. Then:

$$\begin{aligned}
\kappa(pl_\oplus(\omega)) &= \sum_{d\notin\Delta_\omega} \kappa(\varepsilon_d)\\
&= \sum_{\omega\not\models\phi_d} \kappa(\varepsilon_d)\\
&= \kappa\Big(\prod\{\varepsilon_d \mid d\in\Delta \text{ s.t. } \omega\not\models\phi_d\}\Big).
\end{aligned}$$

■

Lemma 10. Let ω and ω' be elements of Ω, and let $bel_{\mathcal{E}}$ be any element of $Bel_{\oplus 1}(\Delta)$. Then, ω is penalty-preferred to ω' if and only if ω is $bel_{\mathcal{E}}$-preferred to ω'.

<u>Proof</u>. We first recall the definition of preference used in penalty logic. Let $\{\Delta_1,\ldots, \Delta_n\}$ be a stratification of Δ. Each default (or, more generally, piece of information) d in Δ_i is associated with the penalty $c(d) = i$, read as the price to pay if d is not satisfied. These penalties induce a complete order on the elements of Ω based on the cost

Lemma 8. Let $\{\Delta_1, \Delta_2, \ldots, \Delta_k\}$ be the stratification given by system **Z**, and let bel_i the εbf built by the LC algorithm at step i. Then, for any default $\alpha\rightarrow\beta$ in Δ,

a) $\alpha\rightarrow\beta$ is tolerated by Δ iff $pl_1(\alpha) = 1$

b) $\alpha\rightarrow\beta$ is not tolerated by $\Delta_i\cup\ldots\Delta_k$ iff $pl_i(\alpha\wedge\beta) = pl_i(\alpha\wedge\neg\beta) = \varepsilon_i$.

c) $\alpha\rightarrow\beta \in \Delta_i$ implies $bel_i \models \alpha\rightarrow\beta$.

d) $bel_i \models \alpha\rightarrow\beta$ and $pl_{i-1}(\alpha\wedge\beta) = pl_{i-1}(\alpha\wedge\neg\beta) = \varepsilon_{i-1}$ (i.e., $bel_{i-1} \not\models \alpha\rightarrow\beta$) implies that $\alpha\rightarrow\beta \in \Delta_i$.

Proof.
(a) We first notice that the rule $\alpha\rightarrow\beta \in \Delta$ is tolerated by Δ iff $\alpha\wedge\phi_\Delta$ is consistent (Benferhat et al., 1992). Recall that $m_1(\Omega) = \varepsilon_1$ and $m_1(\phi_\Delta) = 1-\varepsilon_1$; then, by definition of pl, we have $pl_1(X) = 1$ if $X\cap[\phi_\Delta] \neq \emptyset$, and $pl_1(X) = \varepsilon_1$ otherwise. So, in particular, $pl_1(\alpha) = 1$ iff $[\alpha]\cap[\phi_\Delta] \neq \emptyset$, i.e., iff $\alpha\wedge\phi_\Delta$ is consistent, hence iff $\alpha\rightarrow\beta$ is tolerated by Δ.

(b) Notice that $\alpha\rightarrow\beta$ is not tolerated by $\Delta_i\cup\ldots\Delta_k$ iff α (hence $\alpha\wedge\beta$ and $\alpha\wedge\neg\beta$) is inconsistent with any superset of $\Delta_i\cup\ldots\cup\Delta_k$ iff $pl_i(\alpha\wedge\beta) = pl_i(\alpha\wedge\neg\beta) = \varepsilon_i$.

(c) For i=1, the proposition holds by (a). For i>1, $\alpha\rightarrow\beta \in \Delta_i$ means that $\alpha\wedge\beta$ is consistent with $\Delta_i\cup\ldots\cup\Delta_k$, and that α (hence $\alpha\wedge\beta$ and $\alpha\wedge\neg\beta$) is inconsistent with any superset of $\Delta_i\cup\ldots\cup\Delta_k$. Thus, $pl_i(\alpha\wedge\beta) = \varepsilon_{i-1}$ and $pl_i(\alpha\wedge\neg\beta) = \varepsilon_i$, which means that $bel_i \models \alpha\rightarrow\beta$.

d) By (b) $pl_{i-1}(\alpha\wedge\beta) = pl_{i-1}(\alpha\wedge\neg\beta) = \varepsilon_{i-1}$ means that $\alpha\rightarrow\beta$ is not tolerated by $\Delta_{i-1}\cup\ldots\Delta_k$, hence is not tolerated by any superset of $\Delta_{i-1}\cup\ldots\Delta_k$, therefore $\alpha\rightarrow\beta \notin \Delta_j$ with j<i. Now $bel_i \models \alpha\rightarrow\beta$ implies that $pl_i(\alpha\wedge\beta) >_\infty pl_i(\alpha\wedge\neg\beta)$ then by (b) $\alpha\rightarrow\beta$ is tolerated by $\Delta_i\cup\ldots\Delta_k$ hence $\alpha\rightarrow\beta \in \Delta_i$. ■

Lemma 9. For every $d\in\Delta$, let m_d be the simple support function so that:

$m_d(\Omega) = \varepsilon_d$; $m_d(\phi_d) = 1 - \varepsilon_d$; and $m_d(X) = 0$ otherwise, with $\varepsilon_d \in \mathbb{E}^0$

and let $m_\oplus = \oplus\{m_d \mid d\in\Delta\}$. Then, for any world ω in Ω,

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

and $pl_\oplus(\omega) \approx_\infty 1$ if ω satisfies all the defaults in Δ.

Proof. For a given world ω we have, by definition of plausibility,

$$pl_\oplus(\omega) = \sum_{F:\ \omega\in F} m_H(F) , \qquad (*)$$

$$pl^*(X) = \sum_{i=j}^{n} \varepsilon_i \qquad \text{if } X \cap A_j \neq \emptyset \quad \text{and } X \cap A_{j-1} = \emptyset.$$

These relations are those required by the constraints C_i.

Next we show that any other εbf that satisfy the constraints C will dominate one of the εbf in Λ_C. Consider an εbf bel that satisfies C, and let pl and m be its related plausibility function and basic belief assignment. We show first that if A_1 is not the only focal element of bel with non-infinitesimal mass, than pl dominates the elements of Λ_C. Let B be a focal element of bel with $m(B) \notin \mathcal{E}^0$. If $A \not\subseteq B$, $B \cap A_1 \neq \emptyset$, then $pl(A_1 \cap \overline{B}) \leq 1 - m(B) < 1$ and bel cannot be εLC than bel* as $pl^*(A_1 \cap \overline{B}) = 1 > pl(A \cap \overline{B})$. If $B \cap \overline{A}_1 \neq \emptyset$, then $pl(\overline{A}_1) \geq m(B) > 0$, and we do not have $pl(A_1) >_\infty pl(\overline{A}_1)$, contrary to C_1. So only A_1 can be a focal element with a non infinitesimal basic belief mass, and for all $X \subseteq \Omega$, $X \neq A_1$, $m(X) \in \mathbb{E}^0 \cup \{0\}$. In that case, $\lim_{\eta \to 0} pl(X) = \lim_{\eta \to 0} pl^*(X)$ for all $X \subseteq \Omega$, so we focus on the second requirement of the definition of εLC.

Let the sets $\mathcal{A}_1 = \{X: \ X \subseteq A_1\}$, and $\mathcal{A}_i = \{X: \ X \not\subseteq A_{i-1}, X \subseteq A_i\}$, i = 2,3...n. For i = 1,2...n, let $m_0(A_i) = \Sigma\{m(X): X \in \mathcal{A}_i\}$. By construction m_0 is εLC than m. As already shown, it satisfies C_1, and $m_0(A_i) \in \mathcal{E}^0 \cup \{0\}$ for i = 2,3...n. To satisfy C_i, m_0 must satisfy: $m_0(A_i) \neq 0$ and $m_0(A_i) >_\infty m_0(A_{i+1})$, i =2,3...n-1, in which case $m_0 \in \Lambda_C$. If $m_0(A_i) = 0$ for some $i \in \{2,3\ldots n\}$, then m(X) = 0 for all $X \in \mathcal{A}_i$, in which case $pl(\overline{A}_{i-1} \leftrightarrow A_i) < pl(\overline{A}_i)$, hence it does not satisfy C_i.

Finally, only those elements of Λ_C could be non-dominating εbf that satisfy C. None of them is εLC than any other in Λ_C as they share the same focal elements A_i, so none can be a {0,1}-generalization of the other. So none of them can be taken out of Λ_C. ■

Lemma 7. 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, $\prec_1 \equiv \prec_2$.

Proof. (a) At each step i, we add a new focal element, and give it a mass $\varepsilon_{i-1} - \varepsilon_i$ taken off from Ω. As $\varepsilon_i >_\infty \varepsilon_{i+1}$, the new bel_i still satisfies all the defaults in Δ that were satisfied by bel_{i-1}. In particular, the εbf returned at Step 2 satisfies all of Δ. (b) All the elements of $Bel_{lc}(\Delta)$ are consonant belief functions with the same focal elements. Let ω be a word, and let $\varepsilon_{i-1} - \varepsilon_i$ be the mass given to the smallest focal element that contains ω. Given lemma 3, $pl_{\mathcal{E}}(\omega)$ is of order ε_{i-1} for each $bel_{\mathcal{E}}$ in $Bel_{lc}(\Delta)$. As the ordering $\prec$ only depends on the relative order of magnitude of the involved plausibilities, $\prec$ is the same no matter what $bel_{\mathcal{E}}$ in $Bel_{lc}(\Delta)$ we consider. ■

So, we must have $[\alpha]_{\mathcal{E}} = [\alpha\wedge\beta]_{\mathcal{E}}$, and hence each $bel_{\mathcal{E}}$-preferred world of α satisfies β, thus concluding the proof. ■

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

Proof. For the only-if part, note that infinitesimal probability distributions $P_{\mathcal{E}}$ are a special case of εbf where only singletons are assigned non-zero masses. Thus, EBF(Δ) includes all the infinitesimal probability distributions $P_{\mathcal{E}}$ compatible with Δ, and we have $\alpha \mid\!\sim_{bf} \beta$ only if, for all such distributions, $P_{\mathcal{E}} \models \alpha\rightarrow\beta$. By Lemma 4, the latter condition is equivalent to $\lim_{\eta\rightarrow 0} P_{\mathcal{E}}(\beta|\alpha) = 1$, which is the definition of Adams' ε–consequence. So, we have $\alpha \mid\!\sim_{bf} \beta$ only if $\alpha \mid\!\sim_{\varepsilon} \beta$, and the thesis follows by recalling the equivalence between $\mid\!\sim_{\varepsilon}$ and $\mid\!\sim_{\mathbf{P}}$ (Lehmann and Magidor, 1992, Lemmas 4.7 and 4.9). To prove the if part, we show that each inference relation induced by any $bel_{\mathcal{E}}$ in EBF(Δ) is preferential. Take any $bel_{\mathcal{E}}$, and consider $W = (\Omega, Id, \prec_{\mathcal{E}})$, where Id is the identity function and $\prec_{\mathcal{E}}$ is the $bel_{\mathcal{E}}$-preference relation. Clearly, W is a preferential model. Moreover, by Lemma 5

$$W \models \alpha\rightarrow\beta \;\text{ iff }\; bel_{\mathcal{E}} \models \alpha\rightarrow\beta,$$

that is, W represents the inference relation induced by $bel_{\mathcal{E}}$. By the representation theorem of Kraus et al. (1990, Theorem 5.18), this means that this inference relation is preferential, and therefore it satisfies all the rules of **P**. As this is true for any $bel_{\mathcal{E}}$ in EBF(Δ), then $\mid\!\sim_{bf}$ also satisfies the rules of **P**, and so it contains all the preferential consequences of Δ. ■

The first part can also be proved by showing the transformation between the so-called uniform sequence of probability functions by Adams (1966) and infinitesimal probabilities.

Lemma 6: Consider a set $\{A_1,...,A_n\}$of n nested subsets of Ω, with $A_1 \subseteq A_2 \subseteq ... \subseteq A_n$, $A_1\neq\emptyset$, $A_n=\Omega$. Let $C = \{C_1,...,C_n\}$ be a set of constraints C_i given by:

C_1 : $pl(A_1) = 1 >_\infty pl(\overline{A_1})$, and

C_i : $pl(\overline{A_{i-1}}\cap A_i) >_\infty pl(\overline{A_i})$, i=2,3...n.

Then

$$\Lambda_C = \{bel\text{: } m(A_i) = \varepsilon_i \in \mathbb{E}^0, \; i=2,3...n, \; \varepsilon_i >_\infty \varepsilon_{i+1}, \; i=2,3...n-1, \text{ and } m(A_1) = 1-\sum_{i=2}^{n}\varepsilon_i\}.$$

Proof: For the proof, we take the family Λ_C of εbf, we show that its element satisfy the constraints C, that any εbf that satisfy the constraints C will either belong to Λ_C or will dominate one of the εbf in Λ_C, and none of the εbf in Λ_C dominates another εbf in Λ_C.

Let $bel^* \in \Lambda_C$. By construction, its related plausibility function pl^* satisfies, for all $X\subseteq\Omega$,

$pl^*(X) = 1$ if $X\cap A_1\neq\emptyset$, and

Let $s(\eta)$ denote the last term. The above inequality tells us that $bel_{\mathcal{E}}(\beta|\alpha) \geq 1-s$. (Recall that the $\geq$ is taken pointwise for any $\eta \in (0,1)$.) The hypothesis that $bel_{\mathcal{E}}$ is an εbf-model of $\alpha \to \beta$ means that $pl_{\mathcal{E}}(\alpha\wedge\beta) >_{\infty} pl_{\mathcal{E}}(\alpha\wedge\neg\beta)$, which means (Lemma A4(b))

$$\lim_{\eta\to 0} \frac{pl_{\mathcal{E}}(\alpha\wedge\neg\beta)}{pl_{\mathcal{E}}(\alpha\wedge\beta)} = \lim_{\eta\to 0} s(\eta) = 0.$$

Then, from $bel_{\mathcal{E}}(\beta|\alpha) \geq 1-s$ and $bel_{\mathcal{E}}(\beta|\alpha) \leq 1$ (by definition), we get $\lim_{\eta\to 0} bel_{\mathcal{E}}(\beta|\alpha) = 1$ as desired.

To prove the reverse implication, assume $\lim_{\eta\to 0} bel_{\mathcal{E}}(\beta|\alpha) = 1$. Then, by applying the definition of $bel_{\mathcal{E}}(\beta|\alpha)$,

$$\lim_{\eta\to 0} \left(1 - \frac{pl_{\mathcal{E}}(\alpha\wedge\neg\beta)}{pl_{\mathcal{E}}(\alpha)}\right) = 1, \text{ and so}$$

$$\lim_{\eta\to 0} \frac{pl_{\mathcal{E}}(\alpha\wedge\neg\beta)}{pl_{\mathcal{E}}(\alpha)} = 0,$$

which means $pl_{\mathcal{E}}(\alpha) >_{\infty} pl_{\mathcal{E}}(\alpha\wedge\neg\beta)$. Lemma 2(a) tells us that $pl_{\mathcal{E}}(\alpha) \approx_{\infty} pl_{\mathcal{E}}(\alpha\wedge\neg\beta) + pl_{\mathcal{E}}(\alpha\wedge\beta)$ (we just consider a coarsening of Ω consisting of the two elements $[\alpha\wedge\neg\beta]$ and $[\alpha\wedge\beta]$), and so, by Lemma A9(b), $pl_{\mathcal{E}}(\alpha\wedge\neg\beta) + pl_{\mathcal{E}}(\alpha\wedge\beta) >_{\infty} pl_{\mathcal{E}}(\alpha\wedge\neg\beta)$. Finally, we apply Lemma A8(b) to get $pl_{\mathcal{E}}(\alpha\wedge\beta) >_{\infty} pl_{\mathcal{E}}(\alpha\wedge\neg\beta)$, which concludes the proof. ■

Lemma 5. Let $bel_{\mathcal{E}}$ be an εbf on Ω. For any α, β formulae of $\mathcal{L}$, $bel_{\mathcal{E}} \models \alpha \to \beta$ if, and only if, each bel-preferred world of α satisfies β.

Proof. For the length of the proof, we denote by $[\alpha]_{\mathcal{E}}$ the set of $bel_{\mathcal{E}}$-preferred worlds of α. To prove the if part, assume that for each $bel_{\mathcal{E}}$-preferred world ω of α we have $\omega \models \beta$. First, we show that $[\alpha]_{\mathcal{E}} = [\alpha\wedge\beta]_{\mathcal{E}}$. To see that $[\alpha]_{\mathcal{E}} \subseteq [\alpha\wedge\beta]_{\mathcal{E}}$, suppose there is a ω which is a $bel_{\mathcal{E}}$-preferred model of α but not of $\alpha\wedge\beta$, then, there should exist a world ω' which satisfies $\alpha\wedge\beta$ (and therefore α) such that $\omega' \prec_{\mathcal{E}} \omega$, but this contradicts the fact that ω is a $bel_{\mathcal{E}}$-preferred world of α. To see that $[\alpha]_{\mathcal{E}} \supseteq [\alpha\wedge\beta]_{\mathcal{E}}$, let ω be a $bel_{\mathcal{E}}$-preferred world of $\alpha\wedge\beta$ but not of α, then, there should exist a $bel_{\mathcal{E}}$-preferred world ω' of α such that $\omega' \prec_{\mathcal{E}} \omega$. Since $\omega' \models \beta$, then ω' also satisfies $\alpha\wedge\beta$, but this contradicts the fact that ω is a $bel_{\mathcal{E}}$-preferred world of $\alpha\wedge\beta$. Hence, we must have $[\alpha]_{\mathcal{E}} = [\alpha\wedge\beta]_{\mathcal{E}}$. Now, for each $bel_{\mathcal{E}}$-preferred world ω of $[\alpha\wedge\neg\beta]_{\mathcal{E}}$ there is a $\omega' \in [\alpha\wedge\beta]_{\mathcal{E}}$ such that $\omega' \prec_{\mathcal{E}} \omega$. (If not, then ω' would be a bel-preferred world of α, contradicting the fact that $[\alpha]_{\mathcal{E}} = [\alpha\wedge\beta]_{\mathcal{E}}$.) But then, by definition of $\prec_{\mathcal{E}}$, we have $pl_{\mathcal{E}}(\alpha\wedge\beta) >_{\infty} pl_{\mathcal{E}}(\alpha\wedge\neg\beta)$, which means that $bel_{\mathcal{E}} \models \alpha \to \beta$, thus concluding the proof of the if side. To prove now the only if part, it is enough to show that if $pl_{\mathcal{E}}(\alpha\wedge\beta) >_{\infty} pl_{\mathcal{E}}(\alpha\wedge\neg\beta)$ then $[\alpha]_{\mathcal{E}} = [\alpha\wedge\beta]_{\mathcal{E}}$, for all the worlds in $[\alpha]_{\mathcal{E}}$ satisfy β. Suppose, by way of refutation, that $[\alpha]_{\mathcal{E}} \neq [\alpha\wedge\beta]_{\mathcal{E}}$. Then, there exists a world ω that satisfies $\alpha\wedge\neg\beta$ and which is a $bel_{\mathcal{E}}$-preferred world of α. This means that there is no world ω' that satisfies $\alpha\wedge\beta$ such that $\omega' \prec_{\mathcal{E}} \omega$. But then $pl_{\mathcal{E}}(\alpha\wedge\beta)$ is not preferred to $pl_{\mathcal{E}}(\alpha\wedge\neg\beta)$, which contradicts our hypothesis.

implies that, for each η, the sum of all the masses is 1. Therefore, m_{12} is an ε-mass assignment as per Definition 3. ■

Lemma 3. Let $bel_{\mathcal{E}}$ be an εbf on Ω. For any $X \subseteq \Omega$,

(a) $pl_{\mathcal{E}}(X) \approx_{\infty} \sum_{x \in X} pl_{\mathcal{E}}(x)$.

(b) $pl_{\mathcal{E}}(X) \approx_{\infty} \max\{pl_{\mathcal{E}}(x) \mid x \in X\}$.

Proof. (a) By definition of pl, we have

$$\sum_{x \in X} pl_{\mathcal{E}}(x) \quad = \quad \sum_{x \in X} \quad \sum_{Y:\, x \in Y} m_{\mathcal{E}}(Y) \quad = \quad \sum_{Y \cap X \neq \emptyset} k_Y \cdot m_{\mathcal{E}}(Y) \ ,$$

where $k_Y = |Y \cap X|$, for $Y \subseteq \Omega$, is the number of times that the term $m_{\mathcal{E}}(Y)$ appears in the summation in the middle term. As $m_{\mathcal{E}}(Y) \in \mathbb{E}$ for all $Y \subseteq \Omega$, then, $k_Y \cdot m_{\mathcal{E}}(Y) \approx_{\infty} m_{\mathcal{E}}(Y)$ by Lemma A8(e). By repeatedly applying Lemma A9(c),

$$\sum_{Y \cap X \neq \emptyset} k_Y \cdot m_{\mathcal{E}}(Y) \quad \approx_{\infty} \quad \sum_{Y \cap X \neq \emptyset} m_{\mathcal{E}}(Y) \ .$$

But the right hand side is exactly the expression of $pl_{\mathcal{E}}(X)$, and the thesis thus follows from Lemma A9(e) and the transitivity of $\approx_{\infty}$.

(b) The case when $pl_{\mathcal{E}}(X) = 0$ is trivial, as $pl_{\mathcal{E}}(X) = 0$ if and only if $pl_{\mathcal{E}}(x) = 0$ for all $x \in X$ (this is an immediate consequence of the definition of pl), and $0 \approx_{\infty} 0$ by stipulation. So we assume that $pl_{\mathcal{E}}(X) \neq 0$. But this means that there is some $x \in X$ such that $pl_{\mathcal{E}}(x) \neq 0$, so we can apply Lemma **A11**, and the thesis follows immediately from part (a) above. ■

Lemma 4. Let $bel_{\mathcal{E}}$ be an εbf, and let $\alpha \rightarrow \beta$ be a default rule.

(i) $bel_{\mathcal{E}} \models \alpha \rightarrow \beta$ iff $\max_{\omega \models \alpha \wedge \beta} pl_{\mathcal{E}}(\omega) >_{\infty} \max_{\omega \models \alpha \wedge \neg\beta} pl_{\mathcal{E}}(\omega)$.

(ii) $bel_{\mathcal{E}} \models \alpha \rightarrow \beta$ iff $\lim_{\eta \rightarrow 0} bel_{\mathcal{E}}(\beta|\alpha) = 1$.

Proof. (i) The equivalence follows immediately from Definition 4, Lemma 3(b) and Lemma A9(b).

(ii) Assume first that εbf is an εbf-model of $\alpha \rightarrow \beta$. By definition,

$$bel_{\mathcal{E}}(\beta|\_) = 1 - plE(\neg\beta|\alpha) = 1 - \frac{pl_{\mathcal{E}}(\alpha \wedge \neg\beta)}{pl_{\mathcal{E}}(\alpha)} \ .$$

By monotonicity of plausibility, $pl_{\mathcal{E}}(\alpha) \geq pl_{\mathcal{E}}(\alpha \wedge \beta)$, and then

$$\frac{pl_{\mathcal{E}}(\alpha \wedge \neg\beta)}{pl_{\mathcal{E}}(\alpha)} \ \leq \ \frac{pl_{\mathcal{E}}(\alpha \wedge \neg\beta)}{pl_{\mathcal{E}}(\alpha \wedge \beta)} \ .$$

Appendix B. Proofs of technical lemmas.

Lemma 1. Let $m_{\mathcal{E}}$ be an ε-mass assignment. Then $m_{\mathcal{E}}(A) \in \mathbb{E}^1 \cup \{1\}$ for exactly one element $A \subseteq \Omega$, and $m_{\mathcal{E}}(X) \in \mathbb{E}^0 \cup \{0\}$ for all $X \neq A$.

Proof. Let $A_1, ..., A_k$ be the subsets of Ω, $k = 2^{|\Omega|}$, and let t_i denote $m_{\mathcal{E}}(A_i)$, $i = 1,..., k$, where $t_i \in \mathbb{E}$. Suppose $t_i \in \mathbb{E}^0 \cup \{0\}$ for all i. By definition of infinitesimal, for each t_i we can find an $\eta_i \in (0,1)$ such that $t_i(\eta) < \frac{1}{k}$ for any $\eta \leq \eta_i$. Let η^* be the minimum of these η_i. Then, $\sum_{i=1}^{k} t_i(\eta^*) < 1$. But $m_{\mathcal{E}}$ is an ε-mass assignment, and so we must have $\sum_{i=1}^{k} t_i(\eta) = 1$ for all $\eta \in (0,1)$. We have a contradiction: thus there must be at least one t_i that is not in $\mathbb{E}^0 \cup \{0\}$, that is, $t_i \in \mathbb{E}^1 \cup \{1\}$. We now prove that this element is unique. For suppose that it is not, and that there is $A_j \neq A_i$ such that both t_i and t_j are in $\mathbb{E}^1 \cup \{1\}$. This means that both $\lim_{\eta\to 0} t_i = 1$ and $\lim_{\eta\to 0} t_j = 1$. Then, there are η_i and η_j in (0,1) such that $t_i(\eta) > \frac{1}{2}$ for all $\eta \leq \eta_i$, and that $t_j(\eta) > \frac{1}{2}$ for all $\eta \leq \eta_j$. But then, if we let $\eta^* = \min(\eta_i,\eta_j)$, we have $\sum_{i=1}^{k} t_i(\eta^*) \geq t_i(\eta^*) + t_j(\eta^*) > 1$, again contradicting the hypothesis that $m_{\mathcal{E}}$ is an ε-mass assignment. ■

Lemma 2. Let $\mathcal{E}$ be a finite set of infinitesimals, and let m_1 and m_2 be two ε-mass assignments on Ω. Then $m_{12} = m_1 \oplus m_2$ is an ε-mass assignment, provided that the normalization factor in the combination is 1.

Proof. First note that, by definition of ⊕, the normalization factor k is 1 only if the intersection of any focal element of m_1 with any focal element of m_2 is non-empty. Let A be any non-empty subset of Ω. As the normalization factor k = 1, we have

$$m_{12}(A) = \sum_{B \cap C = A} m_1(B) m_2(C) \quad ,$$

where each non-null $m_1(B)$ and $m_2(C)$ is either 1, ε or (1–ε) with $\varepsilon \in \mathbb{E}^0$. Then, $m_{12}(A)$ is the sum of product of terms in $\mathbb{E}$, and this sum belongs to $\mathbb{E}$ if $m_{12}(A) \leq 1$(Lemma A2(a)).

Let $t'(\eta) = m_{12}(A)$. We need to verify that $t'(\eta) \in [0,1]$ for any $\eta \in (0,1)$. Fix a $\eta \in (0,1)$, and let $m_{1|\eta}$ and $m_{2|\eta}$ denote the (standard) basic belief assignments obtained for this η. Let $m_{12|\eta} = m_{1|\eta} \oplus m_{2|\eta}$. Then, $m_{12|\eta}$ is a basic belief assignment (Shafer, 1976, Theorem 3.1). This implies that $m_{12|\eta}(A) \in [0,1]$; as this is true for all η, then $t'(\eta) = m_{12}(A) \in \mathbb{E}$. Moreover, this

(f) If $s \approx_\infty s'$ and $s' >_\infty t$, then $s >_\infty t$.

Proof. All properties are proved by a direct comparison of their orders. ■

Remark A4. The above properties tell us that we can substitute equivalent ($\approx_\infty$) terms inside $\geq_\infty$, $>_\infty$ and $\approx_\infty$ relations; note that (c) and (d) yield $s+t \approx_\infty s'+t'$ and $s\cdot t \approx_\infty s'\cdot t'$ whenever $s \approx_\infty s'$ and $t \approx_\infty t'$. (e) tells us that we can also substitute equal terms.

Lemma A10. Let $t, t' \in \mathbb{E}\cup\mathbb{R}^+$. If $\kappa(t)$ and $\kappa(t')$ are defined, then $\kappa(t+t') = \min(\kappa(t), \kappa(t'))$.

Proof. The proof is the same as the one for lemma A2.a. relative to $\kappa(t+t')$. Indeed that t and t' belong to $\mathbb{E}\cup\mathbb{R}^+$ was not used. The only requirement was that both $\kappa(t)$ and $\kappa(t')$ were defined. ■

Lemma A11. Let $t_1, ..., t_n \in \mathbb{E}$. Then $\max_{i=1,n} t_i \approx_\infty \sum_{i=1}^{n} t_i$.

Proof. Let $k_i = \kappa(t_i)$ for i = 1,...n. By applying Lemmas A10 and A2(b) recursively, we have $\kappa(\sum_{i=1}^{n} t_i) = \min(k_i: i=1,...n)$ and $\kappa(\max(t_i: i=1,...n)) = \min(k_i: i=1,...n)$. Hence the theorem is proved by Lemma A4(c). ■

Lemma A12. Let $t_1, ..., t_n \in \mathbb{E}$, and let t_j be such $t_j \geq_\infty t_i$ for all $i\neq j$. Then, $\max_{i=1,n} t_i \approx_\infty t_j$.

Proof. The inequalities $t_j \geq_\infty t_i$ for all $i\neq j$ means $\kappa(t_j) \leq \kappa(t_i)$ for all $i\neq j$. So $\kappa(t_j) = \min_{i=1,n} \kappa(t_i)$. By Lemma A2(b), $\kappa(\max_{i=1,n} t_i) = \min_{i=1,n} k_i$. Hence the theorem is proved by Lemma A4(c). ■

Lemma A13. Let $\varepsilon_1, \varepsilon_2, \varepsilon_3, \varepsilon_4 \in \mathbb{E}^0$. Then:

a) $\max(\varepsilon_1, \varepsilon_2) >_\infty \varepsilon_3$ iff $\varepsilon_1 >_\infty \varepsilon_3$ or $\varepsilon_2 >_\infty \varepsilon_3$.

b) if $\varepsilon_1 >_\infty \varepsilon_2$, $\varepsilon_3 >_\infty \varepsilon_4$, then $\max(\varepsilon_1, \varepsilon_3) >_\infty \max(\varepsilon_2, \varepsilon_4)$.

Proof: a) By Lemma A4(b), $\max(\varepsilon_1, \varepsilon_2) >_\infty \varepsilon_3$ iff $\kappa(\max(\varepsilon_1, \varepsilon_2)) < \kappa(\varepsilon_3)$. By Lemma A2(b), $\kappa(\max(\varepsilon_1, \varepsilon_2)) = \min(\kappa(\varepsilon_1), \kappa(\varepsilon_2))$. So $\min(\kappa(\varepsilon_1), \kappa(\varepsilon_2)) < \kappa(\varepsilon_3)$, hence either $\kappa(\varepsilon_1) < \kappa(\varepsilon_3)$ or $\kappa(\varepsilon_2) < \kappa(\varepsilon_3)$, in which case either $\varepsilon_1 >_\infty \varepsilon_3$ or $\varepsilon_2 >_\infty \varepsilon_3$ holds by Lemma A4(b).

b) By Lemma A4(b), $\varepsilon_1 >_\infty \varepsilon_2$ iff $\kappa(\varepsilon_1) < \kappa(\varepsilon_2)$, and $\varepsilon_3 >_\infty \varepsilon_4$ iff $\kappa(\varepsilon_3) < \kappa(\varepsilon_4)$. Hence $\max(\kappa(\varepsilon_1), \kappa(\varepsilon_3)) < \max(\kappa(\varepsilon_2), \kappa(\varepsilon_4))$, what implies by Lemma A4(b) that $\max(\varepsilon_1, \varepsilon_3) >_\infty \max(\varepsilon_2, \varepsilon_4)$. ■

(c) $\kappa(t) < \kappa(s)$. ■

Remark A2. These properties show that the elements of $\mathbb{E}^0$ and $\mathbb{E}^1$ correctly capture the intended meaning of values infinitesimally close to 0 and to 1, respectively. Note that $r \approx_\infty t$ entails that all real numbers are of the same order (no real number is infinitesimally larger than any other real number), and of the same order as the terms in $\mathbb{E}^1$.

Lemma A8. Let $s, t \in \mathbb{E}$. Then

(a) $s \geq_\infty t$ iff $s+t \approx_\infty s$, provided $s+t \in \mathbb{E}$.

(b) $s >_\infty t$ iff $s+t >_\infty t$;

(c) if $s \in \mathbb{E}^0$, then $t >_\infty s \cdot t$, provided that $t \neq 0$;

(d) if $s \in \mathbb{E}^1$, then $t \approx_\infty s \cdot t$;

(e) for any $r \in \mathbb{R}^+$, $t \approx_\infty r \cdot t$, provided $r \cdot t \in \mathbb{E}$.

Proof.

(a) As $\kappa(s) \leq \kappa(t)$, and $\kappa(s+t) = \min(\kappa(s), \kappa(t))$ by Lemma A2, thus $\kappa(s+t) = \kappa(s)$.

(b) $\kappa(s) < \kappa(t)$, and $\kappa(s+t) = \min(\kappa(s), \kappa(t)) = \kappa(s)$, so $\kappa(s+t) < \kappa(t)$.

(c) $\kappa(s) \in (0, \infty)$, $\kappa(t) < \kappa(s.t) = \kappa(s) + \kappa(t)$ provided $\kappa(t) \neq \infty$, what is the case as $t \neq 0$.

(d) $\kappa(s) = 0$, $\kappa(s.t) = \kappa(s) + \kappa(t) = \kappa(t)$.

(e) $\kappa(r) = 0$, $\kappa(r.t) = \kappa(r) + \kappa(t) = \kappa(t)$. ■

Remark A3. The above properties show that $\geq_\infty$ captures the desired behavior of a negligibility relation. More specifically, (a) and (b) show that the sum of elements of $\mathbb{E}$ is of the same order than the larger argument of the sum. (c), (d) and (e) show that the product of an element t of $\mathbb{E}^1$ with elements of $\mathbb{E}^1$ or real numbers is of the same order as t; and with an element of $\mathbb{E}^0$ is of a smaller order. Note that (c) gives $t >_\infty t \cdot t$ for all $t \in \mathbb{E}^0$ as a special case.

Note that if we do not know anything about the relation between t and t', we can neither say much about t+t'. For example, from $t >_\infty s$ and $t' \geq_\infty s'$ we cannot infer $t+t' >_\infty s+s'$. (Take $\kappa(t') < \kappa(t)$, $\kappa(s') < \kappa(s)$ and $\kappa(t') = \kappa(s')$, then $\kappa(t+t') = \min(\kappa(t), \kappa(t')) = \kappa(t') = \kappa(s') = \min(\kappa(s), \kappa(s')) = \kappa(s+s')$, hence $t+t' \approx_\infty s+s'$.) However, from $t >_\infty s$ and $t' >_\infty s'$ we always have $t+t' >_\infty s+s'$.

Lemma A9. Let $s, s', t, t' \in \mathbb{E}$. Then

(a) If $s \approx_\infty s'$ and $t \approx_\infty t'$, then $s \geq_\infty t$ iff $s' \geq_\infty t'$;

(b) If $s \approx_\infty s'$ and $t \approx_\infty t'$, then $s >_\infty t$ iff $s' >_\infty t'$;

(c) If $s \geq_\infty s'$ and $t \geq_\infty t'$, then $s + t \geq_\infty s' + t'$;

(d) If $s \geq_\infty s'$ and $t \geq_\infty t'$, then $s \cdot t \geq_\infty s' \cdot t'$;

(e) if $s = t$, then $s \approx_\infty t$.

(a) By definition, $t \geq_\infty t'$ iff $\lim_{\eta\to0} \frac{t'}{t} \in [0,\infty)$, what is achieved iff $\kappa(t) \leq \kappa(t')$.

(b) From (a), $t >_\infty t'$ means $\kappa(t) \leq \kappa(t')$ and not $\kappa(t') \leq \kappa(t)$, hence $\kappa(t) < \kappa(t')$. Furthermore, $\lim_{\eta\to0} \eta^{\kappa(t')-\kappa(t)} = 0$ in that case.

(c) $t \approx_\infty t'$ means thus $\kappa(t) \leq \kappa(t')$ and $\kappa(t') \leq \kappa(t)$, hence $\kappa(t) = \kappa(t')$. ■

The $\geq_\infty$ relation is extended to $\mathbb{R}^+$, the set of positive reals. For $\alpha, \beta \in \mathbb{R}^+ \cup \mathbb{E}$, we write $\alpha \geq_\infty \beta$ to mean $\lim_{\eta\to0} \frac{\beta}{\alpha} \in [0,\infty)$, and similarly for $\alpha >_\infty \beta$ and $\alpha \approx_\infty \beta$.

Lemma A5. Let $\alpha \in \mathbb{R}^+$ and $t \in \mathbb{E}$.

(a) $\kappa(\alpha) = 0$ and $\lambda(\alpha) = \alpha$.

(b) if $\alpha t < 1$, $\kappa(\alpha t) = \kappa(t)$ and $\lambda(\alpha t) = \alpha\lambda(t)$.

(c) $\alpha \geq_\infty t$.

<u>Proof</u>: (a) $\lim_{\eta\to0} \frac{\alpha}{\eta^0} = \alpha$.

(b) $\lim_{\eta\to0} \frac{\alpha t(\eta)}{\eta^k} = \alpha \lim_{\eta\to0} \frac{t(\eta)}{\eta^k}$, so $\kappa(\alpha t) = \kappa(t)$ and $\lambda(\alpha t) = \alpha\lambda(t)$.

(c) One has $\kappa(t) \geq 0$, and by lemma A5.a, $\kappa(\alpha) = 0$. So $\kappa(\alpha) \leq \kappa(t)$ and by lemma A4.a, $\alpha \geq_\infty t$

■

In the next lemmas, most proofs are trivial thanks to the relation between $\geq_\infty$ and the inequalities on the orders.

Lemma A6. The $\geq_\infty$ relation is a pre-order, i.e., it is reflexive and transitive.

<u>Proof.</u> As $t \geq_\infty t'$ iff $\kappa(t) \leq \kappa(t')$, and $\leq$ is a pre-order on the integers, then $\geq_\infty$ is a pre-order. ■

Remark A1. The expected properties of the induced relations $>_\infty$ and $\approx_\infty$ follow from the properties of $\geq_\infty$. For example, transitivity of $\approx_\infty$ follows from transitivity of $\geq_\infty$. Properties involving different operators also follow easily. E.g., if $t \approx_\infty t'$ and $t' >_\infty t''$, then $t >_\infty t''$.

Lemma A7. Let $s \in \mathbb{E}^0$, $t \in \mathbb{E}^1$, and $r \in \mathbb{R}^+$. Then

(a) $s >_\infty 0$ and $t >_\infty 0$;

(b) $r >_\infty s$ and $r \approx_\infty t$;

(c) $t >_\infty s$.

<u>Proof:</u> We have $\kappa(s) \in (0, \infty)$, $\kappa(t) = 0$, $\kappa(r) = 0$ and $\kappa(0) = \infty$.

(a) $\kappa(s) < \kappa(0)$ and $\kappa(t) < \kappa(0)$.

(b) $\kappa(r) < \kappa(s)$ and $\kappa(r) = \kappa(t)$.

$\sigma = \lim_{\eta\to 0} \frac{t}{\eta^{\kappa(t)}} \cdot \frac{t'}{\eta^{\kappa(t')}} = \lim_{\eta\to 0} \frac{t}{\eta^{\kappa(t)}} \cdot \lim_{\eta\to 0} \frac{t'}{\eta^{\kappa(t')}}$,

and both limits are well defined. Then $\sigma\in \mathbb{R}-\{0\}$ and thus $\kappa(t.t') = \kappa(t)+\kappa(t')$.

Suppose $\kappa(t) = \min(\kappa(t), \kappa(t'))$, and $\kappa(t) \le k^*+k^{**} < \kappa(t)+\kappa(t')$. We can rewrite σ as

$\sigma = \lim_{\eta\to 0} \frac{t}{\eta^{\kappa(t)}} \cdot \lim_{\eta\to 0} \frac{t'}{\eta^{k^*+k^{**}-\kappa(t')}}$. The second term tends to 0, hence $\sigma\notin \mathbb{R}-\{0\}$. By symmetry, the same holds if $\kappa(t') = \min(\kappa(t), \kappa(t'))$.

Suppose $k^*+k^{**}<\min(\kappa(t), \kappa(t'))$, then both terms tend to 0, hence $\sigma\notin \mathbb{R}-\{0\}$.

Suppose $k^*+k^{**} > \kappa(t)+\kappa(t')$, then we can rewrite σ as

$\sigma = \lim_{\eta\to 0} \frac{t}{\eta^{\kappa(t)}} \cdot \lim_{\eta\to 0} \frac{t'}{\eta^{k^*+k^{**}-\kappa(t')}}$. The first term is in $\mathbb{R}-\{0\}$ whereas the second term tends to ∞, hence $\sigma\notin \mathbb{R}-\{0\}$.

(d) Let $\sigma = \lim_{\eta\to 0} \frac{t}{\eta^k} - \lim_{\eta\to 0} \frac{t'}{\eta^k}$ .With $\kappa(t) < \kappa(t')$, then $\sigma\in \mathbb{R}-\{0\}$ iff $k = \kappa(t)$. As $\sigma = \lim_{\eta\to 0} \frac{t-t'}{\eta^k}$, then $\kappa(t-t') = \kappa(t)$. ■

The next lemma shows that the dominant terms contain the information needed to determine the limit of the ratio of two infinitesimals.

Lemma A3: Let $t_1,t_2\in \mathbb{E}$. Whenever $t_2 \neq 0$,

$$\lim_{\eta\to 0} \frac{t_1}{t_2} = \lim_{\eta\to 0} \frac{\lambda(t_1)\eta^{\kappa(t_1)}}{\lambda(t_2)\eta^{\kappa(t_2)}} = \frac{\lambda(t_1)}{\lambda(t_2)} \lim_{\eta\to 0} \eta^{\kappa(t_1)-\kappa(t_2)}.$$

Proof:

$$\lim_{\eta\to 0} \frac{t_1}{t_2} = \lim_{\eta\to 0} \frac{t_1/\eta^{\kappa(t_1)}}{t_2/\eta^{\kappa(t_2)}} \eta^{\kappa(t_1)-\kappa(t_2)} = \lim_{\eta\to 0} \frac{t_1/\eta^{\kappa(t_1)}}{t_2/\eta^{\kappa(t_2)}} \lim_{\eta\to 0} \eta^{\kappa(t_1)-\kappa(t_2)} =$$

$$\frac{\lim_{\eta\to 0} t_1/\eta^{\kappa(t_1)}}{\lim_{\eta\to 0} t_2/\eta^{\kappa(t_2)}} \lim_{\eta\to 0} \eta^{\kappa(t_1)-\kappa(t_2)} = \frac{\lambda(t_1)}{\lambda(t_2)} \lim_{\eta\to 0} \eta^{\kappa(t_1)-\kappa(t_2)}.$$ ■

We show that the definitions of $\geq_\infty$, $>_\infty$ and $\approx_\infty$ can equivalently be expressed by inequalities between orders.

Lemma A4: Let $t, t'\in \mathbb{E}$. Then:

(a) $t \geq_\infty t'$ iff $\kappa(t) \le \kappa(t')$,

(b) $t >_\infty t'$ iff $\kappa(t) < \kappa(t')$ iff $\lim_{\eta\to 0} \frac{t'}{t} = 0$,

(c) $t \approx_\infty t'$ iff $\kappa(t) = \kappa(t')$.

Proof: By Lemma A3, we have $\lim_{\eta\to 0} \frac{t'}{t} = \frac{\lambda(t')}{\lambda(t)} \lim_{\eta\to 0} \eta^{\kappa(t')-\kappa(t)}$ where $\lambda(t)\neq 0$ and $\lambda(t')\neq 0$.

Appendix A: properties of infinitesimals.

We present the formal proofs of those properties of infinitesimals used in this paper and defined in section 2.4. By definition 2, infinitesimals are defined such that their domain is (0,1), their limit for $\eta\to 0$ is 0, and their order is always a well-defined non-negative integer. The set $\mathbb{E}^0$ is the set of infinitesimals, $\mathbb{E}^1 = \{1\text{-}\varepsilon \mid \varepsilon\in \mathbb{E}^0\}$ and $\mathbb{E} = \mathbb{E}^0\cup\mathbb{E}^1\cup\{0\}\cup\{1\}$.

Definition A1: For any continuous function $f:\mathbb{E}\to\mathbb{E}$ whose order $\kappa(f)$ is defined, the limit $\lim_{\eta\to 0}\dfrac{f(\eta)}{\eta^{k(f)}}$ is denoted $\lambda(f)$, the term $\lambda(f).\eta^{\kappa(f)}$ is called the dominant term of f and $\lambda(f)$ is called the coefficient of the dominant term. $\lambda(0) = 0$ and $\kappa(0) = \infty$ by convention.

Lemma A1: For $\varepsilon\in \mathbb{E}^0\cup\{0\}$, $\kappa(1\text{-}\varepsilon) = 0$ and $\lambda(1\text{-}\varepsilon) = 1$.
Proof: We have $\lim_{\eta\to 0}\dfrac{1-\varepsilon(\eta)}{\eta^0} = 1$ for $\varepsilon\in \mathbb{E}^0\cup\{0\}$.

Lemma A2: Let $t,t'\in \mathbb{E}$.
(a) $\kappa(t+t') = \min(\kappa(t), \kappa(t'))$ and if $t+t' \le 1$, $t+t'\in \mathbb{E}$.
(b) $\kappa(\max(t, t')) = \min(\kappa(t), \kappa(t'))$ and $\max(t, t')\in \mathbb{E}$.
(c) $\kappa(t.t') = \kappa(t)+\kappa(t')$ and $t.t'\in \mathbb{E}$.
(d) if $\kappa(t)<\kappa(t')$, then $\kappa(t\text{-}t') = \kappa(t)$ and if $t\text{-}t' \ge 0$, $t\text{-}t'\in \mathbb{E}$.
Proof: The second half of each assertion is valid by noticing that 1) sums, maximum, products and differences preserve continuity, and 2) the ranges and the limits conditions are satisfied for maximum and products and for differences and additions, thanks to the extra requirements. As both t and t' belong to $\mathbb{E}$, their orders are well defined. Therefore the orders are also well defined for their sum, max, product and difference. The first half of each assertion is proved as follow.
(a) Let $\sigma = \lim_{\eta\to 0}\dfrac{t}{\eta^k} + \lim_{\eta\to 0}\dfrac{t'}{\eta^k}$. The value of σ equals 0 if $k<\min(\kappa(t), \kappa(t'))$, $\lambda(t)$ if $k=\kappa(t)<\kappa(t')$, $\lambda(t')$ if $k=\kappa(t')<\kappa(t)$, and $\lambda(t)+\lambda(t')$ if $k=\kappa(t)=\kappa(t')$, and is infinite when $k>\min(\kappa(t), \kappa(t'))$. The order of σ is defined iff $\sigma\in \mathbb{R}-\{0\}$. This is achieved iff $k = \min(\kappa(t), \kappa(t'))$. As $\sigma = \lim_{\eta\to 0}\dfrac{t+t'}{\eta^k}$, $\kappa(t+t') = \min(\kappa(t), \kappa(t'))$.
(b) Let $\sigma = \max(\lim_{\eta\to 0}\dfrac{t}{\eta^k}, \lim_{\eta\to 0}\dfrac{t'}{\eta^k})$. One has $\sigma\in \mathbb{R}-\{0\}$ iff $k = \min(\kappa(t), \kappa(t'))$. Would k be smaller, both term would have a zero limit, and would it be larger, at least one term would be infinite. As $\sigma = \lim_{\eta\to 0}\dfrac{\max(t,t')}{\eta^k}$, $\kappa(\max(t, t')) = \min(\kappa(t), \kappa(t'))$.
(c) Let $\sigma = \lim_{\eta\to 0}\dfrac{t.t'}{\eta^{k^*+k^{**}}}$.
When $k^*+k^{**} = \kappa(t)+\kappa(t')$, σ can be written as

Smets Ph. (1998) The Transferable Belief Model for Quantified Belief Representation. in Handbook of defeasible reasoning and uncertainty management systems. Gabbay D. M. and Smets Ph. Eds. Vol. 1, Kluwer, Doordrecht, 267-301.

Snow P. (1996) Standard probability distributions described by rational default entailment. Personnal communication.

Snow, P. (1999) Diverse confidence levels in a probabilistic semantics for conditional logics. Artificial Intelligence, 113: 269-279

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

Walley, P. (1991) *Statistical reasoning with imprecise probabilities.* Chapman and Hall, London.

Weydert, E. (1995) Defaults and infinitesimals Defeasible Inference by Nonarchimedean Entropy Maximization. In Procs. of 11th. Conf. on Uncertainty in Artificial Intelligence (UAI'95), pp. 540-547.

Wilson, N. (1992) Some theoretical aspects of the Dempster-Shafer theory, PhD thesis, Oxford Polytechnic, May 1992.

Wilson N. (1993) Default logic and Dempster-Shafer theory, 2nd European Conference on Symbolic and Qualitative approaches to reasoning and Uncertainty (ECSQARU-93), (Kruse et al., eds), Springer-Verlag, 372-379.

Wilson, N. (1996) Extended probability. Procs. of hte 12th European Conf. on Artif. Intell. (ECAI'96), John Wiley and Sons Ltd, 667-671.

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

Lehmann, D. (1989) What does aconditional knowledge base entail ? Proc. of the 1st Inter. Conf. onPrinciples of Knowledge Representation and Reasoning (KR'89), R.J. Brachmanet al., eds., Toronto, 1989, 212-222.

Lehmann, D. (1995) Another perspective on Default Reasoning. Annals of Mathematical and Artificial Intelligence 15:61-82.

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

Makinson, D (1989). "General theory of cumulative inference", In: *Non-Monotonic Reasoning* (Proc. of the 2nd Inter. Workshop, Grassau, FRG, June 1988) (M. Reinfrank, J. De Kleer, M.L. Ginsberg, E. Sandewall, eds.), LNAI, Vol. 346, Springer Verlag, Berlin, 1989, p. 1-18.

Nelson, E. (1977) Internal set theory: a new approach to non-standard analysis. Bull. Amer. Math. Soc. 83, 1165-1198.

Parsons, S. and Bourne, R. A. (2000) On proofs in System P. To appearin International Journal of Uncertainty, Fuzziness and Knowledge base Systems, World Scientific Company.

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.

Pinkas, G. (1991). Propositional nonmonotonic reasoning and inconsistency in symmetric neural networks. In Proc. of the 12th IJCAI, pages 525-530, Sydney, Australia, Morgan-Kaufmann.

Reiter, R. (1980) A Logic for default reasoning. Artificial Intelligence 13, 81-132.

Reiter R., Griscuolo G. (1981) On interacting defaults, Proc. nter. Joint Conf. on Artificial Intelligence IJCAI-81, pp. 270-276.

Robinson, A. (1966) *Non-standard Analysis.* Noth-Holland, Amsterdam.

Schurz, G. (1998) Probabilistic semantics for Delgrande's conditional logic and a counterexample to his default logic. Artificial Intelligence, 102: 81-95

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. (1990) The combination of evidence in the Transferable Belief Model. IEEE Trans. PAMI 12, 447-458.

Smets, Ph (1992) The concept of distinct evidence. IPMU 92 Proceedings.pg. 789-794.

Smets, Ph. (1997a). The normative representation of quantified beliefs by belief functions. Artificial Intelligence 92: 229-242.

Smets, Ph. (1997b) The α-junctions: combination operators applicable to belief functions. in Gabbay D., Kruse R., Nonnengart A. and Ohlbach H. J. (eds.) Qualitative and quantitative practical reasoning. Springer, Berlin, pg 131-153.

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.

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

Dubois D., J. Lang, H. Prade (1992) 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.

Dubois D., Prade H. (1995). Conditional objects, possibility theory and default rules. In *Conditionals: From Philosophy to Computer Sciences* (G. Crocco, L. Farinas del Cerro, A. Herzig, eds.) Oxford University Press, pp. 311-346.

Dupin de Saint Cyr, F., Lang, J. and Schiex T. (1994) Penalty logic and its link with Dempster-Shafer theory. In Proc. of 10^{th} Conference on Uncertainty in Artificial Intelligence (UAI-94) , pp. 204-211.

Eiter, T. and Lukasiewicz, T. (2000). Complexity Results for Default Reasoning from Conditional Knowledge Bases. In Principles of Knowledge Representation and Reasoning: Proceedings of the Seventh International Conference (KR2000), pages 62-73, Breckenridge, Colorado, USA, April 2000.

Keisler, H.J. (1976) *Foundations of infinitesimal calculus.* Prindle, Weber and Schmidt, Eds., Boston, MA.

Kennes, R. (1991) Evidential reasoning in a categorial perspective: conjunction and disjunction of belief functions. Procs. of the 7th Conf. on Uncertainty in AI. Los Angeles, CA, 174-181.

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.

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

Geffner, H. (1992). *Default reasoning: Causal and conditional theories*. MIT Press.

Gilio, A. (2000) Precise Propagation of Upper and Lower Probability Bounds in System P. In Proceedings of the Eighth International Workshop on Non-Monotonic Reasoning (NMR'2000), Special Session on Uncertainty Frameworks in Nonmonotonic Reasoning, Breckenridge, Colorado, USA, April 2000.

Goldszmidt, M. and Pearl, J. (1991) System Z^+: A formalism for reasoning with variable-strength defaults. Procs. of the AAAI Conference, Anaheim, CA, pp. 399-404

Goldszmidt M. (1992) Qualitative Probabilities: A Normative Framework for Commonsense Reasoning. PhD Thesis, Technical Report R-190, Cognitive Systems Laboratory, UCLA, Los Angeles.

Goldszmidt M., Pearl J. (1996) Qualitative probabilities for default reasoning, belief revision, and causal modeling. Artificial Intelligence 84: 57-112.

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

Klawonn, F. and Smets, Ph (1991) The dynamics of belief in the Transferable Belief Model. Procs. of the 8th Conf. on Uncertainty in AI. Morgan Kaufmann, San Mateo, CA, 130-137.

Kohlas J. and Monney P.A. (1995) A Mathematical Theory of Hints. An Approach to Dempster-Shafer Theory of Evidence. Lecture Notes in Economics and Mathematical Systems No. 425. Springer-Verlag.

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

Lang J. (1994) Syntax-based default reasoning as probabilistic model-based diagnosis. In Proc. of UAI'94, pp. 391-398.

Finally, we note that the approach presented in this paper is purely semantic: we have defined non-monotonic consequence in terms of satisfaction in a certain class of models. It would be interesting to derive a syntactic characterization of the different εbf-based consequence relations. Also, we should like to find effective algorithms to compute these relations. This is especially true of **LCD**. These tasks are left for future work.

Acknowledgments. The work reported in this paper has been supported in part by the DRUMS II Basic Research Action, funded by the EEC, and in part by the BELON project, funded by the *Communauté Française de Belgique*. Work by the second author was performed while he was with IRIDIA, Université Libre de Bruxelles, under the BELON grant. The authors are indebted to Nic Wilson and an anonymous referee for their in-depth and very constructive suggestions for handling several subtle concepts needed in this paper.

References

Adams, E., W. (1966) Probability and the logic of conditionals, in Aspects of inductive logic, J. Hintikka and P. Suppes (eds), North Holland, Amesterdam, 253-316.

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

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

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, pp. 673-684.

Benferhat, S., Cayrol, C., Dubois, D., Lang, J. and Prade, H. (1993) Inconsistency management and prioritized syntax-based entailment. Procs. of IJCAI’93, pp. 640-645.

Benferhat, S., Dubois, D., Lang, J., Prade, H., Saffiotti A. and Smets P. (1998) A general approach for inconsistency handling and merging information in prioritized knowledge bases. Proc. of KR-98.

Benferhat, S., Dubois, D. and Prade, H. (1999). Possibilistic and Standard Probabilistic Semantics of Conditional Knowledge Bases. Journal of Logic and Computation, 9: 873-895.

Bourne R. A. and Parsons S. (1999) Maximum entropy and variable strength defaults. In Procs. of IJCAI-99, Stockholm.

Boutilier, C. (1992). What is a Default priority?. Proc. of the 9th Canadian Conf. on Artificial Intelligence (AI'92), Vancouver, May 1992, 140-147.

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

Delgrande J.P. & Schaub T.H. (1994). A general approach to specificity in default reasoning. Procs. of the 4rd Conf. on Principles of Knowledge Representation and Reasoning KR’94, pp. 146-157.

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

Interestingly, both **LCD** and Geffner’s system exhibit this good behavior without having to satisfy rational monotonicity.

Table 3 summarizes the essential properties that underlie the systems analyzed in this paper. On the left we list the existing systems, and on the right the equivalent system based on εbf's. The $\mid\!\sim_{lcd}$ system is a special case of the $\mid\!\sim_{\oplus}$ family. We start with the concept of a default base issued by a single source, and apply the auto-deduction principle, in which case we get a system equivalent to **P**. We then add the concept of ε-least commitment applied on Δ, and we get a system equivalent to **Z**. We proceed then by assuming instead that each rule in Δ is issued by a different (distinct) source of information, apply the ε-least commitment principle to each $d \in \Delta$, and combine the individual belief functions by Dempster's rule of combination. Further explicit constraints on the infinitesimals lead to systems equivalent to the penalty system, to the lexicographic system, to Brewka system and to Geffner system. If these explicit constraints are replaced by the application of the least commitment principle to some ε-stratification of the infinitesimals, we obtain the **LCD** system.

Existing Systems	**Main principle**	**εbf Systems**	**Main principle**
P	ε-probabilities	$\mid\!\sim_{bf}$	ε-belief and auto-deduction
Z	tolerance	$\mid\!\sim_{lc}$	+ ε-least commitment on Δ
---	-----	$\mid\!\sim_{\oplus}$	+ ε-least commitment on $d \in \Delta$ and Dempster ⊕ and ...
penalty	paying for unsatisfied defaults	$\mid\!\sim_{\oplus 1}$	$\varepsilon_d = \delta^i$ for $d \in \Delta_i$, $\delta \in \mathcal{E}^0$
lexicographic	equal weight within same strata	$\mid\!\sim_{\oplus 2}$	$\prod_{j=1}^{i-1} \delta j \lvert\Delta j\rvert >_\infty \delta i$, for i = 2 ... n+1
Brewka	preferred sub-theories	$\mid\!\sim_{\oplus 3}$	$\prod_{j=1}^{i-1} \prod_{h=1}^{\lvert\Delta_j\rvert} \varepsilon_{jh} >_\infty \varepsilon_{is}$ for all s=1... $\lvert\Delta_i\rvert$,
Geffner	conditional entailment	$\mid\!\sim_{\oplus 4}$	$\prod_{d' : \varepsilon_{d'} >_\infty \varepsilon_d} \varepsilon_{d'} >_\infty \varepsilon_d$ for all $d \in \Delta$,
---	-----	$\mid\!\sim_{lcd}$	least committed ε-stratification

Table 3: Properties underlying the systems analyzed in this paper

It is important to notice that, in our use of the theory of belief functions, we have actually employed only two mechanisms which are peculiar to this theory: the least-commitment principle, as a way to select minimally informative models; and Dempster’s rule of combination, as a way to aggregate (default) information. Interestingly, we have *not* used numerical values. This gives our treatment a qualitative flavor, and frees us from the delicate problem of having to justify the origin of the numbers that would otherwise be attached to default rules.

	$\mid\!\sim_{bf}$ **P**	$\mid\!\sim_{lc}$ **Z**	$\mid\!\sim_{\oplus 1}$ penalty	$\mid\!\sim_{\oplus 2}$ lex	$\mid\!\sim_{\oplus 3}$ Brewka	$\mid\!\sim_{\oplus 4}$ Geffner	$\mid\!\sim_{lcd}$
KLM rationality postulates	√	√	√	√	√	√	√
No irrelevance		√	√	√	√	√	√
No inheritance block			√	√	√	√	√
Preserve ambiguity	√					√	√
Indep. on # repetitions	√	√			√	√	√
Indep. on # supports	√	√			√	√	√
Rational monotonicity		√	√	√			

Table 2: The εbf-based systems and other non-monotonic systems.

deal with conditional assertions in the sense of Adams (1975) or Lehmann and Magidor (1992). This means that the two approaches can be viewed, in some sense, as complementary.

9 Conclusions

We have shown that we can use (a special class of) belief functions to give semantics to default rules, and to define several notions of non-monotonic consequence. Table 2 summarizes our results. The first row lists the seven systems presented in Section 3 ($\mid\!\sim_{bf}$), Section 4 ($\mid\!\sim_{lc}$), Section 5 ($\mid\!\sim_{\oplus 1}$, $\mid\!\sim_{\oplus 2}$, $\mid\!\sim_{\oplus 3}$ and $\mid\!\sim_{\oplus 4}$), and Section 6 ($\mid\!\sim_{lcd}$), respectively. The second row lists existing systems for which we have proved equivalence with one of our systems. To this respect, the use of ε-belief functions and of Dempster's rule of combination can be seen as a uniform framework to define non-monotonic consequence relations. This framework allowed us to capture several of the systems proposed in the literature; we speculate that it can also be used to model other non-monotonic systems.

The next rows in the table refer to the properties commonly regarded as desirable for a non-monotonic consequence relation. **LCD** satisfies all of these properties except rational monotony. More specifically, it satisfies the rationality postulates of Kraus, Lehmann and Magidor (Theorem 9); it correctly addresses the problems of irrelevance (Example 4 and Lemma 15), of blocking of inheritance (Example 6 and Lemma 16), of ambiguity (Example 8), and of redundancy (Example 9); and it is not sensitive to the number of different rules in a default base that support the same conclusion (Example 10). This good behavior is shared by Geffner's conditional entailment; however, **LCD** may give more intuitively acceptable conclusions in some cases (Example 12).

be well addressed, and considering all possible admissible assignments would result in inferences which are too cautious.

Perhaps the closest relative of **LCD** is the system proposed by Wilson (1992; 1993). Wilson also uses belief functions to give semantics to default reasoning, but he is mainly interested in establishing a strong link between belief functions and **DL.** Very roughly, Wilson's approach can be summarized in three steps as follows (see Wilson, 1993, for more details). The first step is to provide a new definition of what constitutes an extension (called M-extension) of a default theory. Given a default theory (Δ,W) on a language $\mathfrak{L}$ and a sub-set $\Delta' \subseteq \Delta$, Wilson first defines $Cn_{\Delta'}(W)$ as the intersection off all $\Gamma \subseteq \mathfrak{L}$ such that: (i) $W \subseteq \Gamma$, (ii) $Cn(\Gamma) = \Gamma$, and (iii) if $\alpha : \beta / \gamma \in \Delta'$ and $\alpha \in \Gamma$ then $\gamma \in \Gamma$. Next, he defines a sub-base Δ' to be Δ-consistent iff for all $\alpha : \beta / \gamma \in \Delta'$, $\neg\beta \notin Cn_{\Delta'}(W)$. Finally, a set of closed formulas E is said to be an M-extension of (Δ,W) iff there exists a maximally Δ-consistent sub-base Δ' such that $E = Cn_{\Delta'}(W)$. Wilson (1993) shows that each default theory has an M-extension.

The second step in Wilson's construction is to introduce the notion of B-extensions in what he calls the Sources of Evidence Framework (Wilson, 1992). In this framework, we consider a number of sources S_i, each of which gives us a piece of evidence. A source and its evidence corresponds to a simple support function. Let $\underline{a} = (a_1, \ldots, a_m)$ be the vector of real numbers representing the reliabilities of the sources, and let $X = \{x_I : I \subseteq \{1,\ldots,m\}\}$ where x_I represents the event that the sources S_i are reliable for $i \in I$ and the others are unreliable. An SE-structure is defined to be a function $K: X \rightarrow 2^{\mathfrak{L}}$ such that for $I,J \subseteq \{1,\ldots,m\}$ if $J \subseteq I$ then $K(x_J)$ is consistent whenever $K(x_I)$ is consistent. He then defines a probability measure on X, in a similar way as our equation (7, section 5.1). Finally, he defines a B-extension of a given SE-structure K as, roughly speaking, the set of propositions whose beliefs tends to 1, when the reliabilities a_i are pushed in such a way that the belief of all formulas tend to either 0 or 1.

The last step in Wilson's proposal is to show how M-extensions of a default theory can be captured in evidence theory via the Sources of Evidence Framework. The idea is to view each default rule $\alpha_i : \beta_i / \gamma_i$ as a source of information S_i which provides two pieces of information $\alpha_i : \gamma_i / \gamma_i$ and $\neg\beta_i : \perp / \perp$. He shows that a set of closed formulas E is an M-extension of Δ iff E is a B-extension of the SE-structure K^{Δ} defined by for $I \subseteq \{1,\ldots,m\}$, $K^{\Delta}(x_I) = Cn_{\Delta'}(W)$ where $\Delta' = \{d_i \in \Delta, i \in I\}$.

The main common point between Wilson's approach and our proposal is that we both represent default rules as limits of belief functions. Moreover the two approaches view each default rule as being one item of evidence provided by one of several sources of information. However, the two approaches mainly differ in the fact that Wilson's approach captures a variant of Reiter's default logic while in our approach we recover (and also propose) several recent systems which

Proof. As discussed above, in the penguin triangle **LCD** produces a deduction that is not in **DL**. Example 13 shows a case where the opposite is true. ■

LCD is also different from the maximum entropy approach for default reasoning proposed in (Goldszmidt, 1992; Goldszmidt et al., 1993). Let A_ε be the class of probability distributions considered in Section 2.2., namely the class of probability distributions, which are compatible with Δ. The idea in the maximum entropy approach is that, instead of considering all the probability distributions in A_ε as it is done in System **P**, we select one probability distribution, denoted by P^*, which maximizes the following function:

$$H(P) = -\sum P(\omega) \, Log(P(\omega)).$$

Goldszmidt (1992) proposes a semantical algorithm in the same spirit of system **Z** to compute the inference relation based on P^*. The maximal entropy approach partially solves the irrelevance and blocking property inheritance problems. However, it does not solve the ambiguity problem presented in Example 8. Moreover, the algorithm proposed for the maximum entropy approach only deals with default bases which are minimal core sets defined by:

Definition 13: A set Δ is said to be a "minimal-core set" if for each default $\alpha\rightarrow\beta$ of Δ, $\alpha\wedge\neg\beta\wedge\phi_{\Delta-\{\alpha\rightarrow\beta\}}$ is consistent (in the sense of classical logic).

The default bases that contain redundant information are not minimal core sets.

Bourne and Parsons (1999) have solved the problem of restricting to minimal core sets, by adding explicit variable strengths on each default. This extension is very close to System **Z+** (Goldszmidt and Pearl, 1991; 1996). In fact in both systems, the strength *s* associated to a default $p\rightarrow q$ corresponds to a constraint on the admissible rankings defined on the set of interpretations. This constraint says that the rank associated to the best world satisfying $p\wedge q$ should be smaller by at least *s* than the best world satisfying $p\wedge\neg q$ (Note that both systems use Spohn's kappa functions to rank order interpretations, so interpretations with low ranks are preferred). The difference between System **Z+** and Bourne and Parsons' proposal is in the definition of the rank associated to interpretations: one uses the maximum operator while the other uses addition.

Bourne and Parsons approach is also incomparable to **LCD**. The strength associated to defaults in Bourne and Parsons' system are set by the user, and there are no constraints between defaults. The set of plausible conclusions in their system therefore depends on the choice of those strengths. Notice that the "naive" least committed solution of assigning the same rank to all defaults would produce some undesirable effects in this system: redundancies in the default base would not where d_i denotes the i-th default in Δ, is admissible for Δ. Let then $\omega = p\wedge\neg mb\wedge fb\wedge f$ and $\omega' = p\wedge\neg mb\wedge fb\wedge\neg f$. Both interpretations satisfy d_1 and d_2, but while ω satisfies d_3 and falsifies d_4, ω' satisfies d_4 and falsifies d_3. As $d_3 >_\Delta d_4$, then we have $\omega >_\Omega \omega'$. But f is true in ω, and hence $\neg f$ cannot be deduced in the given prioritized structure.

Let us now see what happens with **LCD**. It is easy to see that the least commitment principle gives the partition: $\xi = \{\{\varepsilon_2, \varepsilon_3\}, \{\varepsilon_1, \varepsilon_4\}\}$. We have

$$pl(p\wedge\neg mb\wedge fb\wedge f) \approx_\infty \varepsilon_4, \text{ and}$$
$$pl(p\wedge\neg mb\wedge fb\wedge\neg f) \approx_\infty \varepsilon_3.$$

Since ξ sanctions that $\varepsilon_3 >_\infty \varepsilon_4$, then $p\wedge\neg mb\wedge fb \mathrel{|\!\!\sim}_{lcd} \neg f$ as desired. (Note that this result is also provided by **Z** and the other systems above; the reason why Geffner's system is more cautious in this example is that there are several admissible orders.) ■

We now briefly compare **LCD** with Reiter's default logic **DL**. We only consider normal default logic since in our approach we do not consider default rules of the general form $\alpha :\beta \,/\, \gamma$. The following example shows that **LCD** and **DL** have different behaviors when confronted with an inconsistent set of defaults. Inconsistency here is understood in the sense of Pearl (1988) (or Adams).

Example 13. (Makinson, 1989) Let $\Delta = \{T \to x, x\vee y \to \neg x\}$ and $W = \varnothing$. This default base is classically inconsistent. With **LCD**, we cannot work with this base, as the constraints are $\varepsilon_1 >_\infty \varepsilon_2$ and $\varepsilon_2 >_\infty \varepsilon_1$ and no partition satisfies them. In **DL**, on the contrary, we have one extension, from which x is obtained. ■

DL does not use a specificity criterion to prefer one extension over the other. In the penguin example, with $\Delta = \{p\to b, b\to f, p\to\neg f\}$, **DL** does neither infer f nor $\neg f$ from p. Thus, **DL** does not satisfy the auto-deductivity principle. Variants of **DL** have been proposed that rectify this shortcoming. For instance, Delgrande and Schaub (1994), inspired by Reiter and Criscuolo (1981), transform rules whose antecedents are in a general class into semi-normal defaults, while leaving specific rules unchanged, and use semi-normal default logic to make inferences. This extension remains more cautious than **LCD** in some cases. For example, let $\Delta = \{x\to y\}$, $W = \{\neg y\}$, and suppose $\{\neg y, \neg x\vee y\}$ is consistent; then, we do not get $\neg x$ in Delgrande and Schaub's system, while $\neg x$ follows from $\neg y$ in **LCD**. The following theorem summarizes the relation between **DL** and **LCD**.

Theorem 10. **LCD** consequence is incomparable with **DL**.

terms of $\mathcal{E}$ (i.e., products of elements) which are induced by $\mathbf{C}_\Delta$. The previous example shows that the ordering constraints between terms cannot be disposed of.

It is not easy to say which one of the two conclusions c or ¬c is more intuitive in the last example. This mainly depends on the dependence relation between the two properties a and c when y is true. **LCD** regards these two properties as dependent referring to the second rule of Δ. In a sense, mentioning both properties in the same rule is taken as a sign that if y is exceptional for one property, then this is also (plausibly) the case for the other property, and conversely. Note that if we split d_2 into two rules y→c and y→¬a then **LCD** infers ¬c. Intuitively, the two properties a and c are now considered independent: if y is exceptional with respect to one property, it is not necessarily so for the second property as well.

At first sight, it might seem that the solution provided by **LCD** violates the specificity requirement, since d_3 is the most specific rule in Δ which applies to $y \wedge s \wedge \neg a$, and d_3 entails ¬c. The situation, however, is more complex. Consider the new base $\Delta^* = \Delta \cup \{d_4\}$, where $d_4 = y \wedge \neg a \rightarrow \neg a \wedge c$. Δ and Δ* deduce the same properties about c in **LCD**, since the least commitment principle gives for Δ* the partition $\xi = \{\{\varepsilon_1, \varepsilon_3, \varepsilon_4\}, \{\varepsilon_2\}\}$. However, the specificity principle cannot help us in deciding between c and ¬c in this case: d_3 would lead to ¬c, while d_4 to c, and each one of d_3 and d_4 can be considered as the 'most specific rule' for $y \wedge s \wedge \neg a$. In fact, the specificity principle is clearly defined for antecedents and consequents that are literals (Touretsky, 1984) which is not the case in this example.

The intuitive acceptability of the deductions performed is an important criterion to judge a formal system for commonsense reasoning. The examples in the last sections have shown cases where the conclusion provided by **LCD** is more intuitively acceptable that those provided by the **Z**, Brewka's, Pinkas', and lexicographic systems. We now show a case where **LCD** allows to infer a desirable conclusion that is not provided by Geffner's conditional entailment.

Example 12. Consider the set of defaults: $\Delta = \{p \rightarrow mb \vee fb,\ mb \rightarrow f,\ fb \rightarrow f,\ p \rightarrow \neg f\}$, where p, mb, fb, and f respectively stand for <u>p</u>enguin, <u>m</u>ale <u>b</u>ird, <u>f</u>emale <u>b</u>ird, and <u>f</u>lies. Let then $\mathcal{E} = \{\varepsilon_1, \varepsilon_2, \varepsilon_3, \varepsilon_4\}$ be the set of associated infinitesimals. We wonder whether Tweety, who is a penguin and a female bird but not a male bird, can fly or not — clearly, we expect her not to fly.

Geffner's conditional entailment does not infer that "Tweety does not fly". To show this, it is enough to find an admissible prioritized structure where the conclusion does not follow. Note that the preference relation

$$d_1 >_\Delta d_3 >_\Delta d_4 >_\Delta d_2,$$

Example 11. Consider the set of defaults: $\Delta = \{T \rightarrow a,\ y \rightarrow \neg a \wedge c,\ y \wedge s \rightarrow \neg c\}$ where T denotes the tautology, and let $\mathcal{E} = \{\varepsilon_1, \varepsilon_2, \varepsilon_3\}$ be the set of associated infinitesimals. Both System **Z** and all of the systems considered in Section 5 stratify Δ by considering that $d_1 = T \rightarrow a$ is less prioritary than $d_2 = y \rightarrow \neg a \wedge c$, and that the latter is less prioritary than $d_3 = y \wedge s \rightarrow \neg c$. That is, we have the stratification $S = \{\{d_1\}, \{d_2\}, \{d_3\}\}$. In the case of **LCD**, we have the following constraints:

$\mathbf{C_1}$: $\max\{pl(\omega)|\omega \models a\} >_\infty \max\{pl(\omega)|\omega \models \neg a\}$ i.e., $1 >_\infty \varepsilon_1$
$\mathbf{C_2}$: $\max\{pl(\omega)|\omega \models y \wedge \neg a \wedge c\} >_\infty \max\{pl(\omega)|\omega \models y \wedge (a \vee \neg c)\}$ i.e., $\varepsilon_1 >_\infty \varepsilon_2$
$\mathbf{C_3}$: $\max\{pl(\omega)|\omega \models y \wedge s \wedge \neg c\} >_\infty \max\{pl(\omega)|\omega \models y \wedge s \wedge c\}$ i.e., $\varepsilon_2 >_\infty \max\{\varepsilon_3\varepsilon_1, \varepsilon_3\varepsilon_2\}$

By virtue of $\mathbf{C_2}$ the last constraint simplifies to

$\mathbf{C'_3}$: $\varepsilon_2 >_\infty \varepsilon_3\varepsilon_1$

Then, the least commitment principle gives the partition $\xi = \{\{\varepsilon_1, \varepsilon_3\}, \{\varepsilon_2\}\}$. This solution corresponds to the stratification $S' = \{\{d_1, d_3\}, \{d_2\}\}$, which is different from S. ■

An immediate consequence of this difference is that **LCD** may produce results that are different from those of all the other systems. Consider the question of deciding whether or not c follows from $y \wedge s \wedge \neg a$ given the Δ in the example. All of $\mid\!\sim_{pen}$, $\mid\!\sim_G$, $\mid\!\sim_{lex}$ and $\mid\!\sim_B$ can infer $\neg c$ from $y \wedge s \wedge \neg a$. Roughly, in the presence of the fact $y \wedge s \wedge \neg a$, we apply the rule which is most prioritary according to S, that is d_3, and hence we infer $\neg c$. By contrast, in **LCD** we need to compare

$pl(y \wedge s \wedge \neg a \wedge c) \approx_\infty \varepsilon_1\, \varepsilon_3$, and
$pl(y \wedge s \wedge \neg a \wedge \neg c) \approx_\infty \varepsilon_1\, \varepsilon_2$.

Given the least committed partition above, and by the properties of infinitesimals, we get $\varepsilon_1\, \varepsilon_3 >_\infty \varepsilon_1\, \varepsilon_2$, and then $y \wedge s \wedge \neg a \mid\!\sim_{lcd} c$. This immediately gives us the following.

Theorem 9. The **LCD** consequence is incomparable to all of: Pearl's system **Z**, Brewka's preferred sub-theories, Geffner's conditional entailment, Pinkas' penalty logic, and the lexicographic approach.

<u>Proof.</u> Examples 5 and 6 provide counter-examples to both inclusions between **Z** and **LCD**. The last example above shows incomparability between **LCD** and the other systems. ■

The last example also shows an important fact, which did not show up in our previous examples: the stratification of $\mathcal{E}$ *alone* does not guarantee that the $\mathbf{C_\Delta}$ constraints will be satisfied. For instance, $\mathbf{C_2}$ is not satisfied if we only use the stratification given by the least committed principle. As we have discussed above, LCD-entailment also considers the $>_\infty$ constraints between

ε_1 and ε_2, behave as one default $d_3 = \alpha \rightarrow \beta$ with the associated infinitesimal $\varepsilon_3 = \varepsilon_1\varepsilon_2$: in the auto-deduction constraints, ε_1 and ε_2 will always appear or not appear simultaneously in each term, and can then be uniformly replaced by ε_3.

The following example shows that **LCD** consequence is not sensitive to the number of different rules in a default base that support the same conclusion

Example 10. Consider the Quaker-Republican problem with the extra rule "Generally, ecologists are pacifist": $\Delta = \{q \rightarrow p, r \rightarrow \neg p, e \rightarrow p\}$. The constraints are

$$\begin{array}{lll} \mathbf{C_1}: & \max\{pl(\omega)|\omega \vDash q \wedge p\} >_\infty \max\{pl(\omega)|\omega \vDash q \wedge \neg p\} & \text{i.e., } 1 >_\infty \varepsilon_1 \\ \mathbf{C_2}: & \max\{pl(\omega)|\omega \vDash r \wedge \neg p\} >_\infty \max\{pl(\omega)|\omega \vDash r \wedge p\} & \text{i.e., } 1 >_\infty \varepsilon_2 \\ \mathbf{C_3}: & \max\{pl(\omega)|\omega \vDash e \wedge p\} >_\infty \max\{pl(\omega)|\omega \vDash e \wedge \neg p\} & \text{i.e., } 1 >_\infty \varepsilon_3 \end{array}$$

and we get again the one class least committed partition $\xi = \{\{\varepsilon_1, \varepsilon_2, \varepsilon_3\}\}$. Then,

$$pl(q \wedge e \wedge r \wedge p) \approx_\infty \varepsilon_2 \text{ and } pl(q \wedge e \wedge r \wedge \neg p) \approx_\infty \varepsilon_1\varepsilon_3.$$

Since ε_2 is not comparable with $\varepsilon_1\varepsilon_3$ according to ξ, **LCD** does not deduce that ecologists who are both Quakers and republicans are pacifist, nor that they are not pacifist. ■

8 LCD and the other systems

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

Theorem 8. For a given Δ, if $\alpha \mid\!\sim_{\mathbf{P}} \beta$ then $\alpha \mid\!\sim_{\mathrm{lcd}} \beta$. The converse is not true.

<u>Proof.</u> Lemma 14 proves the inclusion. The examples given in Section 7.2 show that the inclusion is proper. ■

We now study the relation between **LCD** and the systems considered in Section 5. All these systems include a stratification of the default base that essentially corresponds to the one used in System **Z**. What makes the systems different is the way in which this stratification is used to decide entailment. **LCD** also gives rise to a stratification of the rules in the base, which is directly obtained from the partition of the corresponding infinitesimals: for example, the partition $\{\{\varepsilon_2, \varepsilon_3\}, \{\varepsilon_1, \varepsilon_4\}\}$ of $\mathcal{E}$ corresponds to the stratification $\{\{d_2, d_3\}, \{d_1, d_4\}\}$ of Δ. However, and importantly, the stratification produced by **LCD** is not necessarily the same as the one produced by System **Z**. The following example supports this claim.

We get the same partition as in Example 6: $\xi = \{\{\varepsilon_1, \varepsilon_4\}, \{\varepsilon_2, \varepsilon_3\}\}$. Consider now a bird that is a penguin and has metal wings (*sic*). Given the base Δ, we should not be able to say whether or not this beast will fly — we are in a case of ambiguity. We have indeed:

$$pl(b\wedge p\wedge m\wedge f) \approx_\infty \varepsilon_2$$
$$pl(b\wedge p\wedge m\wedge\neg f) \approx_\infty \varepsilon_1\varepsilon_4.$$

As the ξ partition says nothing about the relative magnitude of ε_2 and $\varepsilon_1\varepsilon_4$, we 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 $b\wedge p\wedge m \mid\!\sim_Z \neg f$. ■

7.5. *Syntax sensitivity*

The last desideratum in our list is syntax-independence. Some of the existing systems that go beyond system **P** do not satisfy this requirement: for instance, in the lexicographic approach, discussed in Section 5.3, repetitions of the same default in Δ may change the result. The following example shows that **LCD** is not sensitive to these duplications.

Example 9. 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 the lexicographic approach, we would have $q\wedge r \mid\!\sim p$, while we would prefer to acknowledge the ambiguity and deduce nothing. In **LCD**, we have the constraints

$$\mathbf{C_1}:\quad \max\{pl(\omega)|\omega\models q\wedge p\} >_\infty \max\{pl(\omega)|\omega\models q\wedge\neg p\} \qquad \text{i.e., } 1 >_\infty \varepsilon_1\varepsilon_2$$
$$\mathbf{C_2}:\quad \max\{pl(\omega)|\omega\models q\wedge p\} >_\infty \max\{pl(\omega)|\omega\models q\wedge\neg p\} \qquad \text{i.e., } 1 >_\infty \varepsilon_1\varepsilon_2$$
$$\mathbf{C_3}:\quad \max\{pl(\omega)|\omega\models r\wedge\neg p\} >_\infty \max\{pl(\omega)|\omega\models r\wedge p\} \qquad \text{i.e., } 1 >_\infty \varepsilon_3$$

Using the least commitment principle, we get the one class partition $\{\{\varepsilon_1, \varepsilon_2, \varepsilon_3\}\}$. Then

$$pl(q\wedge r\wedge p) \approx_\infty \varepsilon_3,\ pl(q\wedge r\wedge\neg p) \approx_\infty \varepsilon_1\varepsilon_2,$$

and we have neither $q\wedge r \mid\!\sim_{lcd} p$ nor $q\wedge r \mid\!\sim_{lcd} \neg p$, as desired. It is important to note that, if we considered infinitesimals in the same class as being equivalent ($\approx_\infty$) rather than unconstrained ($\sim_\infty$), then the ambiguity would not be preserved, and our consequence relation would be sensitive to duplications — in fact, we would have $\varepsilon_3 >_\infty \varepsilon_1\varepsilon_2$ by Lemma A8(c) and Lemma A9(b), and hence $q\wedge r \mid\!\sim_{lcd} p$. ■

It is easy to realize that **LCD** is not sensitive to duplications in general. In fact, any two instances of the same default, say $d_1 = \alpha\rightarrow\beta$ and $d_2 = \alpha\rightarrow\beta$, with the associated infinitesimals

Example 7. Consider the defaults "Generally, Swedes are blond" and "Generally, Swedes are tall", represented by $\Delta = \{s \rightarrow b,\ s \rightarrow t\}$. What can be said about short Swedish? It seems reasonable to expect that they are still blond. The two auto-deductions constraints are $1 >_\infty \varepsilon_1$ and $1 >_\infty \varepsilon_2$, which are always satisfied by Lemma A7(b). Hence, all terms are free, and by the least commitment principle we put them all in a single class; that is, we have a single least committed partition: $\xi = \{\{\varepsilon_1, \varepsilon_2\}\}$. For any bel in $Bel_{lcd}(\Delta)$,

$$pl(s \wedge \neg t \wedge b) \approx_\infty \varepsilon_1,\ pl(s \wedge \neg t \wedge \neg b) \approx_\infty \varepsilon_1 \varepsilon_2,$$

and since $\varepsilon_1 >_\infty \varepsilon_1\varepsilon_2$ by Lemma A8(c), we conclude that short Swedes are blond, i.e., $s \wedge \neg t \mathrel{|\!\!\sim}_{lcd} b$. Notice that this result does not follow in **Z**. ■

We can generalize the above examples, and show that, in a given context α, all the defaults in the subset $Free(\Delta \cup \{\alpha\})$ of Δ which is not responsible for the inconsistency can be used in the inference process. The subset $Free(\Delta \cup \{\alpha\})$ is unique and is formally defined in the following way:

$$Free(\Delta \cup \{\alpha\}) = \{\beta \rightarrow \gamma \in \Delta \ / \ \nexists A \subseteq \Delta,\ A \text{ minimally inconsistent with } \alpha, \text{ and } \beta \rightarrow \gamma \in A\}$$

where $A \subseteq \Delta$ is minimally inconsistent with α iff:

- $\alpha \wedge \phi_A \vdash \bot$;
- $\forall \beta \rightarrow \gamma \in A,\ \alpha \wedge \phi_{A-\{\beta \rightarrow \gamma\}} \nvdash \bot$.

Lemma 16. If $\alpha \wedge Free(\Delta \cup \{\alpha\}) \vdash \beta$ then $\alpha \mathrel{|\!\!\sim}_{lcd} \beta$.

Note that Lemma 16 does not hold for System **Z**: Example 6 above provides a simple counter-example (where $b \rightarrow l$ belongs to $Free(\Delta \cup \{p\})$).

7.4. Ambiguity preservation

The next desideratum listed in the introduction is 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 8. Let $\Delta = \{b \rightarrow f,\ p \rightarrow \neg f,\ p \rightarrow b,\ m \rightarrow f\}$, where the last default means "Generally, objects with <u>m</u>etal-wings fly". The constraints that the elements of $Bel_{lcd}(\Delta)$ must satisfy are the three in Example 4, plus

$$\mathbf{C_4}: \quad \max\{pl(\omega) | \omega \vDash m \wedge f\} >_\infty \max\{pl(\omega) | \omega \vDash m \wedge \neg f\} \qquad \text{i.e., } 1 >_\infty \varepsilon_4$$

We can prove that **LCD** correctly addresses the irrelevance problem in general. Let In(Δ) be the set of propositional symbols which appear in Δ and Out(Δ) = V - In(Δ) be the set of propositional symbols which do not appear in Δ, where V denotes the set of all propositional symbols of the language. In the example 4, we have V={p,b,f,r}, In(Δ)={p,b,f}, Out(Δ)={r}. Let $\mathcal{L}_{Out(\Delta)}$ be the set of all propositional formulas composed of propositional symbols which do not appear in Δ, and let $\mathcal{L}_{in(\Delta)}$ be the set of all propositional formulas composed of propositional symbols which appear in Δ.

Lemma 15. Let $\delta \in \mathcal{L}_{out(\Delta)}$ and $\alpha, \beta \in \mathcal{L}_{in(\Delta)}$. Then $\alpha \mid\!\sim_{lcd} \beta$ implies $\alpha \wedge \delta \mid\!\sim_{lcd} \beta$.

7.3. Blocking of inheritance

Several systems, including Pearl's system **Z**, suffer from the problem of inheritance blocking: if a class C contains a property p that conflicts with the one inherited from a super-class C', than no property at all is inherited from C' (even properties unrelated to p). The canonical example is built by adding to the usual penguin problem the default b → l (read "generally, birds have legs"). From this, system **Z** cannot deduce that penguins have legs too, i.e., $p \not\mid\!\sim_{\mathbf{Z}} l$.[12] The next two examples show that **LCD** does not suffer from this problem.

Example 6. Let Δ = {b→f, p→¬f, p→b, b→l}, and let $\mathcal{E} = \{\varepsilon_1, \varepsilon_2, \varepsilon_3, \varepsilon_4\}$ be the associated infinitesimals. Let bel be any element of $Bel_{lcd}(\Delta)$. The constraints that bel must satisfy are the same as in Example 4, plus

$$\mathbf{C_4}: \quad \max\{pl(\omega) | \omega \models b \wedge l\} >_\infty \max\{pl(\omega) | \omega \models b \wedge \neg l\} \qquad \text{i.e., } 1 >_\infty \varepsilon_4$$

It is easy to verify that the following is the only minimally committed partition satisfying these constraints: $\xi = \{\{\varepsilon_1, \varepsilon_4\}, \{\varepsilon_2, \varepsilon_3\}\}$. To see if penguins have legs, note that

$$\begin{aligned} pl(p \wedge l) &\approx_\infty \max\{pl(p \wedge l \wedge b \wedge f), pl(p \wedge l \wedge b \wedge \neg f), pl(p \wedge l \wedge \neg b \wedge f), pl(p \wedge l \wedge \neg b \wedge \neg f)\} \\ &\approx_\infty \max\{\varepsilon_2, \varepsilon_1, \varepsilon_2\varepsilon_3, \varepsilon_2\} \approx_\infty \varepsilon_1 \end{aligned}$$

$$\begin{aligned} pl(p \wedge \neg l) &\approx_\infty \max\{pl(p \wedge \neg l \wedge b \wedge f), pl(p \wedge \neg l \wedge b \wedge \neg f), pl(p \wedge \neg l \wedge \neg b \wedge f), pl(p \wedge \neg l \wedge \neg b \wedge \neg f)\} \\ &\approx_\infty \max\{\varepsilon_2\varepsilon_4, \varepsilon_1\varepsilon_4, \varepsilon_2\varepsilon_3\varepsilon_4, \varepsilon_2\varepsilon_4\} \approx_\infty \varepsilon_1\varepsilon_4 \end{aligned}$$

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

---

[12] Goldszmidt and Pearl (1991) have suggested an extension of **Z**, called $\mathbf{Z}^+$, which correctly handles this example. Unfortunately, $\mathbf{Z}^+$ does not solve the problem of inheritance blocking in general: if we add the rule "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 does not appear in **LCD**.

and $\varepsilon_2 \sim_\infty \varepsilon_1\varepsilon_4\varepsilon_5$. ■

The failure of **LCD** to satisfy rational monotonicity should not be seen as a drawback, but as an indication that rational monotonicity does not necessarily apply to all situations. The last example provides a case where the conclusions produced by **LCD** could be regarded as more intuitively acceptable than those produced by accepting rational monotonicity. To see why, consider the context "$b \wedge p \wedge m$" — that is, consider an individual who is both b, p and m. It is reasonable to expect that neither y nor $\neg y$ can be derived in this context. Similarly, it is reasonable to expect neither f nor $\neg f$ to be derived. However, we should expect to have $\neg y \vee f$. In fact, either the individual flies or it does not. If it does, we have trivially $\neg y \vee f$; and if it does not, then it cannot be a Yogi, and we again have $\neg y \vee f$. As we have seen, all these results are indeed produced by **LCD**.[11] Consider now the extended context "$b \wedge p \wedge m \wedge y$", where we also know that the individual is a Yogi. **LCD** does not produce the inference $b \wedge p \wedge m \wedge y \mid\!\sim \neg y \vee f$, thus violating rational monotonicity. But this is again a reasonable behavior, since we have y in the premises, and we have no reason to infer f. Interestingly, the use of rational monotonicity in this example is even more questionable if we use it in combination with the other rules of system **P**. In fact, we have (by Reflexivity and Right Weakening) $b \wedge p \wedge m \wedge y \mid\!\sim y$; and from this and from $b \wedge p \wedge m \wedge y \mid\!\sim \neg y \vee f$ we have (by And) $b \wedge p \wedge m \wedge y \mid\!\sim f$, whose intuitive validity can be easily challenged.

7.2. Irrelevance

The following example shows that **LCD** correctly addresses the irrelevance problem.

Example 4 (continued). Let Δ be as above, and consider a new property "red" (r) unrelated to b, p and f. We expect that red birds fly (recall that this is not the case in system **P**, that is, $b \wedge r \not\mid\!\sim_{\mathbf{P}} f$). For any bel in Bel_{lcd}, we have (we apply Lemma 3, Lemma 9, and Lemma A12)

$$pl(b \wedge r \wedge f) \approx_\infty \max\{pl(b \wedge r \wedge f \wedge p), pl(b \wedge r \wedge f \wedge \neg p)\} \approx_\infty \max\{\varepsilon_2, 1\} \approx_\infty 1$$
$$pl(b \wedge r \wedge \neg f) \approx_\infty \max\{pl(b \wedge r \wedge \neg f \wedge p), pl(b \wedge r \wedge \neg f \wedge \neg p)\} \approx_\infty \max\{\varepsilon_1, \varepsilon_1\} \approx_\infty \varepsilon_1$$

As any εbf must satisfy $1 >_\infty \varepsilon_1$ by Lemma A7(b), we have $pl(b \wedge r \wedge f) >_\infty pl(b \wedge r \wedge \neg f)$, which implies $b \wedge r \mid\!\sim_{lcd} f$ as desired. ■

---

[11] These arguments depend on the assumed dependence relation between p and y. For instance if y and p denote equivalent propositions then inferring $\neg y \vee f$ is controversial. **LCD** regards propositions for which there is not explicitly stated relation in the database as being independent — this is a common implicit assumptions in logical systems. In our case, all we know about y is that it flies; as we have no information to infer a relation between p and y, we assume they are independent. Thus, knowing that an individual does not fly, we plausibly prefer to deduce that it is not a y rather than an exceptional y.

$$\frac{\alpha \mathrel{|\!\sim} \beta \qquad \alpha \not\mathrel{|\!\sim} \neg\gamma}{\alpha\wedge\gamma \mathrel{|\!\sim} \beta}$$

This rule has been originally proposed by Lehmann (1989) as a candidate rule for minimizing the amount of information lost when we add a new consistent item of information γ to a pre-existing α. Although no definite reason has been given for the necessity of this rule in a non-monotonic reasoning system, rational monotonicity is usually regarded as desirable, and it is validated by many current extensions of system **P**; two notable exceptions are Brewka's system and Geffner's conditional entailment. **LCD** also fails to satisfy this rule, as shown by our next example.

Example 5. Let $\Delta = \{b\to f,\ p\to\neg f,\ p\to b,\ m\to f,\ y\to f\}$, where the first three defaults are as in Example 2, and the two last rules mean "Generally, objects with metal-wings fly" and "Generally, Yogis fly." The constraint induced by Δ are:

- **C1**: $\max\{t_\omega \mid \omega\models b\wedge f\} >_\infty \max\{t_\omega \mid \omega\models b\wedge\neg f\}$ i.e., $1 >_\infty \varepsilon_1$
- **C2**: $\max\{t_\omega \mid \omega\models p\wedge\neg f\} >_\infty \max\{t_\omega \mid \omega\models p\wedge f\}$ i.e., $\max\{\varepsilon_1, \varepsilon_3\} >_\infty \varepsilon_2$
- **C3**: $\max\{t_\omega \mid \omega\models p\wedge b\} >_\infty \max\{t_\omega \mid \omega\models p\wedge\neg b\}$ i.e., $\max\{\varepsilon_1, \varepsilon_2\} >_\infty \varepsilon_3$
- **C4**: $\max\{t_\omega \mid \omega\models m\wedge f\} >_\infty \max\{t_\omega \mid \omega\models m\wedge\neg f\}$ i.e., $1 >_\infty \varepsilon_4$
- **C5**: $\max\{t_\omega \mid \omega\models y\wedge f\} >_\infty \max\{t_\omega \mid \omega\models y\wedge\neg f\}$ i.e., $1 >_\infty \varepsilon_5$

There is one minimally committed partition that satisfies the auto-deduction constraints, given by: $\xi=\{\{\varepsilon_1, \varepsilon_4, \varepsilon_5\}, \{\varepsilon_2, \varepsilon_3\}\}$. From this, we see that :

(i) We have $b\wedge p\wedge m \mathrel{|\!\sim}_{lcd} \neg y\vee f$. In fact, for any bel in Bel_{lcd}:

$pl(b\wedge p\wedge m\wedge(\neg y\vee f)) \approx_\infty \max\{pl(b\wedge p\wedge m\wedge\neg y\wedge f), pl(b\wedge p\wedge m\wedge\neg y\wedge\neg f), pl(b\wedge p\wedge m\wedge y\wedge\neg f)\}$

$\approx_\infty \max\{\varepsilon_2, \varepsilon_1\varepsilon_4, \varepsilon_2\}$;

$pl(b\wedge p\wedge m\wedge y\wedge\neg f)) \approx_\infty \varepsilon_1\varepsilon_4\varepsilon_5$;

and $\varepsilon_1\varepsilon_4 >_\infty \varepsilon_1\varepsilon_4\varepsilon_5$.

(ii) We do *not* have $b\wedge p\wedge m \mathrel{|\!\sim}_{lcd} \neg y$. In fact, for any bel in Bel_{lcd}:
$pl_\xi(b\wedge p\wedge m\wedge\neg y) \approx_\infty \max\{pl(b\wedge p\wedge m\wedge\neg y\wedge f), pl(b\wedge p\wedge m\wedge\neg y\wedge\neg f)\} \approx_\infty \max\{\varepsilon_2, \varepsilon_1\varepsilon_4\}$;
$pl_\xi(b\wedge p\wedge m\wedge y) \approx_\infty \max\{pl(b\wedge p\wedge m\wedge y\wedge f), pl(b\wedge p\wedge m\wedge y\wedge\neg f)\} \approx_\infty \max\{\varepsilon_2, \varepsilon_1\varepsilon_4\varepsilon_5\}$;
and $\varepsilon_2 \sim_\infty \varepsilon_1\varepsilon_4$ since from the partition $\xi=\{\{\varepsilon_1, \varepsilon_4, \varepsilon_5\}, \{\varepsilon_2, \varepsilon_3\}\}$ we can neither deduce $\varepsilon_2 >_\infty \varepsilon_1\varepsilon_4$ nor $\varepsilon_1\varepsilon_4 >_\infty \varepsilon_2$.

(iii) From (i) and (ii), according to the rule of rational monotonicity we should have $b\wedge p\wedge m\wedge y \mathrel{|\!\sim}_{lcd} \neg y\vee f$. However, this is not the case. In fact, for any bel in Bel_{lcd}:

$pl(b\wedge p\wedge m\wedge y\wedge f) = \varepsilon_2$,

$pl(b\wedge p\wedge m\wedge y\wedge\neg f) = \varepsilon_1\varepsilon_4\varepsilon_5$, Concerning the first point, we emphasize that in our approach the constraints are directly given on infinitesimals, and therefore the number of variables involved in the constraints is at most equal to $|\Delta|$. This means that it is always possible to make explicit the optimal solutions since the size of the solution is not very large. In the preferential models-based approach to default reasoning, by contrast the constraints are often given on the set of interpretations, and the size of solutions is very large.

Concerning the second point, the LCD-entailment is defined at the semantical level, hence looking for the syntactic counterpart is an important point. One possible way to find syntactic counterparts is to explore the compilation techniques developed in (Benferhat et al., 1998). Their work provides a syntactic inference for reasoning from stratified knowledge bases Σ using utilitarist semantics (like the one based on penalty logic). The idea is to transform Σ into a new knowledge base Σ' such that if we apply the egalitarist semantics (like the one based on possibilistic logic) on Σ' we get the same results as if we apply the utilitarist semantics on Σ, and hence we can use the syntactic inference from the approaches based on egalitarist semantics. Exploring this possibility is part of our future work.

7 Analysis of the LCD consequence relation

We now study in more detail the patterns of reasoning which are captured by **LCD**. To do this, we consider the desiderata listed in the Introduction, and show how **LCD** addresses them.

7.1. Rationality and specificity

The fact that **LCD** respects the principle of specificity was illustrated by Example 4. It is easy to see that it also respects the KLM postulates of rationality (Kraus et al., 1990). In fact, all the elements of $Bel_{lcd}(\Delta)$ are εbf-model of Δ, then **LCD** relation is at least as strong as $\mid\sim_{bf}$, and therefore it satisfies the KLM properties by Theorem 1. We have thus proved the following.

Lemma 14. If $\alpha \mid\sim_{\mathbf{P}} \beta$ then $\alpha \mid\sim_{lcd} \beta$.

The converse of the previous lemma is false (see Section 7.2. for a counter-example). So, **LCD** is strictly stronger that **P**. The KLM postulates are commonly accepted as the minimal core of any system of non-monotonic inference. Another rule that has found wide, although not unanimous consensus, is the rule of rational monotonicity:

$\alpha \mathrel{|\!\sim}_{lcd} \beta$ iff for any $bel_{\mathcal{E}}$ in $Bel_{lcd}(\Delta)$, $bel_{\mathcal{E}} \models \alpha \rightarrow \beta$.

Example 4 (continued). It is immediate to verify that **LCD** correctly addresses the classical penguin problem. Let $bel_{\mathcal{E}}$ be any element of $Bel_{lcd}(\Delta)$. From Lemma 9, $pl_{\mathcal{E}}(b \wedge p \wedge f) \approx_\infty \varepsilon_2$ and $pl_{\mathcal{E}}(b \wedge p \wedge \neg f) \approx_\infty \varepsilon_1$. We have seen above that, for any $bel_{\mathcal{E}}$ in $Bel_{lcd}(\Delta)$, the set $\mathcal{E}$ must be such that $\varepsilon_1 >_\infty \varepsilon_2$. Hence, $b \wedge p \mathrel{|\!\sim}_{lcd} \neg f$ as expected. ■

In practice, the algorithm needed to determine if $\alpha \mathrel{|\!\sim}_{lcd} \beta$ requires only the knowledge of the $>_\infty$ relation between the ε_i's of $\mathcal{E}$ and some terms (products) built from these ε_i's. It consists in constructing the constraints $\mathbf{C}_\Delta$, building the set $\mathcal{P}$ of ε-stratifications compatible with $\mathbf{C}_\Delta$, finding in $\mathcal{P}$ the set $\mathcal{P}^*$ of minimally committed ε-stratifications, finding the $>_\infty$ inequalities that must be satisfied. Building explicitly the set of $bel_{\mathcal{E}}$ is not necessary.

For notation simplicity, we write the inequalities constraints that result from $\mathbf{C}_\Delta$ and the ε-stratification using ε_i symbols, with the understanding that they denote proto-infinitesimals to which infinitesimals will be assigned, infinitesimals that will satisfy all the inequalities.

When the $>_\infty$ inequalities cannot determine if $t_1 >_\infty t_2$ or $t_2 >_\infty t_1$ or $t_1 \approx_\infty t_2$, we denote this indeterminacy by $t_1 \sim_\infty t_2$, meaning that each of the three relations can be obtained from some interpretation used to build $Bel_{lcd}(\Delta)$.

An important property of **LCD** is that, if both Δ and the alphabet of $\mathcal{L}$ underlying Δ are finite, then **LCD**-consequence is decidable. To see this, first notice that if Δ is finite, so is $\mathbb{Z}$. We can then enumerate all the ε-stratifications of $\mathbb{Z}$, select those that satisfy the $\mathbf{C}_\Delta$ constraints (which are finite), and select among these the ones which are minimally committed. Second, as the number of propositional variables is finite, so is the number of possible truth assignments ω. Then, the set of all t_ω terms is also finite. Thus, we can enumerate all the possible $>_\infty$ orderings on this set, and select those that: (i) satisfy the $\mathbf{C}_\Delta$ constraints, and (ii) are consistent with some of the ε-stratifications selected above. Finally, in order to check if $\alpha \mathrel{|\!\sim}_{lcd} \beta$, we need to check that $\max_{\omega \models \alpha \wedge \beta} t_\omega >_\infty \max_{\omega \models \alpha \wedge \neg \beta} t_\omega$ holds for all the orderings between terms so selected. We have so proved the following:

Theorem 7. Let $\mathcal{L}$ be a propositional language on a finite alphabet, and Δ a default base on $\mathcal{L}$. Then, the $\mathrel{|\!\sim}_{lcd}$ relation is decidable.

The definition of effective technique to compute the set of LCD-consequences is left for further research. Here, we only give two brief comments on the computation of the least committed stratifications and the computation of the plausible consequences once the least committed stratifications are computed.

Note that any solution which satisfies $\mathbf{C_2}$ and $\mathbf{C_3}$ should be such that $\varepsilon_1 >_\infty \varepsilon_2$ and $\varepsilon_1 >_\infty \varepsilon_3$. Indeed, assume that $\varepsilon_2 \geq_\infty \varepsilon_1$ (resp. $\varepsilon_3 \geq_\infty \varepsilon_1$) then to satisfy $\mathbf{C_2}$ (resp. $\mathbf{C_3}$) we should have $\varepsilon_3 >_\infty \varepsilon_2$ (resp. $\varepsilon_2 >_\infty \varepsilon_3$), which means that we should have $\varepsilon_3 >_\infty \varepsilon_1$ (resp. $\varepsilon_2 >_\infty \varepsilon_1$), and this means the impossibility of satisfying $\mathbf{C_3}$ (resp. $\mathbf{C_2}$).

Therefore since any solution should be such that $\varepsilon_1 >_\infty \varepsilon_2$ and $\varepsilon_1 >_\infty \varepsilon_3$ the two constraints $\mathbf{C_2}$ and $\mathbf{C_3}$ are satisfied, and there is no constraints between ε_2 and ε_3. Hence, we only get three ε-stratifications of $\mathcal{Z}$ that are compatible with these constraints (depending on whether : $\varepsilon_2 >_\infty \varepsilon_3$, $\varepsilon_2 \approx_\infty \varepsilon_3$ or $\varepsilon_3 >_\infty \varepsilon_2$):

$\xi = \{\{\zeta_1\}, \{\zeta_2, \zeta_3\}\}$,

$\xi' = \{\{\zeta_1\}, \{\zeta_2\}, \{\zeta_3\}\}$,

$\xi'' = \{\{\zeta_1\}, \{\zeta_3\}, \{\zeta_2\}\}$.

It is easy to see that ξ is strictly less committed than both ξ' and ξ'' — intuitively, this is because ξ does not impose any unnecessary order between ζ_2 and ζ_3. So, ξ is the only minimally committed ε-stratification. Any interpretation that satisfies the ε-stratification ξ produces a set $\mathcal{E}$ such that: $\varepsilon_1 >_\infty \varepsilon_2$, $\varepsilon_1 >_\infty \varepsilon_3$ and there is no constraint between ε_2 and ε_3. Any such $\mathcal{E}$ is also compatible with the $\mathbf{C}_\Delta$ constraints. Thus, the set $\mathrm{Bel}_{lcd}(\Delta)$ consists of all the εbf's in $\mathrm{Bel}_\oplus(\Delta)$ built from such $\mathcal{E}$'s . ■

6.3. *The* ***LCD*** *consequence relation*

We define our new consequence relation using the elements in $\mathrm{Bel}_{lcd}(\Delta)$. It is easy to see that the $\mathrm{Bel}_{lcd}(\Delta)$ set "behaves well" for this goal. First, $\mathrm{Bel}_{lcd}(\Delta)$ is a subset of EBF(Δ), that is, any element of $\mathrm{Bel}_{lcd}(\Delta)$ is an εbf-model of Δ: this is true by construction, since any element of $\mathrm{Bel}_{lcd}(\Delta)$ must satisfy the auto-deduction constraints. Second, $\mathrm{Bel}_{lcd}(\Delta)$ is a proper subset of EBF(Δ): for example, consider any εbf in $\mathrm{Bel}_\oplus(\Delta)$ obtained by the ξ' partition in Example 3; this εbf satisfies the $\mathbf{C}_\Delta$ constraints, hence it belongs to EBF(Δ), but it is not in $\mathrm{Bel}_{lcd}(\Delta)$ because ξ' is not minimally committed. Third, if the default base Δ is consistent, then $\mathrm{Bel}_{lcd}(\Delta)$ is non-empty. To see this, take the $\mathcal{E}$ defined for the penalty order in Section 5.2 (or, for that matter, any of the approaches in Section 5): if Δ is consistent, this $\mathcal{E}$ satisfies the $\mathbf{C}_\Delta$ constraints (because $\mid\!\sim_{pen}$ is auto-deductive), and then the εbf built by (7) from this $\mathcal{E}$ belongs to $\mathrm{Bel}_{lcd}(\Delta)$.

We put $\mathrm{Bel}_{lcd}(\Delta)$ inside (10, section 5.1) to obtain our new definition of consequence relation. We name **LCD** this consequence relation, and denote it by $\mid\!\sim_{lcd}$.

Figure 2: Order constraints induced on the terms t_ω by ξ" in example 3.

We summarize the above arguments in the following definition of $Bel_{lcd}(\Delta)$.

Definition 13. Let Δ be a default base with $\mathcal{Z}$ an associated set of proto-infinitesimals and $\mathbf{C}_\Delta$ the set of constraints induced by Δ. Let $\mathcal{P}$ be the set of ε-stratifications of $\mathcal{Z}$ compatible with $\mathbf{C}_\Delta$. Let $\mathcal{P}^*$ be the set of minimally committed elements of $\mathcal{P}$. Let $\mathbb{I}$ be the set of interpretations $\mathcal{I}$ of $\mathcal{Z}$ that satisfy $\mathbf{C}_\Delta$ and an ε-stratification of $\mathcal{Z}$ in $\mathcal{P}^*$. Then we define $Bel_{lcd}(\Delta)$ to be the set of εbf $bel_{\mathcal{E}}$ such that there exist an $\mathcal{I} \in \mathbb{I}$ and $bel_{\mathcal{E}}$ is built from $\mathcal{E} = \mathcal{I}(\mathcal{Z})$ by (7, section 5.1).

Symbolically, this definition and those underlying it can be written as:

given a default base $\Delta = \{d : \alpha_d \rightarrow \beta_d\}$ and $\mathcal{Z}$ its associated set of proto-infinitesimals, let

$\mathcal{P} = \{\xi$ ε-stratification of $\mathcal{Z}, \exists \mathcal{I} : \mathcal{I} \models \xi, \mathcal{I} \models \mathbf{C}_\Delta\}$,

$\mathcal{P}^* = \{\xi : \xi \in \mathcal{P}, \nexists \xi' : \xi' \in \mathcal{P}, \xi'$ strictly less committed than $\xi\}$,

$\mathbb{I} = \{\mathcal{I} : \exists \xi \in \mathcal{P}^*, \mathcal{I} \models \xi, \mathcal{I} \models \mathbf{C}_\Delta\}$

$Bel_{lcd}(\Delta) = \{bel_{\mathcal{E}} : \mathcal{E} = \mathcal{I}(\mathcal{Z}), \mathcal{I} \in \mathbb{I}, bel_{\mathcal{E}}$ induced by $m_\oplus = \oplus\{m_d \mid d \in \Delta\}, (m_d(\Omega) = \varepsilon_d, m_d(\phi_d) = 1 - \varepsilon_d, m_d(X) = 0$ otherwise$)\}$.

This lengthy definition just means that $Bel_{lcd}(\Delta)$ is the set of all εbf that can be built from a set of infinitesimals that are compatible with the constraints imposed by the auto-deduction principle, and that are "least committed".

Example 4. Let again $\Delta = \{b \rightarrow f, p \rightarrow \neg f, p \rightarrow b\}$ be a default base, and $\mathcal{Z} = \{\zeta_1, \zeta_2, \zeta_3\}$ a corresponding set of proto-infinitesimals. The requirement of auto-deduction gives us the following three constraints on $\mathcal{I}$:

$\mathbf{C_1}$: $\max\{t_\omega \mid \omega \models b \wedge f\} >_\infty \max\{t_\omega \mid \omega \models b \wedge \neg f\}$ i.e., $1 >_\infty \varepsilon_1$

$\mathbf{C_2}$: $\max\{t_\omega \mid \omega \models p \wedge \neg f\} >_\infty \max\{t_\omega \mid \omega \models p \wedge f\}$ i.e., $\max\{\varepsilon_1, \varepsilon_3\} >_\infty \varepsilon_2$

$\mathbf{C_3}$: $\max\{t_\omega \mid \omega \models p \wedge b\} >_\infty \max\{t_\omega \mid \omega \models p \wedge \neg b\}$ i.e., $\max\{\varepsilon_1, \varepsilon_2\} >_\infty \varepsilon_3$

where $\varepsilon_i = \mathcal{I}(\zeta_i)$. The constraint $1 >_\infty \varepsilon_1$ is always satisfied and using Lemma **A13** the two last constraints are equivalent to:

$\mathbf{C_2}$: $\varepsilon_1 >_\infty \varepsilon_2$ or $\varepsilon_3 >_\infty \varepsilon_2$

$\mathbf{C_3}$: $\varepsilon_1 >_\infty \varepsilon_3$ or $\varepsilon_2 >_\infty \varepsilon_3$

minimize the number of classes, and to assign as many proto-infinitesimals as possible to classes with low index. Since by Definition 10 the interpretations of the proto-infinitesimals in lower classes are larger, using minimally committed ε-stratifications is a way to capture our informal requirement that the infinitesimals be "as large as possible".

We are now ready to define our new subclass of $Bel_{\oplus}(\Delta)$. We consider all the interpretations $\mathtt{I}$ of the set $\mathtt{Z} = \{\zeta_1, ..., \zeta_n\}$ of proto-infinitesimals and that $\mathtt{I}$ satisfies the auto-deduction constraints $\mathbf{C}_\Delta$. Among these, we consider those that satisfy an ε-stratification which is minimally committed among those that are compatible with $\mathbf{C}_\Delta$. We are interested in the εbf's built by (7, section 5.1) from any such set $\mathcal{E}$; we denote by $Bel_{lcd}(\Delta)$ the family of all these εbf's. The name "lcd" comes from the two main mechanisms used to build this family: "lc" for the least-commitment principle (applied twice), and "d" for Dempster's rule of combination.

Example 3. (Continued). Let $\mathcal{E}'$ and $\mathcal{E}''$ be the two following ε-stratifications: $\xi' = \{\xi'_0=\{\zeta_2,\zeta_3\}, \xi'_1=\{\zeta_4\}, \xi'_2=\{\zeta_5\}, \xi'_3=\{\zeta_1\}\}$; and $\xi'' = \{\xi''_0=\{\zeta_2,\zeta_3,\zeta_5\}, \xi''_1=\{\zeta_4,\zeta_1\}\}$, which are compatible with $\mathbf{C}_\Delta$. Of course, these are not the only ε-stratifications that are compatible with the $\mathbf{C}_\Delta$ constraints in this example, (e.g., consider the ε-stratification $\xi''' = \{\xi'''_0=\{\zeta_2\},\xi'''_1=\{\zeta_3\}, \xi'''_2=\{\zeta_4,\zeta_5\}, \xi'''_3=\{\zeta_1\}$). We can easily check that ξ'' is the only minimally committed one. Indeed, suppose that this was not the case; then there would exist an ε-stratification $\xi= \{\xi_0,\dots, \xi_m\}$ where at least either $\zeta_1$or ζ_4 is in ξ_0, but this is impossible since any interpretation that satisfies this ε-stratification either falsifies $\mathbf{C}_1$ or $\mathbf{C}_4$. The ε-stratification ξ'' induces a set of ordering constraints among the terms t_ω. Figure 2 graphically represents these constraints (dashed lines), together with those that can be deduced from the auto-deduction principle and from the properties of infinitesimals (solid lines). In the picture, an arrow from t to t' means that $t >_\infty t'$ holds (transitivity arrows are not drawn). ■

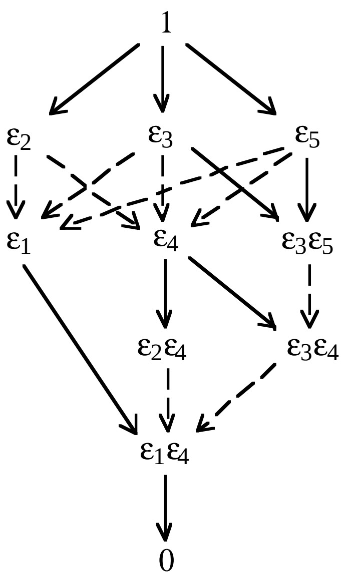

below (Example 9), not making these elements equivalent allows us to avoid some potential problems of syntax dependence.

In general, constraints between infinitesimals should be regarded as constraints over the possible interpretations $\mathtt{I}$. This is true in particular for the constraints $\mathbf{C_d}$ that result from the auto-deduction principle (section 6.1). The writing

$$\mathbf{C_1}: \quad \max\{\varepsilon_2, \varepsilon_5, \varepsilon_4\} >_\infty \varepsilon_1$$

should be actually taken to mean that the interpretation $\mathtt{I}$ should satisfy

$$\mathbf{C_1}: \quad \max\{\mathtt{I}(\zeta_2), \mathtt{I}(\zeta_5), \mathtt{I}(\zeta_4)\} >_\infty \mathtt{I}(\zeta_1).$$

For sake of simplicity, however, we keep the simplified writing as used in section 6.1.

Definition 11: An *interpretation* $\mathtt{I}$ *satisfies a set of constraints* $\mathbf{C}_\Delta$, written $\mathtt{I}\models\mathbf{C}_\Delta$, iff all the constrains $\mathbf{C_d}\in\mathbf{C}_\Delta$ are true statements. A*n ε-stratification* ξ *of* $\mathtt{Z}$ *satisfies* $\mathbf{C}_\Delta$ if each interpretation $\mathtt{I}$ satisfying ξ satisfies $\mathbf{C}_\Delta$. A*n ε-stratification* ξ *of* $\mathtt{Z}$ *is compatible with* $\mathbf{C}_\Delta$ if there exists an interpretation satisfying ξ which satisfies $\mathbf{C}_\Delta$.

In general, interpretations that satisfy ε-stratifications which are just compatible with $\mathbf{C}_\Delta$, but which do not satisfy $\mathbf{C}_\Delta$, do not necessarily make true *all* the constraints of $\mathbf{C}_\Delta$. As an example, consider $\Delta = \{b\rightarrow a, s\rightarrow a, b\wedge s\rightarrow\neg a\}$ ("if one takes a bath, he stays alive", "if one shaves (with an electric razor), he stays alive", "if one takes a bath while shaving, he does not stay alive"), and let $\mathtt{Z} = \{\zeta_1, \zeta_2, \zeta_3\}$ be the corresponding set of proto-infinitesimals. An interpretation $\mathtt{I}$ satisfies the auto-deduction constraints $\mathbf{C}_\Delta$ whenever: $1 >_\infty \mathtt{I}(\zeta_1)$, $1 >_\infty \mathtt{I}(\zeta_2)$, $\mathtt{I}(\zeta_1)\mathtt{I}(\zeta_2) >_\infty \mathtt{I}(\zeta_3)$. We can check that there are ε-stratifications of $\mathtt{Z}$ which are compatible with $\mathbf{C}_\Delta$ but there is no ε-stratification which satisfies $\mathbf{C}_\Delta$. This means that the constraints that result from the ε-stratification of $\mathtt{Z}$ alone cannot guarantee that all $\mathbf{C}_\Delta$ constraints are true statements. Therefore, we also need to consider the $>_\infty$ constraints between terms (i.e., products of elements) of $\mathcal{E}$.

Let us now introduce formally the least committed ε-stratification:

Definition 12. Let $\mathtt{Z} = \{\zeta_1,..., \zeta_n\}$ be a set of proto-infinitesimals, and let $\xi = \{\xi_0,\ldots, \xi_m\}$ and $\xi' = \{\xi'_0,\ldots,\xi'_{m'}\}$ be two ε-stratifications of $\mathtt{Z}$. We say that ξ is *less committed* than ξ' iff, for all $\zeta_i\in\mathtt{Z}$, $\zeta_i\in\xi_j$ and $\zeta_i\in\xi'_k$ imply $j\leq k$. We say that ξ is *strictly less committed* than ξ' if at least one inequality is strict. If $\mathcal{F}$ is a family of ε-stratifications of $\mathtt{Z}$, we call an element ξ of $\mathcal{F}$ *minimally committed* in $\mathcal{F}$ if there is no $\xi' \neq \xi$ in $\mathcal{F}$ such that ξ' is strictly less committed than ξ.

Intuitively, an ε-stratification is less committed than another one if it places the proto-infinitesimals in classes ξ_j with a lower index j. The least commitment principle comes down to “desirable” sets of infinitesimals. We have already applied this principle in deciding the form of the individual m_d’s in section 5.1. We now consider a second facet of the principle, and require that the $>_\infty$ ordering between terms of $\mathcal{E}$ be 'minimally committed'.

Defining what counts as a minimally committed ordering requires some care. We shall consider orderings of the infinitesimals that will be included in a set $\mathcal{E}$ such that, on the induced rule ordering: (i) rules in higher classes are more exceptional than those in lower class; and (ii) two rules in the same class are not comparable with respect to 'normality'. Note that this ordering corresponds to a stratification of Δ in the sense that it splits Δ into several strata. To apply the principle of least commitment, then, we shall look for orderings (stratifications) in which each rule is as uncommitted as possible. In terms of the infinitesimals in $\mathcal{E}$, this means that ε_d should belong to the highest possible class, provided that the auto-deduction constraints are not violated. In the example above, we would prefer $\mathcal{E}$” to $\mathcal{E}$’, as ε_5 and ε_1 are in “higher” classes in $\mathcal{E}$” than in $\mathcal{E}$’.

Formally, the task consists in building an appropriate $\mathcal{E}$, hence in selecting infinitesimals that satisfy some constraints based on some $>_\infty$ ordering.

Definition 9: Suppose a default base $\Delta = \{d_i : i = 1,...n\}$. Let $\mathtt{Z}$ be a set of variables ζ_i, i=1...n, one per default in Δ. We call these variables ζ_i's *proto-infinitesimals* (as infinitesimals will be assigned to them). An *interpretation* $\mathtt{I}$ is a mapping from proto-infinitesimals to infinitesimals. Let $\mathtt{I}(\zeta_i)$ denote the infinitesimal assigned by $\mathtt{I}$ to the proto-infinitesimals ζ_i: so $\mathtt{I}(\zeta_i) \in \mathbb{E}^0$. Let $\mathtt{I}(\mathtt{Z})$ be the set $\mathcal{E} = \{\varepsilon_1,\ldots, \varepsilon_n\}$so that $\varepsilon_i = \mathtt{I}(\zeta_i)$ for all ζ_i in $\mathtt{Z}$.

Definition 10. Let $\mathtt{Z} = \{\zeta_1, \ldots, \zeta_n\}$ be a set of proto-infinitesimals. An *ε-stratification* of $\mathtt{Z}$ is defined as a partition $\xi = \{\xi_0,\ldots, \xi_m\}$ of $\mathtt{Z}$ (partition is used here in the sense that $\xi_0 \cup \ldots \cup \xi_m = \mathtt{Z}$, $\forall i \neq j$, $\xi_i \cap \xi_j = \emptyset$ and $\forall i$, $\xi_i \neq \emptyset$). *An interpretation* $\mathtt{I}$ *satisfies an ε- stratification* $\xi = \{\xi_0,\ldots, \xi_m\}$ of $\mathtt{Z}$, written $\mathtt{I} \models \xi$, if for any ζ_i, ζ_j in $\mathtt{Z}$ with $\zeta_i \in \xi_i$ and $\zeta_j \in \xi_j$,
i < j implies $\mathtt{I}(\zeta_i) >_\infty \mathtt{I}(\zeta_j)$.

Intuitively, an ε-stratification induces (through the interpretation that satisfies it) a ranking of the elements of $\mathcal{E}$ by orders of magnitude: all infinitesimals assigned to a proto-infinitesimals in a class ξ_i are infinitesimally larger than those assigned to a proto-infinitesimals in any lower class ξ_j, (i.e., i < j). If $\mathtt{Z}$ is the set of proto-infinitesimals associated to a default base Δ, any ε-stratification of $\mathtt{Z}$ induces a stratification on Δ by putting all the rules whose corresponding proto-infinitesimals are in a given class ξ_i into the same layer.

Of course, several interpretations can satisfy the same ε-stratification. Furthermore, nothing is said about the ordering between infinitesimals assigned to the proto-infinitesimals that belong to the same class: their ordering depends on the particular interpretation. As we shall see

C$_3$: $1 >_\infty \varepsilon_3$

C$_4$: $\max\{\varepsilon_1, \varepsilon_2, \varepsilon_5\} >_\infty \varepsilon_4$

C$_5$: $1 >_\infty \varepsilon_5$.

Note that **C₂**, **C₃** and **C₅** are trivially satisfied by virtue of Lemma A7(b). So, any set $\mathcal{E} = \{\varepsilon_1, \varepsilon_2, \varepsilon_3, \varepsilon_4, \varepsilon_5\}$ of infinitesimals such that **C₁** and **C₄** are satisfied is a solution of $\mathbf{C}_\Delta$; for any such $\mathcal{E}$, the $bel_{\mathcal{E}}$ built by (7, section 5.1) is an εbf-model of Δ. ■

Note that in general the $\mathbf{C}_\Delta$ system of constraints does not determine a unique $>_\infty$ ordering that must be satisfied by the terms of $\mathcal{E}$. In the example above, any set $\mathcal{E}$' such that

$$\varepsilon_2 \approx_\infty \varepsilon_3 >_\infty \varepsilon_4 >_\infty \varepsilon_5 >_\infty \varepsilon_1$$

satisfies the $\mathbf{C}_\Delta$ constraints; but the same is true for any set $\mathcal{E}$” such that

$$\varepsilon_2 \approx_\infty \varepsilon_3 \approx_\infty \varepsilon_5 >_\infty \varepsilon_4 \approx_\infty \varepsilon_1.$$

This observation suggests that the auto-deduction principle may not constrain the $Bel_\oplus(\Delta)$ family to a sufficient extent, and the consequence relation based on this principle alone may still be too cautious. The next example shows that this is indeed the case.

Example 3. (Continued) Given the Δ considered above, we ask the question of whether or not penguins that have feathers and are birds can fly, that is, whether we can conclude “f” or “¬f” from “p∧fa∧b”. Clearly, we would like to be able to conclude “¬f” but not “f”. Let $bel_{\mathcal{E}}$ be any εbf built by (7, section 5.1). By Lemma 9, we have:

$pl_{\mathcal{E}}(p\wedge fa\wedge b\wedge f) \approx_\infty \varepsilon_4$

$pl_{\mathcal{E}}(p\wedge fa\wedge b\wedge\neg f) \approx_\infty \varepsilon_5$.

Let $\mid\!\sim_{ad}$ be the consequence relation based on all the εbf in $Bel_\oplus(\Delta)$ whose $\mathcal{E}$ satisfies the $\mathbf{C}_\Delta$ constraints. Take bel’ and bel” in $Bel_\oplus(\Delta)$ respectively built from the two ordered sets $\mathcal{E}$’ and $\mathcal{E}$” given above. Both sets satisfy the $\mathbf{C}_\Delta$ constraints. In $\mathcal{E}$’, we have $\varepsilon_4 >_\infty \varepsilon_5$, and then bel’ gives us “f” but not “¬f”. In $\mathcal{E}$”, we have $\varepsilon_5 >_\infty \varepsilon_4$, and then bel” gives us “¬f” but not “f”. Therefore, we have neither $p\wedge fa\wedge b \mid\!\sim_{ad} f$ nor $p\wedge fa\wedge b \mid\!\sim_{ad} \neg f$. ■

6.2. The least committed principle.

In the last example, the solution obtained from $\mathcal{E}$’ is clearly undesirable. Our next step is to strengthen our constraints by using the least commitment principle in order to select only

ω	p	fa	b	f	d_1	d_2	d_3	d_4	d_5	C_1	C_2	C_3	C_4	C_5
1	F	F	F	F										
2	F	F	F	T										
3	F	F	T	F			√		√			R		R
4	F	F	T	T			√					R		L
5	F	T	F	F		√					R			
6	F	T	F	T		√					R			
7	F	T	T	F					√		L	L		R
8	F	T	T	T							L	L		L
9	T	F	F	F	√					R			L	
10	T	F	F	T	√			√		R			R	
11	T	F	T	F			√		√	L		R	L	R
12	T	F	T	T			√	√		L		R	R	L
13	T	T	F	F		√				L	R		L	
14	T	T	F	T		√		√		L	R		R	
15	T	T	T	F					√	L	L	L	L	R
16	T	T	T	T				√		L	L	L	R	L

Table 1: Construction of the $\mathbf{C}_\Delta$ constraints of Example 3.

Each row in the table represents a possible truth assignment to our four propositional variables or world. A √ mark in the d_i column indicates that the world does *not* satisfy the default d_i, that is, ϕ_{d_i} is false in that world. For each world ω, the term t_ω is given by the product of the ε_d's for the defaults d marked by a √ in the ω row. E.g., $t_3 = \varepsilon_3\varepsilon_5$. The $\mathbf{C_i}$ column is used to write the auto-deduction constraint for the default $d_i = \alpha_i \rightarrow \beta_i$: we put an L mark for each world that satisfies $\alpha_i \wedge \beta_i$, and an R mark for each world that satisfies $\alpha_i \wedge \neg\beta_i$. We then write the $\mathbf{C_i}$ constraint as follows: for the left hand side, we take the maximum of the terms t_ω over the words ω marked by L in the $\mathbf{C_i}$ column, where t_ω is obtained as above; we proceed in a similar way for the right hand side using the R worlds. In our example, we obtain the following five constraints:

$\mathbf{C_1}$: $\max\{\varepsilon_3\varepsilon_5, \varepsilon_3\varepsilon_4, \varepsilon_2, \varepsilon_2\varepsilon_4, \varepsilon_5, \varepsilon_4\} >_\infty \max\{\varepsilon_1, \varepsilon_1\varepsilon_4\}$

$\mathbf{C}_2$: $\max\{\varepsilon_5, 1, \varepsilon_4\} >_\infty \max\{\varepsilon_2, \varepsilon_2\varepsilon_4\}$

$\mathbf{C}_3$: $\max\{\varepsilon_5, 1, \varepsilon_4\} >_\infty \max\{\varepsilon_3\varepsilon_5, \varepsilon_3, \varepsilon_3\varepsilon_4\}$

$\mathbf{C}_4$: $\max\{\varepsilon_1, \varepsilon_3\varepsilon_5, \varepsilon_2, \varepsilon_5\} >_\infty \max\{\varepsilon_1\varepsilon_4, \varepsilon_3\varepsilon_4, \varepsilon_2\varepsilon_4, \varepsilon_4\}$

$\mathbf{C}_5$: $\max\{\varepsilon_3, 1, \varepsilon_3\varepsilon_4, \varepsilon_4\} >_\infty \max\{\varepsilon_3\varepsilon_5, \varepsilon_5\}$.

The properties of infinitesimals given in Appendix A allow us to simplify these constraints. For example, from Lemma A8(c), we can deduce: $\varepsilon_2 >_\infty \varepsilon_2\varepsilon_4$; $\varepsilon_4 >_\infty \varepsilon_3\varepsilon_4$; and $\varepsilon_5 >_\infty \varepsilon_3\varepsilon_5$. Hence, by Lemma A12, $\max\{\varepsilon_2, \varepsilon_5, \varepsilon_4\} \approx_\infty \max\{\varepsilon_3\varepsilon_5, \varepsilon_3\varepsilon_4, \varepsilon_2, \varepsilon_2\varepsilon_4, \varepsilon_5, \varepsilon_4\}$, and we can replace the former for the latter in $\mathbf{C_1}$ by virtue of Lemma A9(b). By similar arguments, we obtain the following simplified set of constraints:

$\mathbf{C_1}$: $\max\{\varepsilon_2, \varepsilon_5, \varepsilon_4\} >_\infty \varepsilon_1$

$\mathbf{C}_2$: $1 >_\infty \varepsilon_2$

The second facet is new, and it concerns the choice of the $\mathcal{E}$ set: we want this set to be "as free as possible" with respect to the $>_\infty$ ordering between its elements and terms; as we shall see shortly, this corresponds to requiring that the elements of $\mathcal{E}$ be "as large as possible" while retaining the auto-deduction property.

In the rest of this section, we show how we can translate the above principles into a precise set of $>_\infty$ constraints on $\mathcal{E}$. These constraints identify a specific subfamily of $Bel_\oplus(\Delta)$ which we will use to define our new consequence relation.

6.1 *The auto-deduction principle*

By Definition 4 and Lemma 4, the auto-deduction principle means that, for each default d = $\alpha \to \beta$ in Δ, we must have

$$\max_{\omega \models \alpha\wedge\beta} pl_\oplus(\omega) \quad >_\infty \quad \max_{\omega \models \alpha\wedge\neg\beta} pl_\oplus(\omega).$$

This constraint is equivalent to

$$\mathbf{C_d:} \quad \max_{\omega \models \alpha\wedge\beta} t_\omega \quad >_\infty \max_{\omega \models \alpha\wedge\neg\beta} t_\omega \ ,$$

where the t_ω's are the terms defined by (8) in section 5.1 above, and we have used Lemma A9(b) to substitute equivalent infinitesimals inside inequalities. Note that we can equivalently write:

$\mathbf{C_d}$: for all ω s.t. $\omega \models \alpha\wedge\neg\beta$, there is a ω' s.t. $\omega' \models \alpha\wedge\beta$ and $t_{\omega'} >_\infty t_\omega$.

Given Δ, the auto-deduction principle translates to a set of constraints $\mathbf{C}_\Delta = \{\mathbf{C_d} \mid d \in \Delta\}$ between the terms of $\mathcal{E}$; any $bel_\mathcal{E}$ built from a set $\mathcal{E}$ that satisfies these constraints also satisfies the auto-deduction principle; said differently, this $bel_\mathcal{E}$ is an εbf-model of Δ.

Example 3. Let Δ = {p→fa∨b, fa→b, b→fa, p→¬f, b→f}, where p, fa, b, f stand for "penguin", "feathered animal", "bird" and "fly", and let $\mathcal{E} = \{\varepsilon_1, \varepsilon_2, \varepsilon_3, \varepsilon_4, \varepsilon_5\}$ be the corresponding set of infinitesimals. Table 1 gives all what is needed to compute the $\mathbf{C}_\Delta$ constraints.

Note that the complexity results of several non-monotonic inference discussed in this Section can be for instance found in Eiter and Luckasiewicz (2000).

6 Using Both the LC Principle and Dempster's Rule.

We have seen how to define several non-monotonic consequence relations based on the $Bel_{\oplus}(\Delta)$ family of εbf's by imposing some a-priori constraints over (products of) the weights ε_d associated to the different defaults in Δ. Until now, we have considered ad-hoc constraints inspired by existing non-monotonic systems. In this section, we take a different route, and generate the constraints starting from two general principles: the combined εbf should satisfy all the defaults in Δ; and it should be least committed, in a sense that we shall soon define. As we shall show, the resulting consequence relation, that we name **LCD**, is distinct from all the current non-monotonic systems that are aware of; moreover, **LCD** adequately addresses all the desiderata listed in the Introduction.

Recall that the $Bel_{\oplus}(\Delta)$ family in the previous section was built from a given set $\Delta = \{d_1,\ldots, d_n\}$ of defaults and an *arbitrary* set $\mathcal{E} = \{\varepsilon_1,\ldots, \varepsilon_n\}$ of infinitesimals associated to these defaults. Also recall that the properties of the induced consequence relation $\mid\!\sim_{\oplus}$ critically depend on certain properties of $\mathcal{E}$, namely, on the $>_{\infty}$ ordering between its elements and terms. We now propose to restrict the $Bel_{\oplus}(\Delta)$ family by restricting the choice of the $\mathcal{E}$ set in such a way that the two following principles are satisfied:

- *Auto-deduction principle.* For any $bel_{\mathcal{E}}$ in $Bel_{\oplus}(\Delta)$, we want that $bel_{\mathcal{E}} \models \alpha\to\beta$ for each rule $\alpha\to\beta$ in Δ; said differently, we want $Bel_{\oplus}(\Delta) \subseteq EBF(\Delta)$.
- *Least commitment principle.* We want each $bel_{\mathcal{E}}$ to be minimally committed among the $bel_{\mathcal{E}}$ resulting from the auto-deduction principle, that is, we want it to convey as little information as possible while still satisfying all the defaults in Δ.

The auto-deduction principle amounts to requiring that our $bel_{\mathcal{E}}$ be εbf-models of Δ; this principle was already satisfied by the non-monotonic consequence relations considered in the previous section. The least-commitment principle deserves more comments. Given the way in which $Bel_{\oplus}(\Delta)$ is built, this principle includes two distinct facets. The first one concerns the form of the individual m_d's used in the combination, and has already been discussed in Section 5.1: it translates to the requirement that each m_d is a simple support function with focus ϕ_d,

$m_d(\Omega) = \varepsilon_d$; $m_d(\phi_d) = 1 - \varepsilon_d$; and $m_d(X) = 0$ otherwise; with $\varepsilon_d \in \mathcal{E}$.

but satisfied by ω', there exists a default d' satisfied by ω but falsified by ω' such that $d' >_\Delta d$.[10] The triple $(\Omega, >_\Omega, \Delta, >_\Delta)$ is called a *prioritized admissible structure.* Finally, Geffner defines a conditional entailment relation $\mid\!\sim_G$ in the usual way by: $\alpha \mid\!\sim_G \beta$ iff, for each admissible prioritized structure $(\Omega, >_\Omega, \Delta, >_\Delta)$, β holds in all the $>_\Omega$-preferential models of α.

We can capture Geffner's conditional entailment in our framework. For each admissible preference relation $>_\Delta$ on Δ, we build the set $\mathcal{E}$ by associating to each $d \in \Delta$ an infinitesimal $\varepsilon_d \in \mathcal{E}$ such that:

1) for all $d,d' \in \Delta$, if $d >_\Delta d'$ then $\varepsilon_{d'} >_\infty \varepsilon_d$; and

2) for all $d \in \Delta$, $\prod_{d':\varepsilon_{d'} >_\infty \varepsilon_d} \varepsilon_{d'} >_\infty \varepsilon_d$.

The second default requires that, if the infinitesimal corresponding to a default is (infinitesimally) smaller than the infinitesimals corresponding to other defaults, it also be (infinitesimally) smaller than their product.

Let $>_\Delta$ be an admissible preference relation on Δ. We denote by $Bel_\oplus(\Delta, >_\Delta)$ the family of εbf in $Bel_\oplus(\Delta)$, associated to $>_\Delta$, whose parameter $\mathcal{E}$ satisfies the above constraints. Each εbf in $Bel_\oplus(\Delta, >_\Delta)$ encodes the admissible prioritized structure associated to $>_\Delta$:

Lemma 13. Let $>_\Delta$ be an admissible preference relation on Δ, and $(\Omega, >_\Omega, \Delta, >_\Delta)$ be an admissible prioritized structure. Then, for any ω and ω' in Ω, and $\omega >_\Omega \omega'$ if and only if ω is $bel_{\mathcal{E}}$-preferred to ω' for any $bel_{\mathcal{E}}$ in $Bel_\oplus(\Delta, >_\Delta)$.

We denote by:

$$Bel_{\oplus 4}(\Delta) = \bigcup \{Bel_\oplus(\Delta, >_\Delta) \text{ s.t. } >_\Delta \text{ is admissible with } \Delta\}$$

the family of εbf in $Bel_\oplus(\Delta)$ obtained by taking into account all the orders $>_\Delta$ which are admissible with Δ. Finally, we let $\mid\!\sim_{\oplus 4}$ be the consequence relation obtained by using $Bel_{\oplus 4}(\Delta)$ in (10, section 5.1).

Theorem 6. For a given Δ, $\alpha \mid\!\sim_G \beta$ if, and only if, $\alpha \mid\!\sim_{\oplus 4} \beta$.

<u>Proof</u>. Immediate consequence of Lemma 13. ■

---

[10] Note that the total ordering on Δ induced by the stratification given by LC algorithm (or System Z) is admissible with Δ. Moreover, when $>_\Delta$ is again the stratification given by LC algorithm, then the definition of $>_\Omega$ given by Geffner is exactly the same as the one induced by Brewka's preferred sub-theories.

enumeration) in the i-th layer Δ_i. We build $\mathcal{E}$ by associating to each d_{ij} an infinitesimal ε_{ij} such that, for any i,

$$\prod_{j=1}^{i-1} \prod_{h=1}^{|\Delta_j|} \varepsilon_{jh} >_\infty \varepsilon_{is} \text{ for all } s = 1, ..., |\Delta_i|,$$

where $|\Delta_i|$ is number of defaults in Δ_i. The main difference between these constraints and those used to recover the lexicographical system is that here ε_{ij} and ε_{ih} are not constrained for all $j \neq h$, while in the lexicographical systems they are considered as equivalent.

We denote by $Bel_{\oplus 3}(\Delta)$ the family of εbf's obtained by (7) in section 5.1 for different choices of the ε_{ij}'s infinitesimals, provided that they satisfy the above constraints. Clearly $Bel_{\oplus 2}(\Delta) \subseteq Bel_{\oplus 3}(\Delta)$.

The following lemma shows the ordering based on $Bel_{\oplus 3}(\Delta)$ induces the same ordering on Ω as the one used by Brewka.

Lemma 12. Let ω and ω' be elements of Ω. Then, ω is B-preferred to ω' if and only if ω is $bel_{\mathcal{E}}$-preferred to ω' in each $bel_{\mathcal{E}}$ of $Bel_{\oplus 3}(\Delta)$.

Finally, we let $\mid\!\sim_{\oplus 3}$ be the consequence relation obtained by using $Bel_{\oplus 3}(\Delta)$ in (10, section 5.1).

Theorem 5. Let $\mid\!\sim_B$ be the inference relation of Brewka's preferred sub-theories system. Then, for any given Δ, $\alpha \mid\!\sim_B \beta$ if, and only if, $\alpha \mid\!\sim_{\oplus 3} \beta$.

Proof. See the appendix. ■

5.5 Geffner's conditional entailment

The fourth approach that we consider here is Geffner's conditional entailment (1992). Geffner's approach differs from the previous ones in the following points:

• Several orders between defaults are considered rather than just one. Each of them is called admissible with the default base Δ.

• The order between default rules is in general a partial order (in the previous approaches, it was a complete pre-ordering).

An order between elements of a default base Δ, denoted by $>_\Delta$, is said to be *admissible* with Δ if, whenever a sub-base $D \subseteq \Delta$ does not tolerate a default $d \in \Delta$, then D contains another default d' such that $d >_\Delta d'$. (Tolerate here is taken in the sense of Pearl's system **Z**.) Each admissible order $>_\Delta$ on Δ induces an order over Ω in the following way: $\omega >_\Omega \omega'$ iff, for each default d falsified by ω

The second case that we consider is the lexicographic system defined by Benferhat et al. (1993), and Lehmann (1995), and set by Dubois et al. (1992) in the framework of possibilistic logic. A similar approach has also been considered in diagnosis by De Kleer (1990) and Lang (1994). The main idea is to again start from a stratification $\{\Delta_1,\ldots, \Delta_k\}$ of Δ, and regard all defaults in the i-th layer as being equally important, and being more important than any set of defaults in subsequent layers (see the proof of Lemma 11 in Appendix B for a complete description of the lexicographic system).

We can reproduce the behavior of the lexicographic system by our approach by imposing a corresponding order on the ε_d's. We again start from the LC (or **Z**) stratification $\{\Delta_1,\ldots, \Delta_k\}$ of Δ, and associate an infinitesimal δ_i to each layer Δ_i such that

$$\prod_{j=1}^{i-1} \delta_j^{|\Delta j|} >_\infty \delta_i \,, \quad \text{for } i = 2, ..., n+1$$

where $|\Delta_j|$ is the number of defaults in Δ_j.[9] We than build $\mathcal{E}$ by letting $\varepsilon_d = \delta_i$ for each i and each $d\in\Delta_i$. Clearly, any such ε_d is an infinitesimal. We denote by $Bel_{\oplus 2}(\Delta)$ the family of εbf's obtained from these ε_d's by applying (7, section 5.1). These εbf's differ in the choice of the δ_i's infinitesimals, provided that they satisfy the above constraints. Each εbf in $Bel_{\oplus 2}(\Delta)$ induces the same ordering on Ω as the one used in the lexicographic approaches.

Lemma 11. Let ω and ω' be elements of Ω, and let $bel_{\mathcal{E}}$ be any element of $Bel_{\oplus 2}(\Delta)$. Then, ω is lex-preferred to ω' if and only if ω is $bel_{\mathcal{E}}$-preferred to ω'.

Finally, we let $\mid\!\sim_{\oplus 2}$ be the consequence relation obtained by using $Bel_{\oplus 2}(\Delta)$ in (10, section 5.1).

Theorem 4. Let $\mid\!\sim_{lex}$ be the inference relation of the lexicographic system. Then, for any given Δ, $\alpha \mid\!\sim_{lex} \beta$ if, and only if, $\alpha \mid\!\sim_{\oplus 2} \beta$.

Proof. Immediate, as $\mid\!\sim_{\oplus 2}$ and $\mid\!\sim_{lex}$ are preferential relations based on the same order. ■

5.4 Brewka's preferred sub-theories

The third approach that we consider are the *preferred sub-theories* originally proposed by Brewka (1989), and later independently introduced in (Dubois et al., 1992) in the setting of possibilistic logic. This approach has also been used by Boutilier (1992) in system **Z** to define a non-monotonic inference relation, and by Baral et al. (1992) to combine belief bases. The core of the approach is the same as for the lexicographic one, the main difference being the definition of the ordering on Ω (see the proof of Lemma 12 in Appendix B for details).

To recover Brewka's ordering in our framework, we let $\{\Delta_1,\ldots,\Delta_k\}$ be the LC (or **Z**) stratification of a given base Δ. Let then d_{ij} be the j-th default (according to an arbitrary

---

[9] A similar condition has been proposed by Snow (1996) to define a class of probability distributions called "Atomic bound systems". A probability distribution is an Atomic bound system iff there exists a linear ordering $>_S$ on Ω, such that for each $\omega\in\Omega$, $P(\omega) > \sum_{\omega>_S\omega'} P(\omega')$. With a logarithmic transformation, this condition is very close to the one given above between infinitesimals.

However, all of these approaches also apply to default reasoning, since the latter can be seen as a particular case of inconsistency handling — the observation of an exceptional situation leads to an inconsistent knowledge base, which needs to be revised. Technically, each one of the approaches considered here leads to a preferential order over the interpretations, which in turn induces a non-monotonic consequence relation over the formulas. In each case, we show that the same order can be recovered using εbf's. Symmetrically, our reformulation show that εbf's can be used to deal with the general case of knowledge base revision in face of inconsistency.

In penalty logic, each formula of a knowledge base Σ is associated with a number, called penalty, which represents a kind of price to pay if this formula is not satisfied. The sum of the penalties of all the formulae that are not satisfied by an interpretation ω is called the *cost* of ω; this cost induces a preference over interpretations (see the proof of Lemma 10 in Appendix B for details).

We can easily reproduce the behavior of penalty logic in our framework. Let Δ be a default base, and let $\{\Delta_1,\ldots, \Delta_k\}$ be a stratification of Δ that respects specificity; for example, the stratification generated by the LC algorithm, or equivalently by Pearl's system **Z**. We build the set $\mathcal{E} = \{\varepsilon_1,..., \varepsilon_n\}$ of infinitesimals associated to Δ as follows: we take an arbitrary infinitesimal δ and let, for each $d \in \Delta_i$, $\varepsilon_d = \delta^i$. Intuitively, we need the exponential to recover the additive behavior of penalties from the multiplicative behavior of our ε_d's (see Lemma 9). Note that any such ε_d obviously satisfies the definition of infinitesimals in Sec. 2.4. We denote by $Bel_{\oplus 1}(\Delta)$ the family of εbf's obtained from these ε_d's by applying (7); these εbf's differ in the choice of the δ infinitesimal. Each εbf in this family induces the same ordering on Ω as the one used in penalty logic.

Lemma 10. Let ω and ω' be elements of Ω, and let $bel_{\mathcal{E}}$ be any element of $Bel_{\oplus 1}(\Delta)$. Then, ω is penalty-preferred to ω' if and only if ω is $bel_{\mathcal{E}}$-preferred to ω'.

Finally, we let $\mid\!\sim_{\oplus 1}$ be the non-monotonic consequence relation obtained by using $Bel_{\oplus 1}(\Delta)$ for $Bel_{\oplus}(\Delta)$ in (10) in section 5.1.

Theorem 3. Let $\mid\!\sim_{pen}$ be the inference relation of penalty logic. Then, for any given Δ, $\alpha \mid\!\sim_{pen} \beta$ if, and only if, $\alpha \mid\!\sim_{\oplus 1} \beta$.

Proof. By Lemma 5, $\mid\!\sim_{\oplus 1}$ can be defined in a preferential way. So, both $\mid\!\sim_{\oplus 1}$ and $\mid\!\sim_{pen}$ are preferential relations, and the thesis follows by noting that they are based on the same order. ∎

5.3 The lexicographic system

$m_3(\neg p \vee b) = 1 - \varepsilon_3$, $m_3(\Omega) = \varepsilon_3$.

According to (7) we have, for all X, $m_\oplus(X) = (m_1 \oplus m_2 \oplus m_3)(X)$. The value of $pl_\oplus$ for any proposition of interest can be computed by (9). For instance, we have: $pl_\oplus(b \wedge p \wedge f) \approx_\infty \varepsilon_2$, $pl_\oplus(b \wedge p \wedge \neg f) \approx_\infty \varepsilon_1$, $pl_\oplus(b \wedge \neg p \wedge f) \approx_\infty 1$, and $pl_\oplus(\neg b \wedge p \wedge f) \approx_\infty \varepsilon_1 \cdot \varepsilon_2$. From this, we see that $bel_\oplus \models b \wedge p \rightarrow \neg f$ iff $\varepsilon_1 >_\infty \varepsilon_2$; thus, we have $b \wedge p \mathrel{|\!\!\sim}_\oplus \neg f$ if and only if the $\varepsilon_1 >_\infty \varepsilon_2$ relation is guaranteed to hold for all the elements of $Bel_\oplus(\Delta)$. ■

The last example shows that the $\mathrel{|\!\!\sim}_\oplus$ consequence relation built in this way is not interesting if we let all the ε_d to be equal — for example, we could not deduce that birds who are penguins do not fly. More generally, it shows that the properties of $\mathrel{|\!\!\sim}_\oplus$ are determined by the $\geq_\infty$ relations that exist between the elements ε_d of $\mathcal{E}$, and the induced $>_\infty$ and $\approx_\infty$ relations — that is, on the relative stiffness of the defaults in Δ. In fact, the properties of $\mathrel{|\!\!\sim}_\oplus$ depend in general on the ordering between values of $pl_\oplus$, that is, by (9), between terms.

In the rest of this paper, we explore the properties of different consequence relations obtained from (10) by imposing different types of constraints on the $\mathcal{E}$ parameters; these constraints are typically in the form of a set of $\geq_\infty$-inequalities between terms. In the following subsections, we impose some a priori constraints inspired by four existing systems, and we obtain consequence relations, which are equivalent to these systems. In the next section, we will use constraints coming from general principles, and obtain a new consequence relation that is incomparable with all the existing systems we are aware of and that go beyond system **P**. The ability of εbf's to capture several existing approaches in a uniform framework is an important property; however, the reader mainly interested in the new system can directly jump to Section 6.

5.2 Penalty-based order

We now start our analysis of special cases of the $\mathrel{|\!\!\sim}_\oplus$ consequence relations. In each case, we define a class of sets $\mathcal{E}$ of infinitesimals using a construction inspired by an existing system, and consider the consequence relation obtained by (10) when we restrict attention to the εbf's $bel_\mathcal{E}$ in $Bel_\oplus(\Delta)$ which are built from these sets. In each case, we shall show that any such $bel_\mathcal{E}$ induces the same $bel_\mathcal{E}$-preference relation $\prec_\mathcal{E}$ on Ω (see Definition 7), and that this preference relation is the same as the one of the inspiring system. Hence, the εbf-consequence relation so defined is equivalent to the one of the target system. The first case that we consider is inspired by the so-called "penalty logic" proposed by Pinkas (1991), and applied by Dupin et al. (1994) to handle inconsistencies in knowledge bases. It should be noted that penalty logic, as well as the other systems considered below, has originally been proposed as an inconsistency-handling technique.

The εbf's built by (7) have in general a much more complex structure than the ones used in LC: the focal elements are not nested, and values may include products of several ε_d and $(1-\varepsilon_d)$ for different d. Luckily, as we only care for the order of magnitude of the infinitesimals, the plausibility value of each world can be computed in a simple way as follows.

Lemma 9. Let $bel_\oplus$ be an εbf built from Δ and $\mathcal{E}$ as above. Then, for any world ω in Ω,

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

and $pl_\oplus(\omega) \approx_\infty 1$ if ω satisfies all the defaults in Δ.

In what follows, we shall use t_ω to denote the product of ε_d's that gives the plausibility of world ω once Δ and $\mathcal{E}$ have been fixed; that is, for any $\omega \in \Omega$ we let

$$t_\omega = \prod_{\substack{d\in\Delta \\ \omega\not\models\phi_d}} \varepsilon_d , \tag{8}$$

with the stipulation that $t_\omega = 1$ if ω satisfies all the defaults in Δ. We call t_ω the *term* associated to ω; note that, for any ω, $t_\omega \in \mathbb{E}$. Lemma 9 together with Lemma 3 allow us to compute the (order of the) plausibility value of any formula α by:

$$pl_\oplus(\alpha) \approx_\infty \max_{\omega\models\alpha} t_\omega. \tag{9}$$

Given a default base Δ, we denote by $Bel_\oplus(\Delta)$ the family of all εbf's that can be built from Δ using Dempster's rule according to (7); the elements of $Bel_\oplus(\Delta)$ differ in the choice of the set $\mathcal{E}$ of infinitesimals associated to the defaults in Δ. We can use $Bel_\oplus(\Delta)$ to define a new notion of non-monotonic consequence, denoted by $\mid\!\sim_\oplus$, as follows

$$\alpha \mid\!\sim_\oplus \beta \;\text{ iff }\; \text{for any } bel_{\mathcal{E}} \text{ in } Bel_\oplus(\Delta),\ bel_{\mathcal{E}} \models \alpha\rightarrow\beta, \tag{10}$$

where Δ represents, as usual, the given background knowledge.

Example 2. Let again $\Delta = \{b\rightarrow f,\ p\rightarrow\neg f,\ p\rightarrow b\}$, and let $\mathcal{E} = \{\varepsilon_1, \varepsilon_2, \varepsilon_3\}$ be the infinitesimals associated to the defaults in Δ. The ε-mass assignments m_d corresponding to these defaults 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

Then, it makes sense to represent each default by one (epsilon-) belief function, and to combine these belief functions by Dempster's rule to obtain a representation of the aggregate effect of all the defaults in Δ.[8] We can then define entailment by looking at the conditionals that are satisfied by the combined belief function. The characteristics of the consequence relation so obtained depend on certain conditions over the relative weights of the defaults. As we show below, we can re-obtain several of the non-monotonic relations that have been proposed in the literature by imposing the right conditions. In section 6, we will use more general conditions and obtain a new consequence relation that has several advantages over the current proposals.

5.1 Entailment based on Dempster's rule

In order to build the combined εbf mentioned above, we first have to decide which belief function to use to represent each individual default rule. Sticking to the arguments used in the last section, we propose to use the ε-least committed εbf that satisfies that default. More precisely, for each default rule $d = \alpha \rightarrow \beta$, we want the corresponding ε-mass assignment m_d to be minimally committed in the sense explained in Section 4 (see Definition 8): that is, m_d should belong to the set Λ_d of non-dominant εbf that satisfy the default d. From Lemma 6, this means that each m_d must have the form

$$m_d(\Omega) = \varepsilon_d;\ \ m_d(\phi_d) = 1 - \varepsilon_d;\ \ \text{and } m_d(X) = 0 \text{ otherwise,}$$

where ε_d is an infinitesimal in $\mathbb{E}^0$. Given a default base $\Delta = \{d_1,...,d_n\}$, we consider a set $\mathcal{E} = \{\varepsilon_1,..., \varepsilon_n\}$ of infinitesimals, and associate each default d_i to ε_i; we write ε_d to denote the infinitesimal associated to the default $d \in \Delta$. Intuitively, ε_d accounts for the "violability" of rule d. From these Δ and $\mathcal{E}$ we build a combined ε-mass assignment $m_\oplus$ by Dempster's rule as follows:

$$m_\oplus = \oplus\{m_d \mid d \in \Delta\}, \tag{7}$$

where each m_d is as above. We denote by $bel_\oplus$ and $pl_\oplus$ the corresponding ε-belief and ε-plausibility functions. Note that, no matter what $\mathcal{E}$, each m_d is an ε-mass assignment. As far as Δ is consistent, $\cap_{d\in\Delta} [\phi_d] = [\phi_\Delta] \neq \emptyset$, so Lemma 2 condition is satisfied and $m_\oplus$ is thus an ε-mass assignment.

---

[8] The construction given below should extend to the case where each source S_i of information provides a base Δ_i of defaults instead of just one default. We could use the ε-least-commitment principle to build a representative εbf for each Δ_i as discussed in the previous section, and then combine these representatives by Dempster's rule. The exploration of this issue is left for future work.

step i; that is, for each $\alpha \rightarrow \beta$ of Δ, $\alpha \rightarrow \beta \in \Delta'_i$ if and only if $bel_i \models \alpha \rightarrow \beta$. The following lemma shows that the Δ_i sub-base generated by Pearl's system contains exactly the same defaults.

Lemma 8. Let $\{\Delta_1, \Delta_2, \ldots, \Delta_k\}$ be the stratification given by system **Z**, and let bel_i the εbf built by the LC algorithm at step i. Then, for any default $\alpha \rightarrow \beta$ in Δ,

a) $\alpha \rightarrow \beta$ is tolerated by Δ iff $pl_1(\alpha) = 1$.

b) $\alpha \rightarrow \beta$ is not tolerated by $\Delta_i \cup \ldots \Delta_k$ iff $pl_i(\alpha \wedge \beta) = pl_i(\alpha \wedge \neg\beta) = \varepsilon_i$.

c) $\alpha \rightarrow \beta \in \Delta_i$ implies $bel_i \models \alpha \rightarrow \beta$.

d) $bel_i \models \alpha \rightarrow \beta$ and $pl_{i-1}(\alpha \wedge \beta) = pl_{i-1}(\alpha \wedge \neg\beta) = \varepsilon_{i-1}$ (i.e., $bel_{i-1} \not\models \alpha \rightarrow \beta$) implies that $\alpha \rightarrow \beta \in \Delta$.

The fact that the ranking produced by the LC algorithm coincides with the one of system **Z** tells us that the LC consequence precisely coincides with Pearl's system, and those systems equivalent to system **Z** like the one based on possibilistic logic (Dubois et al., 1994; Benferhat et al., 1992).

Theorem 2. For a given Δ, $\alpha \mathrel{|\!\!\sim}_{lc} \beta$ if, and only if, $\alpha \mathrel{|\!\!\sim}_{Z} \beta$.

Proof. Immediate consequence of Lemma 8. ■

The stratification produced by Pearl's algorithm or by our LC algorithm induces an ordering that respects the specificity requirement. This is a useful property, and we will use this ordering for defining several non-monotonic systems in the next section. It should be emphasized that, although LC and **Z** produce the same stratification of Δ, our approach differs from Pearl's in the choice of the primitive notion. In system **Z,** one starts from an a priori definition of "tolerance", while in the LC approach we start from the notion of being "less committed". Which of the two notions provides a more natural starting point is a matter of opinion.

5 Using Dempster's rule of Combination

The second route that we propose to explore to obtain a stronger consequence relation is to strengthen (4, section 3.2) 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. (A similar construction has been suggested by Wilson, see section 8.) Suppose that we regard each default in Δ as being one item of evidence provided by one of several *distinct* sources of information, each associated with a *weight* that indicates its reliability, i.e., the relative "stiffness" of the default rule it provides.

εbf, but were not satisfied by bel_{i-1}. An ordered partition of Δ as the one above is often called a *stratification*, or *ranking*, of Δ.

The idea to stratify the default base Δ was first suggested by Pearl (1990). To overcome the problem of irrelevance in Adams' ε-consequence relation, Pearl proposed a default reasoning system, called **Z**, based on a ranking of default rules that respects the notion of specificity. Pearl defines three relations between a default rule $\alpha \rightarrow \beta$ and a possible world ω: the world *verifies* the rule if $\omega \vDash \alpha \wedge \beta$; it *satisfies* it if $\omega \vDash \neg\alpha \vee \beta$; and it *falsifies* it if $\omega \vDash \alpha \wedge \neg\beta$. 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 $\Delta = \{\alpha_i \rightarrow \beta_i,\ i = 1,..., m\}$ 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 (stratification) $\{\Delta_1, \Delta_2, \ldots, \Delta_k\}$ such that rules in Δ_i are tolerated by all rules in $\Delta_i \cup \ldots \cup \Delta_k$. Correspondingly, he defines a ranking Z over the rules of Δ by letting Z(d) = i for each rule d in Δ_i.

From this ranking, Pearl defines a non-monotonic consequence relation $\mid\!\sim_Z$ in the following way. First, he builds a ranking function κ on the worlds from the Z ranking by

$$\kappa(\omega) = \max\{Z(d_i) + 1 \mid \omega \vDash \alpha_i \wedge \neg\beta_i,\ d_i = \alpha_i \rightarrow \beta_i \in \Delta\},$$

and κ(ω)=1 if ω satisfies all the rules of Δ.
Intuitively, κ(ω) is lower for worlds that only violate rules that are in lower sub-bases — i.e., that are less specific. Second, he induces from κ a ranking function z on formulas by

$$z(\alpha) = \min\{\kappa(\omega) \mid \omega \vDash \alpha\}.$$

z(α) is low if all the models of α have low rank. So, z(α) can be read as a degree of "abnormality" of α with respect to the rules in Δ. Note that z is a disbelief function in the sense of Spohn (1988). Indeed, we have $z(\alpha\vee\beta)=\min(z(\alpha), z(\beta))$. Finally, the $\mid\!\sim_Z$ relation is defined by

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

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

The above stratification coincides with the one produced by the LC algorithm.[7] To see why, consider a default base Δ, and let $\{\Delta_1, \Delta_2, \ldots, \Delta_k\}$ be the stratification of Δ built by Pearl's algorithm, and $\{\Delta'_1, \Delta'_2, \ldots, \Delta'_k\}$ be the one built by LC. As we noticed above, each Δ'_i sub-base contains exactly the default rules that are satisfied by bel_i, the εbf built by the LC algorithm at

[7] Upon reflection, this result is not so surprising. The εbf produced by the LC algorithm is a consonant belief function, hence it corresponds to a possibility measure; and, as shown by (Benferhat et al, 1992), possibility measures can be used in a natural way to generate a **Z**-equivalent ranking of a default base.

of EBF(Δ), and each element in the family induces the same ordering $\prec$ on the worlds in Ω. The latter property means that we can decide entailment by looking at just one element of $Bel_{lc}(\Delta)$.

Lemma 7. 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, $\prec_1 \equiv \prec_2$.

4.3 *The LC consequence relation*

We can use the set $Bel_{lc}(\Delta)$ to define a new non-monotonic consequence relation: we call *lc-consequence* this relation, and denote it by $\mid\!\sim_{lc}$. The definition of $\mid\!\sim_{lc}$ is similar to (4) in section 3.2, except that we now restrict the attention to the εbf-models that are in $Bel_{lc}(\Delta)$:

$$\alpha \mid\!\sim_{lc} \beta \quad \text{iff} \quad \text{for any } bel_{\mathcal{E}} \text{ in } Bel_{lc}(\Delta),\ bel_{\mathcal{E}} \models \alpha \rightarrow \beta. \tag{6}$$

Note that, as $Bel_{lc}(\Delta)$ is a subset of EBF(Δ), lc-consequence is stronger than bf-consequence; that is, if $\alpha \mid\!\sim_{bf} \beta$, then also $\alpha \mid\!\sim_{lc} \beta$. Theorem 1 then tells us that lc-consequence is also stronger than system **P**. The following example shows that it is *strictly* stronger than **P**; in fact, we shall show in the next subsection that lc-consequence precisely coincides with Pearl's system **Z**.

Example 1 (continued). We can use the bel_2 εbf built above to verify 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 $\varepsilon_1 >_\infty \varepsilon_2$ by (5). Hence, by Lemma 2, $b \wedge p \mid\!\sim_{lc} \neg f$. Moreover, if r is a new propositional variable (e.g., "red"), we have $b \wedge r \mid\!\sim_{lc} f$: in fact, $pl_2(b \wedge r \wedge \neg f) = \varepsilon_1$, $pl_2(b \wedge r \wedge f) = 1$, and $1 >_\infty \varepsilon_1$. That is, and differently from system **P**, irrelevant properties like "red" do not block the inheritance of the "fly" property■

4.4 *A note on stratification*

In building the focal elements, the LC algorithm builds a collection of nested sets $\Gamma = \{\Gamma_1, \Gamma_2, \ldots, \Gamma_k\}$: we start with $\Gamma_1 = \Delta$, and then build each Γ_i by eliminating some defaults from Γ_{i-1}. The collection Γ induces a partition $\{\Delta_1, \Delta_2, \ldots, \Delta_k\}$ of Δ in the obvious way: we let $\Delta_i = \Gamma_i - \Gamma_{i+1}$, i = 1, ..., k–1, and $\Delta_k = \Gamma_k$, . For instance, in Example 1 above, we have $\Delta_1 = \{b \rightarrow f\}$ and $\Delta_2 = \{p \rightarrow b, p \rightarrow \neg f\}$. It is easy to realize that Δ_i contains exactly the defaults that are satisfied by the bel_i

Step 0. Let $i = 0$, $\Gamma_0 = \Delta$, $sat_0 = \varnothing$, m_0 s.t. $\{m_0(\Omega) = 1; m_0(X) = 0$ otherwise$\}$.

Step 1. Repeat until $\Gamma_i = \varnothing$

1a. Let $i = i+1$

1b. Let $\Gamma_i = \Gamma_{i-1} - sat_{i-1}$

1c. Let bel_i be the εbf given by:
$\{m_i(\Omega) = \varepsilon_i ; m_i(\phi_{\Gamma_i}) = m_{i-1}(\Omega) - \varepsilon_i ; m_i(X) = m_{i-1}(X)$ otherwise$\}$

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

1e. If $sat_i = \varnothing$ then Fail.

Step 2. Return bel_{i-1}.

Figure 1: The LC algorithm.

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 $\mathcal{E} = \{\varepsilon_1, \varepsilon_2, \varepsilon_2\}$ such that (5) is satisfied. We have $\Gamma_1 = \Delta$, and bel_1 given by: $m_1(\phi_{\Gamma_1}) = 1-\varepsilon_1$, $m_1(\Omega) = \varepsilon_1$, and $m_1(X) = 0$ otherwise, with $[\phi_{\Gamma_1}] = [(\neg b \vee f) \wedge (\neg p \vee b) \wedge (\neg p \vee \neg f)] = [b \wedge \neg p \wedge f] \cup [\neg b \wedge \neg p \wedge f] \cup [\neg b \wedge \neg p \wedge \neg f]$. Hence:

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

From this, we compute the set sat_1 of defaults which are satisfied by bel_1: $sat_1 = \{b \to f\}$ (using Lemma 3, and recalling that $1 >_\infty \varepsilon_1$ by Lemma A7(b)). We iterate, and get Γ_2 by removing $b \to f$ from Γ_1. Now, $[\phi_{\Gamma_2}] = [b \wedge \neg p \wedge f] \cup [\neg b \wedge \neg p \wedge f] \cup [b \wedge \neg p \wedge \neg f] \cup [\neg b \wedge \neg p \wedge \neg f] \cup [b \wedge p \wedge \neg f]$, and bel_2 is given by

$$m_2(\Omega) = \varepsilon_2;\ \ m_2(\phi_{\Gamma_2}) = \varepsilon_1-\varepsilon_2;\ \ m_2(\phi_{\Gamma_1}) = 1-\varepsilon_1;\ \ m_2(X) = 0 \text{ otherwise.}$$

Then,

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

All the defaults in Γ_2 are now satisfied (as $\varepsilon_1 >_\infty \varepsilon_2$ by (5)) and the algorithm ends returning bel_2. This is a consonant ε-belief function with focal elements $[\phi_{\Gamma_1}] \subseteq [\phi_{\Gamma_2}]$. ■

We denote by $Bel_{lc}(\Delta)$ the family of all the εbf’s that can be built by the above procedure; the elements of this family differ in the choice of the $\mathcal{E}$ set of infinitesimals, provided that (5) is satisfied. The following Lemma shows that this family “behaves well” for our goals: it is a subset because otherwise $\alpha \rightarrow \beta$ would be conflicting, and that $[\alpha \wedge \neg\beta] \cap [\phi_\Delta] = \emptyset$, so $pl_1(\alpha \wedge \beta) = 1-\varepsilon_1$ and $pl_1(\alpha \wedge \neg\beta) = \varepsilon_1$, with $1-\varepsilon_1 >_\infty \varepsilon_1$). If there are still conflicts, however, this εbf may not satisfy some of the defaults in Δ, so is not an εbf-model of Δ (see Example 1 below for an illustration). In this case, we consider the set Δ' of all the defaults that are not satisfied, and build a second focal element $[\phi_{\Delta'}]$. By construction, $[\phi_\Delta] \subseteq [\phi_{\Delta'}]$. We then move "almost all" the ε_1 mass currently on Ω to $[\phi_{\Delta'}]$, leaving a "small" ε_2 on Ω; that is, we build the εbf given by: $m_1(\phi_\Delta) = 1-\varepsilon_1$, $m_1(\phi_{\Delta'}) = \varepsilon_1-\varepsilon_2$, $m_1(\Omega) = \varepsilon_2$, and $m_1(X) = 0$ otherwise, with ε_2 an infinitesimal such that $\varepsilon_1 >_\infty \varepsilon_2$. By reasoning as above, it is easy to see that when there are no conflicts in Δ' this new εbf is an εbf-model of Δ' (and of Δ). If there are still conflicts in Δ', we iterate the procedure until the εbf so built satisfies all the defaults in Δ.

More precisely, let n be the cardinality of Δ, and let $\mathcal{E} = \{\varepsilon_1, ..., \varepsilon_n\}$ be a set of infinitesimals such that

$$\varepsilon_i >_\infty \varepsilon_{i+1} \text{ and } \varepsilon_i > \varepsilon_{i+1} \text{ for all } i = 1, ..., n-1. \tag{5}$$

(Recall that $\varepsilon_i > \varepsilon_{i+1}$ means $\varepsilon_i(\eta) > \varepsilon_{i+1}(\eta)$ for all $\eta \in (0,1)$.) Note that such a set always exists: for instance, we can take $\varepsilon_i = \delta^i$ for some arbitrary infinitesimal δ. From this set, we build an εbf using the LC algorithm shown in Figure 1. We start from a vacuous belief assignment (Step 0). At each step i, we add a new focal element $[\phi_{\Gamma_i}]$ corresponding to the words that satisfy the defaults in Δ not already satisfied by the current εbf, and give it a mass $\varepsilon_{i-1}-\varepsilon_i$, taken off from the mass currently given to Ω (with the convention $\varepsilon_0=1$). We repeat this process until the obtained εbf satisfies all the defaults in Δ. Failing to find an εbf occurs if Δ (more precisely, ϕ_Δ) is inconsistent (Step 1e). It is easy to see that the εbf returned in step 2 is indeed an εbf according to Definition 4; in particular, the requirement that $\varepsilon_i > \varepsilon_{i+1}$ for all i guarantees that, for any $\eta \in (0,1)$, all masses are in [0,1]. Notice that the focal elements are nested, the inner one being $[\phi_\Delta]$, and therefore the εbf so built is consonant. By Lemma 6, then this εbf belongs to Λ_C, the set of εLC belief functions that satisfy the set C of constraints imposed by the auto-deductivity requirement: $pl(\alpha \wedge \beta) >_\infty pl(\alpha \wedge \neg\beta)$ for each $\alpha \rightarrow \beta \in \Delta$.

Definition 8: ε-least commitment: Let bel_1 and bel_2 be two belief functions (classical or εbf). We say that bel_1 is εLC than bel_2 if:

1: $\lim_{\eta\to 0} pl_1(A) \geq \lim_{\eta\to 0} pl_2(A)$ for all $A \subseteq \Omega$, and at least one inequality is strict,

or

2: $\lim_{\eta\to 0} pl_1(A) = \lim_{\eta\to 0} pl_2(A)$ for all $A \subseteq \Omega$, and m_1 is a strict {0,1}-generalization of m_2 for every η in a neighborhood of 0.

To illustrate the meaning of the second requirement, consider two εbf bel_i with $m_i(A) = 1-\varepsilon_i$, $m_i(\Omega) = \varepsilon_i \in \mathbb{E}^0$, i=1,2, for some $A \subseteq \Omega$. Their limits are equal and both share the same focal elements A and Ω, so the first requirement does not apply. The second requirement is not satisfied as none of the two belief functions is a {0,1}-generalization of the other. Thus none of bel_1 and bel_2 is εLC than the other, even when $\varepsilon_2(\eta) > \varepsilon_1(\eta)$ for all η.

The εLC principle can be used to build εbf's that satisfy some constraints and that are minimally committed. Let C be a set of constraints that must be satisfied by the εbf in EBF(Δ). We denote by Λ_C the set of the non dominating elements of EBF(Δ) that satisfy the constraints in C where 'non-dominating' is derived from the use of the definition of εLC. The next lemma illustrates the construction of a Λ_C which εbf's have nested focal elements and that will be used later on.

Lemma 6: Consider a set $\{A_1,...,A_n\}$of n nested subsets of Ω, with $A_1 \subseteq A_2 \subseteq ... \subseteq A_n$, $A_1 \neq \emptyset$, $A_n=\Omega$. Let $C = \{C_1,...,C_n\}$ be a set of constraints C_i given by:

C_1 : $pl(A_1) = 1$, and

C_i : $pl(\overline{A}_{i-1} \cap A_i) >_\infty pl(\overline{A}_i)$, i=2,3...n.

Then

$$\Lambda_C = \{bel\text{: } m(A_i) = \varepsilon_i \in \mathbb{E}^0,\ i=2,3...n,\ \varepsilon_i >_\infty \varepsilon_{i+1},\ i=2,3...n\text{-}1, \text{ and } m(A_1) = 1\text{-}\sum_{i=2}^{n}\varepsilon_i\}.$$

4.2 A peeling algorithm based on ε-least-commitment

We show how to build an εbf-model of Δ that satisfies the ε-least-commitment principle. A similar approach has been proposed by Benferhat et al (1992) using possibility measures. Intuitively, we proceed as follows. We start by the vacuous εbf, given by: $m_0(\Omega)=1$, and $m_1(X)=0$ elsewhere. This is the least committed of all the εbf, but clearly it is not an εbf-model of a non-empty Δ, as we do not have $pl_0(\alpha\wedge\beta) >_\infty pl_0(\alpha\wedge\neg\beta)$ for all $\alpha\to\beta \in \Delta$. So, we build the focal element $[\phi_\Delta]$ consisting of the worlds where all the defaults in Δ are satisfied, and move "almost all" the mass currently on Ω to $[\phi_\Delta]$, leaving a small ε_1 on Ω. That is, we build the εbf given by: $m_1(\phi_\Delta) = 1-\varepsilon_1$, $m_1(\Omega) = \varepsilon_1$, and $m_1(X) = 0$ elsewhere, where ε_1 is an infinitesimal. When there are no conflicting defaults in Δ, this εbf is an εbf-model of Δ (take any $\alpha\to\beta \in \Delta$, and note that $[\alpha\wedge\beta] \cap [\phi_\Delta] \neq \emptyset$

The notion of a (classical) belief function being minimally informative is defined in (Smets, 1988). The *least-commitment principle* states that, in order to model by a belief function bel an item of information as "all what is known", we should use the least committed belief function that is compatible with that information. By recalling the definition of least commitment in Section 2, this means that, for any $X \subseteq \Omega$, the value of bel(X) should be as small as possible — said differently, the agent should not give any proposition more belief than justified by what he/she knows. Recall that, when the item of information is represented by a propositional formula α, the least committed belief function representing it is the simple support function that gives a mass 1 to $[\alpha]$ and 0 anywhere else.

An item of information is usually compatible with a family of belief functions, denoted $\mathcal{B}$, and the application of the least-commitment principle consists ideally in finding the element of $\mathcal{B}$ that is less committed than the other elements of $\mathcal{B}$. Unfortunately usually such a unique element does not exist. The application of the least-commitment principle produces only the least committed set $\Lambda_{\mathcal{B}}$ of the non dominating elements of $\mathcal{B}$ (see Section 2.3)

Once εbf are used, the concept of least commitment must be extended into the concept of ε–least commitment (εLC).[6] Since all we are interested in with εbf is their behavior when η tends to 0, we focus on the εbf when η is in the neighborhood of 0 (to say that a property holds when η is in the neighborhood of 0 means that there is a $\delta>0$ such that the property holds for every $0<\eta<\delta$). One way to define εLC is based on a special form of 'generalization', where all coefficients belong to {0, 1} (Klawonn and Smets, 1991). Consider two basic belief assignments m_1 and m_2 where the masses of m_1 are obtained from the masses of m_2 in such a way such that, for each $A \subseteq \Omega$, the mass $m_2(A)$ is reallocated by m_1 to some superset of A. Let $\Pi = \{\Pi_1, ...\Pi_n\}$ be a partition of the subsets of Ω, and for each element Π_i of Π, let B_i be a set of Ω that contains all the subsets of Ω in Π_i. For all these B_i's, let

$$m_1(B_i) = \sum_{A \in \Pi_i} m_2(A) \qquad\qquad \text{if } \mathcal{A}_B \neq \emptyset,$$

and let all the other values of m_1 be null. So $m_2(A)$ is reallocated by m_1 to a superset of A. A consequence of this relation between m_1 and m_2 is that $pl_1(A) \geq pl_2(A)$ for all $A \subseteq \Omega$, what means that bel_1 is less committed than bel_2. m_1 is called a {0,1}-generalization of m_2, and if $m_1 \neq m_2$, then it is 'strict'.

In the following definition of εLC, the first requirement encompasses the case when infinitesimals are not involved, the second one is necessary when εbf are involved and their limits are equal when η tends to 0.

---

[6] The shorthand εLC is used both for 'ε-least commitment' and for 'ε-least committed'.

Lemma 5. Let $bel_{\mathcal{E}}$ be an εbf on Ω. For any α, β formulae of $\mathcal{L}$, $bel_{\mathcal{E}} \models \alpha \rightarrow \beta$ if, and only if, each $bel_{\mathcal{E}}$-preferred world of α satisfies β.

Characterizing bf-consequence in a preferential framework is useful for relating our systems to other existing ones by making use of known results. For example, we can show that a specific εbf-based system is equivalent to an existing preferential system by showing that they share the same preference relation: we shall make extensive use of this technique in Section 5. As a more immediate result, we show that bf-consequence is 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, $\alpha \mid\!\sim_{\mathbf{P}} \beta$.

Proof. (Sketch, details in Appendix B.) [⇒] Infinitesimal probability distributions are a special case of εbf where only singletons are assigned a non-zero mass. Thus, bf-consequence is a subset of Adams' ε-consequence, which is equivalent to **P**-consequence. [⇐] Lemma 4 tells us that the inference relation induced by any $bel_{\mathcal{E}}$ in EBF(Δ) is preferential, and therefore it satisfies all the rules of **P**. As this is true for any $bel_{\mathcal{E}}$, then $\mid\!\sim_{bf}$ satisfies the rules of **P**, and so it contains all the preferential consequences of Δ. ■

This result tells us two things. First, we can use (epsilon) belief functions to give semantics to the system **P**, just as Adams used (epsilon) probabilities. Second, $\mid\!\sim_{bf}$ suffers from the same problems as system **P**, namely, the problems of irrelevance and of blocking of inheritance; that is, bf-consequence is too weak with respect to our desiderata. In the rest of this paper, we explore several ways to define bolder consequence relations by restricting our attention in (4), (section 3.2 above) to just *some* of the εbf-models in EBF(Δ).

4 Using the Least-Commitment Principle

We have seen that any ε-belief function induces a specific (preferential) order over the possible worlds in Ω. In particular, any εbf-model of Δ induces one such order which is compatible with all the defaults in Δ (in the sense specified by the ⊨ relation). One way to select some of the εbf-models of Δ is by imposing some constraints on this order. A reasonable constraint is to require it to be minimally informative: intuitively, we want to look at the consequences of "only knowing" Δ (and nothing more).

4.1. ε-least commitment

definitions 3, 4 and 5: together, these constitute the core of our approach to defining non-monotonic systems using (epsilon-) belief functions.

3.3 A preferential view

Non-monotonic consequence relations are often characterized in terms of preferential semantics (Shoham, 1988; Kraus et al., 1990). In a nutshell, a preferential model is a structure W = (S, f, <) where S is a finite set of states, f maps each state to a world, and < is a strict partial order over S. The < order is meant to capture the idea that a state is "more normal" than another one; < is sometimes called a *preference relation*. Given a preferential model W, a state s is called a minimal, or *preferred*, state for a formula ϕ iff: (i) $f(s) \models \phi$, and (ii) there is no state s' in S such that s'<s and $f(s') \models \phi$. We may call s a <-preferred state for ϕ if we need to make the order < explicit. A formula ψ is a consequence of ϕ in W, denoted by $W \models \phi \rightarrow \psi$, iff each preferred state s for ϕ is such that $f(s) \models \psi$.

Our belief function semantics can be given an equivalent characterization in a preferential framework. More specifically, we can associate each εbf $bel_{\mathcal{E}}$ with a preferential order among worlds in Ω as follows.

Definition 6. Let $bel_{\mathcal{E}}$ be an εbf on Ω, and let ω_1 and ω_2 be two worlds in Ω. We say that ω_1 is *$bel_{\mathcal{E}}$-preferred to* ω_2, and write $\omega_1 \prec_{\mathcal{E}} \omega_2$, iff $pl_{\mathcal{E}}(\omega_1) >_{\infty} pl_{\mathcal{E}}(\omega_2)$. We call the strict partial order $\prec_{\mathcal{E}}$ the *$bel_{\mathcal{E}}$-preference relation.*

The order $\prec_{\mathcal{E}}$ can be used to define the set of preferred worlds, according to $bel_{\mathcal{E}}$, of a given formula.

Definition 7. Let $bel_{\mathcal{E}}$ be an εbf on Ω, and α a formula of $\mathcal{L}$. A world ω of Ω is a *$bel_{\mathcal{E}}$-preferred world* for α if: (1) ω satisfies α; and (2) there is no other world ω' that satisfies α such that $\omega' \prec_{\mathcal{E}} \omega$.

Recall that, by definition, the εbf-model relation $\models$ is characterized by the $>_{\infty}$ relation between values of plausibility. As the latter is closely related to $\prec_{\mathcal{E}}$, we may expect that the $\models$ relation can be defined in a preferential way via $\prec_{\mathcal{E}}$. The following lemma shows that this is indeed the case.

Lemma 4. Let $bel_{\mathcal{E}}$ be an εbf, and let $\alpha \rightarrow \beta$ be a default rule.

(i) $bel_{\mathcal{E}} \models \alpha \rightarrow \beta$ iff $\max_{\omega \models \alpha \wedge \beta} pl_{\mathcal{E}}(\omega) >_{\infty} \max_{\omega \models \alpha \wedge \neg \beta} pl_{\mathcal{E}}(\omega)$.

(ii) $bel_{\mathcal{E}} \models \alpha \rightarrow \beta$ iff $\lim_{\eta \rightarrow 0} bel_{\mathcal{E}}(\beta|\alpha) = 1$.

The notion of εbf-model can be extended to a full base Δ of default rules in the obvious way:

Definition 5. We say that $bel_{\mathcal{E}}$ is an εbf-model of Δ, written $bel_{\mathcal{E}} \models \Delta$, iff $bel_{\mathcal{E}}$ is an εbf-model of each rule in Δ. We denote by EBF(Δ) the set of all the εbf-models of Δ.

Thus, EBF(Δ) is a subset of the set of all the ε-belief functions $bel_{\mathcal{E}}$ such that $bel_{\mathcal{E}} \models \alpha \rightarrow \beta$ for any default $\alpha \rightarrow \beta$ in Δ.

Our next step is to use εbf models to define a notion of *entailment* for default bases, i.e., to define which conditional assertions $\alpha \rightarrow \beta$ are entailed by a default base Δ. A natural way to do this is via the usual definition of logical entailment: Δ entails $\alpha \rightarrow \beta$, written $\Delta \models \alpha \rightarrow \beta$, iff every εbf that is an εbf-model of Δ is also an εbf-model of $\alpha \rightarrow \beta$. That is,

$$\Delta \models \alpha \rightarrow \beta \quad \text{iff} \quad \text{for any } bel_{\mathcal{E}} \text{ in EBF}(\Delta),\ bel_{\mathcal{E}} \models \alpha \rightarrow \beta. \tag{3}$$

By construction, every rule in Δ is entailed from Δ itself, a property we will call the "auto-deductivity property".

Recall that the default base Δ is meant to represent our background knowledge. Once Δ is fixed, (3) tells us which formulae are to be considered "normal consequences" of any given formula α — namely, all the formulae β such that $\Delta \models \alpha \rightarrow \beta$ holds. One goal in this paper is to study the properties of the notion of "normal consequence" so obtained. We call *bf-consequence* this consequence relation, and denote it by $\mid\!\sim_{bf}$. It follows from (3) that, for a given Δ and for any α, β ∈ $\mathcal{L}$,

$$\alpha \mid\!\sim_{bf} \beta \quad \text{iff} \quad \text{for any } bel_{\mathcal{E}} \text{ in EBF}(\Delta),\ bel_{\mathcal{E}} \models \alpha \rightarrow \beta. \tag{4}$$

As we show in the next subsection, the $\mid\!\sim_{bf}$ relation is not entirely satisfactory, and we shall define several more specific non-monotonic consequence relations in the rest of this paper. In each case, the definition will look like (4), with the EBF(Δ) set replaced by some smaller set of εbf-models. These subsets will be obtained by imposing constraints on the form of the εbf's and/or on the ordering of the infinitesimals that they contain. In all cases, we shall use formula (4) and

Lemma 1 .[5] Let $m_{\mathcal{E}}$ be an ε-mass assignment. Then $m_{\mathcal{E}}(A) \in \mathbb{E}^1 \cup \{1\}$ for exactly one element $A \subseteq \Omega$, and $m_{\mathcal{E}}(X) \in \mathbb{E}^0 \cup \{0\}$ for all $X \subseteq \Omega$, $X \neq A$.

Moreover, under certain conditions, the result of applying Dempster's rule of combination to εbf's is itself an εbf, as stated by the following.

Lemma 2. Let m_1 and m_2 be two ε-mass assignments on Ω. Then $m_{12} = m_1 \oplus m_2$ is an ε-mass assignment, provided that the normalization factor in the combination is 1.

The hypothesis on the normalization factor means that m_1 and m_2 have no disjoint focal elements; as we shall see in Section 5, this is always true in our framework provided that the default base Δ is (classically) consistent.

Finally, we state a couple of properties that will be useful for working with εbf's.

Lemma 3. Let $bel_{\mathcal{E}}$ be an εbf on Ω. For any $X \subseteq \Omega$,

(a) $pl_{\mathcal{E}}(X) \approx_\infty \sum_{x \in X} pl_{\mathcal{E}}(\{x\})$.

(b) $pl_{\mathcal{E}}(X) \approx_\infty \max_{x \in X} pl_{\mathcal{E}}(\{x\})$.

3.2 εbf-entailment

We can use εbf's to give semantics to default rules in a spirit similar to Adams' use of infinitesimal probabilities. From now on, and unless stated otherwise, the frame of discernment Ω will be assumed to be the set of all the distinct truth assignments (or worlds) for our propositional language $\mathcal{L}$. Recall that we use [α] to denote the set $\{\omega \in \Omega \mid \omega \models \alpha\}$; and that we use the abbreviations $m_{\mathcal{E}}(\alpha)$ for $m_{\mathcal{E}}([\alpha])$, and $m_{\mathcal{E}}(\omega)$ for $m_{\mathcal{E}}(\{\omega\})$ (similarly for $bel_{\mathcal{E}}$ and $pl_{\mathcal{E}}$).

Definition 4. Let $bel_{\mathcal{E}}$ be an εbf, and let $\alpha \rightarrow \beta$ be a default rule. We say that $bel_{\mathcal{E}}$ is an *εbf-model* of $\alpha \rightarrow \beta$, and write $bel_{\mathcal{E}} \models \alpha \rightarrow \beta$, iff $pl_{\mathcal{E}}(\alpha \wedge \beta) >_\infty pl_{\mathcal{E}}(\alpha \wedge \neg \beta)$.

We can give εbf-models other equivalent definitions, like those stated by the next lemma: the first one has a more operational flavor; the second one uses conditional belief functions, and thus it is closer to the spirit of Adams' proposal.

[5] The proofs of all lemmas and theorems are reported in Appendix B.

defaults Δ if, for any ε-belief function bel that satisfies all the default rules in Δ (i.e., $bel(\beta_i|\alpha_i)$ is close to 1 for each rule $\alpha_i \rightarrow \beta_i$), $bel(\beta|\alpha)$ is close to 1. As we will see, this definition produces the same results as Adams' system.

3.1 Epsilon-belief functions

Our first step is to define ε-belief functions. Intuitively, these are functions whose values are infinitesimal, and which are belief functions when the argument η of these infinitesimals is fixed. Our aim is to study the limit behavior of these belief functions as η tends to 0, in which case their values get close to 0 or 1. We shall do this using the $>_\infty$ and $\approx_\infty$ relations introduced above.

Definition 3. Let Ω be a finite non-empty set. An *ε-mass assignment* on Ω is a function $m_{\mathcal{E}}: 2^\Omega \rightarrow \mathbb{E}$ such that, for all $\eta \in (0,1)$,

(i) $m_{\mathcal{E}}(\emptyset)(\eta) = 0$;

(ii) $\sum_{X \subseteq \Omega} m_{\mathcal{E}}(X)(\eta) = 1$.

The value of $m_{\mathcal{E}}$ depends on η, but for any fixed $\eta \in (0,1)$, $m_{\mathcal{E}}$ is a (standard) basic belief assignment on Ω. In fact, for all η, all the values $m_{\mathcal{E}}(X)(\eta)$ are in [0,1] due to our definition of infinitesimals (recall that, if $t \in \mathbb{E}$, then for any $\eta \in (0,1)$, $t \geq 0$ and $t \leq 1$); and they add up to 1 by definition. In practice, we will only consider the behavior of $m_{\mathcal{E}}$ when η tends to 0. As η approaches 0, all masses approach either 0 or 1. More precisely, the masses of all focal elements are infinitesimal, except for one whose limit is 1.

If $m_{\mathcal{E}}$ is an ε-mass assignment, for every η, we compute the belief function induced by $m_{\mathcal{E}}(.)(\eta)$ via (2, section 2.3), The set of belief functions so obtained is denoted by $bel_{\mathcal{E}}$, and called an *ε-belief function* (εbf, for short). We denote by $pl_{\mathcal{E}}$ the corresponding ε-plausibility function. Notice that the index ε of $m_{\mathcal{E}}$ is meant here to simply suggest the special nature of the masses. Later in the paper, we shall study εbf's which are built from some specific set of infinitesimals $\mathcal{E} \subseteq \mathbb{E}$. The index ε then will stand for this specific set.

The restrictions that we have imposed on infinitesimals as defined in Section 2.4 guarantee that an εbf exists. For example, $m_{\mathcal{E}}(A) = 1-\varepsilon$, $m_{\mathcal{E}}(\Omega) = \varepsilon$, and $m_{\mathcal{E}}(X) = 0$ elsewhere, for some $A \subseteq \Omega$ and $\varepsilon \in \mathbb{E}^0$. We call an assignment of this form an *ε-simple support function*.

pre-order $\geq_\infty$ defined as follows: for any $t_1, t_2 \in \mathbb{E}^0$, $t_1 \geq_\infty t_2$ iff $\lim_{\eta\to 0} \frac{t_2}{t_1} \in [0,\infty)$, what is equivalent to $\kappa(t_1) \leq \kappa(t_2)$ (see Lemma A4.a). Intuitively, $t_1 \geq_\infty t_2$ says that t_2 is at least as small as t_1.

The $\geq_\infty$ pre-order can be naturally extended to elements of $\mathbb{E}$ and to $\mathbb{R}^+$, the set of positive real numbers. For the particular case $t_1 = t_2 = 0$, we postulate $0 \geq_\infty 0$.

The $\geq_\infty$ relation induces two other binary relations on $\mathbb{E}\cup\mathbb{R}^+$ as follows:

• $t_1 >_\infty t_2$ iff $t_1 \geq_\infty t_2$ and not $t_2 \geq_\infty t_1$ (read "t_2 is infinitesimally smaller than t_1") or equivalently iff $\kappa(t_1) < \kappa(t_2)$,

• $t_1 \approx_\infty t_2$ iff both $t_1 \geq_\infty t_2$ and $t_2 \geq_\infty t_1$ (read "t_2 is as small as t_1") or equivalently iff $\kappa(t_1) = \kappa(t_2)$

When $t_1 \approx_\infty t_2$, we also say that t_1 and t_2 are *of the same order*.

The inequality $>_\infty$ is irreflexive and transitive, and $\approx_\infty$ is reflexive, symmetric and transitive, i.e., it is an equivalence relation (see the Appendix for proofs).

We also use the standard mathematical relations and operators on elements of $\mathbb{E}$, meaning pointwise application. For example, by $t_1 > t_2$ we mean $t_1(\eta) > t_2(\eta)$ for every $\eta \in (0,1)$.

In this paper, we use the ordered set $\langle \mathbb{E}\cup\mathbb{R}^+, >_\infty\rangle$ to distinguish different orders of infinitesimals. Note that 0 lays at the bottom of this structure, and all the elements of $\mathbb{E}^1\cup\mathbb{R}^+$ at the top. More precisely, for any $t\in\mathbb{E}^0$ and $a\in\mathbb{R}^+$, we have $a >_\infty t$ and $t >_\infty 0$ — that is, elements of $\mathbb{E}^0$ are infinitesimally larger than 0 (in the sense of $>_\infty$), but infinitesimally smaller than any real number. Moreover, if $t'\in\mathbb{E}^1$, we have $t' >_\infty t$ and $t' \approx_\infty a$, that is elements of $\mathbb{E}^1$ are infinitesimally larger then the elements of $\mathbb{E}^0$, and they are of the same order as positive real numbers. (In fact, all real numbers are of the same order: $a \approx_\infty b$ for all $a,b\in\mathbb{R}^+$; intuitively, no positive real number is infinitesimally larger than any other.) Appendix A reports more properties.

3 A Belief Function Semantics for Defaults

In this section, we adapt Adams' ε-semantics to a belief function framework. First, we introduce the notion of ε-belief functions, whose values are either (infinitesimally) close to 0 or close to 1. Next, we interpret a default rule $\alpha_i\to\beta_i$ as meaning that the conditional belief $bel(\beta_i|\alpha_i)$ is close to 1. Finally, we define a consequence relation $\mid\!\sim$ in a natural way: $\alpha \mid\!\sim \beta$ follows from a base of class of belief functions that we are interested in. Our treatment is grounded in standard analysis; technical details are given in Appendix A.

We define first the *order* of a function and the concept of an *infinitesimal* as used in this paper.

Definition 1: Let f be a continuous function from $\mathbb{R}$ to $\mathbb{R}$. A non negative integer k is called the *order of f*, denoted $\kappa(f)$, if and only if $\lim_{\eta\to 0} \frac{f(\eta)}{\eta^k} \in \mathbb{R}\text{-}\{0\}$ when such a limit exists.

If there is a k that satisfies the previous definition, then this k is unique (to see this, note that the above limit is infinite for any n >k, and it is zero for any n < k). If there is no such k, we say that the order of f is undefined. We also extend the notion of order to real constants $c \in R$ by postulating: $\kappa(c) = 0$ if $c \neq 0$; and $\kappa(0) = \infty$ (addition, minimum and exponentiation are extended in order to include infinity in the obvious way).

Definition 2: An *infinitesimal* ε is a real continuous function from (0,1) to (0,1) such that :

1) $\lim_{\eta\to 0} \varepsilon(\eta) = 0$,
2) the order $\kappa(\varepsilon)$ is defined.

We limit ourselves to the infinitesimals with limited domains and ranges and that admit an order, which is a positive integer, since these provide all we need in this paper. Our definition of $\kappa(\varepsilon)$ fits with the intuitive idea that the larger the order of ε, the 'closer' ε is from 0 in the neighborhood of 0.

We denote by $\mathbb{E}^0$ the set of all infinitesimals so defined.

We then define $\mathbb{E}^1 = \{1-\varepsilon \mid \varepsilon \in \mathbb{E}^0\}$. Intuitively, $\mathbb{E}^1$ contains elements which are infinitesimally close to 1, i.e., for any $t \in \mathbb{E}^1$, $\lim_{\eta\to 0} t(\eta) = 1$. Finally, we define the set $\mathbb{E}$ by

$$\mathbb{E} = \mathbb{E}^0 \cup \mathbb{E}^1 \cup \{0\} \cup \{1\}.$$

By construction, the elements of $\mathbb{E}$ are functions t of η whose value tends to either 0 (for $t \in \mathbb{E}^0$) or to 1 (for $t \in \mathbb{E}^1$) as η tends to 0, plus the 0 and 1 constants. These are the objects that we shall use as values for our extreme belief functions.

Although the elements of $\mathbb{E}^0$ are "small" in the above sense, some elements are "smaller" than others. For instance, if ε is an infinitesimal, then ε^2 is smaller than ε, meaning that the ratio $\frac{\varepsilon^2}{\varepsilon}$ is itself infinitesimal: we say that ε^2 is *negligible* with respect to ε. The ability to rank infinitesimals into classes of "smallness" will be pivotal to our study. To do so, we equip the set $\mathbb{E}^0$ with a total of the two pieces of evidence is represented by the belief function $bel_1 \oplus bel_2$ obtained by Dempster's rule of combination, and described by the basic belief assignment

$$m_1 \oplus m_2(A) = \frac{1}{c} \sum_{B \cap C = A} m_1(B).m_2(C) \quad A \subseteq \Omega.$$

$$\text{where } c = 1 - \sum_{B \cap C = \emptyset} m_1(B).m_2(C) \quad .$$

The c factor is called the normalization factor. Note that the $\oplus$ operator is commutative and associative, so the order in which we combine pieces of evidence is irrelevant.

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 specifically supporting the subset X of Ω, now supports $X \cap A$. This transfer of belief masses is described by the following relation, where bel(.|A) and pl(.|A) denote the conditional belief and plausibility functions obtained after revising m with the new evidence A:

$$bel(X|A) = \frac{bel(X \cup A^c) - bel(A^c)}{1 - bel(A^c)} \quad ,$$

$$pl(X|A) = \frac{pl(X \cap A)}{pl(A)}.$$

A formal justification of both Dempster's rule of combination and Dempster's rule of conditioning can be found in (Smets, 1990, 1997b) and (Klawonn and Smets, 1991). Distinctness is defined in (Smets, 1992).

2.4. Infinitesimals and the induced negligibility relation

In this work, we consider "extreme" belief functions which values are "infinitesimally close" to 0 or to 1, much in the spirit of Adams' system. Before we embark in the description of these belief functions, we need to give precise meaning to the expression "infinitesimally close to 0 or to 1". The study of such expressions is the object of non-standard analysis: we address the reader to (Robinson, 1965), (Keisler, 1976) and (Nelson, 1977) for some foundational work; and to (Lehmann and Magidor, 1992), (Weydert, 1995) and (Wilson, 1996) for applications to probability theory. For the limited goals of this paper, however, we do not need a full theory of non-standard analysis. In what follows, we outline the elements that we need in order to precisely define the where A^c denotes the set-theoretic complement of A with respect to Ω.

It is important to emphasize that our approach is based on the transferable belief model (Smets, 1997a), and not on upper and lower probabilities (Walley, 1991). Accordingly, the term bel(A) should be read as the belief given by the agent to the fact that the actual world belongs to A, and not as the lower envelop of a family of probability functions representing the agent's belief about which world is the actual world.

If all the focal elements of a basic belief assignment are singletons, then bel is a probability measure, and bel = pl. If 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*; in which case, for all $A, B \subseteq \Omega$,

$$\text{bel}(A \cap B) = \min(\text{bel}(A), \text{bel}(B)); \text{ and}$$
$$\text{pl}(A \cup B) = \max(\text{pl}(A), \text{pl}(B)).$$

In this case, bel is a necessity measure and pl is a possibility measure (Zadeh, 1978; Dubois and Prade, 1988). If m has at most one focal element $A \neq \Omega$, $A \neq \emptyset$, i.e.,

$$m(A) = s, \ m(\Omega) = 1-s, \ m(\text{elsewhere}) = 0, \quad s \in [0,1],$$

then its related belief function is called a *simple support function*.

Belief functions can be partially ordered by their strength. Let bel_1 and bel_2 be two belief functions over Ω. We say that bel_1 is *less committed* than bel_2 (or bel_2 dominates bel_1 where 'dominating' means 'more informative') iff, for all $A \subseteq \Omega$, $\text{pl}_1(A) \geq \text{pl}_2(A)$ with at least one strict inequality. When $m_1(\emptyset) = m_2(\emptyset) = 0$, as required in this paper, this definition is equivalent to $\text{bel}_1(A) \leq \text{bel}_2(A)$ for all $A \subseteq \Omega$ with at least one strict inequality. Intuitively, bel_1 poses less stringent constraints than bel_2 on which world could be the actual one. The *vacuous belief function*, defined by the basic belief masses $m(\Omega) = 1$ and $m(A) = 0$ for all $A \neq \Omega$, is the least committed of all the belief functions on Ω.

Given a family $\mathfrak{B}$ of belief functions, we define the *least committed set* $\Lambda_{\mathfrak{B}}$ as those belief functions in $\mathfrak{B}$ that are not dominating another belief function in $\mathfrak{B}$:

$$\Lambda_{\mathfrak{B}} = \{ \text{bel}: \text{bel} \in \mathfrak{B}, \nexists \text{bel}' \in \mathfrak{B}, \text{bel}' \text{ less committed than bel} \}.$$

The theory of belief functions provides mechanisms to account for the *dynamics* of belief states, that is, the way belief states are modified in the light of evidence. Let E be a piece of evidence that bears on Ω. The impact of E is represented by a belief function that describes the agent's beliefs on Ω given E (and nothing else). Suppose that the agent receives two distinct pieces of evidence E_1 and E_2, and let bel_1 and bel_2 be the induced belief functions. The combined effect

2.3 Belief functions

The semantics for default reasoning that we develop in this paper is based on the formalism of belief functions. In order to make this paper self-contained, we recall here the basic definitions. For a more complete exposition, we refer the reader to (Shafer, 1976; Smets, 1988; Smets and Kennes, 1994, Kohlas and Monney, 1995, Smets, 1998).

Let Ω be a finite set of worlds, one of them being the actual world. A (normal[4]) *basic belief assignment* on Ω is a function m: $2^\Omega \rightarrow [0,1]$ that satisfies,

$$m(\emptyset) = 0,$$

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

The term m(A), called the *basic belief mass* given to $A \subseteq \Omega$, 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 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 be equivalently represented by the function bel: $2^\Omega \rightarrow [0,1]$, called a *belief function*, defined, for any $A \subseteq \Omega$, by

$$\mathrm{bel}(A) = \sum_{B:\, B\subseteq A} m(B) \quad . \tag{2}$$

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

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

The term pl(A) quantifies the degree of plausibility, of potential support, that the actual world belongs to A. It contains those parts of beliefs given to propositions that do not contradict A. Note that

$$\mathrm{pl}(A) = 1 - \mathrm{bel}(A^c),$$

---

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

problems of irrelevance and of inheritance blocking: in the bird and penguin example, **P** cannot deduce that a given *red* bird flies, and it cannot deduce that penguins have legs.

2.2 Probabilistic semantics for default rules

Adams (1966, 1975), and later Pearl (1988) and Lehmann and Magidor (1992), have suggested a probabilistic interpretation where a default rule $\alpha \rightarrow \beta$ is read as the constraint $P(\beta|\alpha) > 1-\varepsilon$, with P a probability distribution over Ω.[2] Given a set of defaults Δ and a given ε, they construct a class of probability distributions A_ε such that, for each distribution P in A_ε and each default $\alpha_i \rightarrow \beta_i$ in Δ, $P(\beta_i|\alpha_i) > 1-\varepsilon$ (recall that we write α as an abbreviation for $[\alpha]$). When A_ε is empty, the default base Δ is inconsistent. For example, $\Delta = \{\alpha \rightarrow \beta, \alpha \rightarrow \neg\beta\}$is inconsistent. Note that if Δ is classically inconsistent, it is also inconsistent. In the main part of this paper, we only deal with consistent default bases. In that case, 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, for each default $\alpha_i \rightarrow \beta_i$ in Δ, $P(\beta_i|\alpha_i)$ is very high. Parsons and Bourne (2000) and Gilio (2000) have computed optimal bounds associated with proofs for default consequences of System **P**. Adams and Pearl interpret ε as a quantity that can be arbitrary small, while Lehmann and Magidor interpret it as an infinitesimal positive number. In section 2.4, we give a formal definition of the infinitesimals used in this paper. In general, probabilistic semantics for default rules are based on probability values that are either close to 0 or close to 1. Pearl (1988) and Dubois and Prade (1995) have shown that if we allow arbitrary (not necessarily extreme) values, then the corresponding consequence relation does not satisfy the rationality postulates defined by (Kraus et al., 1990). Alternative probabalistic semantics for system **P** have been proposed in Snow(1999), Schurz (1998), and standard probabilistic semantics for System **P** has been proposed in (Benferhat et al., 1999).

Lehmann and Magidor (1992) have shown that ε-consequence is equivalent to system **P**.[3] This means that Adams' construction gives a (probability-based) semantics for **P**; unfortunately, this also means that this construction suffers from the same problems as **P**, namely, irrelevance and inheritance blocking.

---

[2] Interestingly, Dubois and Prade (1995) pointed out the similarity between a default rule $\alpha \rightarrow \beta$ and the *conditional object* $\beta|\alpha$, which can be seen as a symbolic counterpart to the conditional probability $P(\beta|\alpha)$.

[3] The equivalence holds only for consistent default bases. In the case of inconsistencies, Adams has also considered a further rule ""from $\alpha \mid\!\sim \perp$ infers $\beta \mid\!\sim \delta$" (which cannot be recovered from System P). In our paper, we have only considered finite and consistent default bases and hence the technical differences between Adams proposal and Lehmann's framework (concerning the handling of inconsistent default bases) disappear.

$d = \alpha \rightarrow \beta$, we denote by ϕ_d the formula of $\mathfrak{L}$ obtained by replacing $\rightarrow$ by the material implication, that is, $\phi_d = \neg\alpha \vee \beta$. If E is a set of defaults, we let $\phi_E = \wedge_{d \in E} \phi_d$. A *default base* is a multiset $\Delta = \{\alpha_i \rightarrow \beta_i,\ i=1,...,n\}$ of defaults. We emphasize that a default base is a multiset rather than a set, i.e., $\{\alpha \rightarrow \beta\}$ is different from $\{\alpha \rightarrow \beta,\ \alpha \rightarrow \beta\}$.

A default base can contain defaults that hint at opposite conclusions given the same premise. For example, consider 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"); given the fact $b \wedge p$, the first default supports the conclusion "f", while the third one supports "$\neg$f" — in fact, the formula $\phi_\Delta \wedge (b \wedge p)$ is classically inconsistent. We say that the defaults in Δ are *conflicting*. If the formula ϕ_Δ is itself inconsistent, we say that Δ is *classically inconsistent*. In general, no non-trivial result can be derived from a classically inconsistent default base.

We use default bases to represent background knowledge about what is normally the case. Given a base Δ, we are interested in defining a consequence relation $\mid\sim_\Delta$ between formulae of $\mathfrak{L}$ that tells us which consequences we can "reasonably" draw from given facts, given the background knowledge Δ. Our goal is to define $\mid\sim_\Delta$ so that it fulfils the desiderata listed in the Introduction. For example, given the base Δ above, we would like to have $b \mid\sim f$ and $b \wedge p \mid\sim \neg f$, but not $b \wedge p \mid\sim f$ (we omit the Δ subscript when this is clear from the context). This example shows that $\mid\sim$ should be non-monotonic; other desirable formal properties for the $\mid\sim$ consequence relation have been discussed, for instance, by Gabbay (1985), Kraus, Lehmann and Magidor (1990), Lehmann and Magidor (1992) and Gärdenfors and Makinson (1994).

In particular, Kraus, Lehmann and Magidor (1990) have proposed a set of postulates, known as the KLM postulates, that are commonly regarded as the minimal core of any "reasonable" non-monotonic system, and defined a non-monotonic system, called **P** (for "Preferential"), based on the following six postulates:

1. Reflexivity:	$\alpha \mid\sim \alpha$
2. Left Logical Equivalence (LLE):	from $\alpha \models \alpha'$, $\alpha' \models \alpha$ and $\alpha \mid\sim \beta$ deduce $\alpha' \mid\sim \beta$
3. Right Weakening (RW):	from $\beta \models \beta'$ and $\alpha \mid\sim \beta$ deduce $\alpha \mid\sim \beta'$
4. OR:	from $\alpha \mid\sim \gamma$ and $\beta \mid\sim \gamma$ deduce $\alpha \vee \beta \mid\sim \gamma$
5. Cautious Monotony (CM):	from $\alpha \mid\sim \beta$ and $\alpha \mid\sim \gamma$ deduce $\alpha \wedge \beta \mid\sim \gamma$
6. Cut:	from $\alpha \wedge \beta \mid\sim \gamma$ and $\alpha \mid\sim \beta$ deduce $\alpha \mid\sim \gamma$

From these rules, a consequence relation $\mid\sim_{\mathbf{P}}$ can be defined for any given Δ by: $\phi \mid\sim_{\mathbf{P}} \psi$ if and only if $\phi \mid\sim \psi$ can be derived from Δ using the rules of system **P**. Unfortunately, system **P** turned out to be too weak to be satisfactory: for instance, it suffers from the above mentioned our desiderata, and we devote the rest of the paper to study ways to strengthen it. In Section 4, we use the least-commitment principle, and obtain a system equivalent to Pearl's system **Z**. In Section 5, we use Dempster's rule of combination, and obtain a family of systems that includes systems equivalent to several popular non-monotonic systems. These results are important, because they show that ε-belief functions provide a general framework for default reasoning that captures several existing systems (and possibly others) as particular cases. Each system is obtained by imposing certain constraints over the ε-belief functions. In Section 6, we study the system obtained when these constraints come from two general principles: least commitment and auto-deduction. We call **LCD** this new system. In Section 7, we analyze the behavior of **LCD**, and show that it correctly addresses the problems of specificity, of irrelevance, of inheritance blocking, of ambiguity, and of redundancy. Section 8 is devoted to a comparison of **LCD** with other popular systems; we show that **LCD** is strictly stronger than Kraus, Lehmann and Magidor's system **P**, but it is incomparable with all of the major extensions of **P** currently found in the literature. Section 9 contains some final remarks and suggestions for future work. More technical details are collected in two appendices.

2 Background

2.1 Default reasoning

We are interested in representing *default rules* (or, simply, *defaults*) of the form

normally, if we have α, then β is the case, (1)

where α and β are formulae of some underlying language $\mathfrak{L}$. In this paper, we assume that $\mathfrak{L}$ is a classical propositional language constructed from a finite set **V** of propositional symbols (denoted by lower case letters p, q, r,…) and the usual connectives $\wedge$ (conjunction), $\vee$ (disjunction), and $\neg$ (negation). The elements of $\mathfrak{L}$, or *formulae*, will be denoted by Greek letters α, β, δ,…. An *interpretation* for $\mathfrak{L}$ is an assignment of a truth-value in {T, F} to each formula of $\mathfrak{L}$ 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 a world ω *satisfies* a formula α, and write $\omega \models \alpha$, iff α is true in ω. We denote by $[\alpha]_\Omega$ the set of all the worlds in Ω that satisfy α. Throughout this paper, we shall normally work with objects of the type [α] (where Ω is fixed and understood), and identify [α] with the *proposition* denoted by the formula α. We sometimes use α as a shortcut for [α] when the context prevents any ambiguity. Finally, we say that α entails p, and write $\alpha \models \beta$ iff all the worlds that satisfy α also satisfy β, i.e., $[\alpha] \subseteq [\beta]$.

We write the default rule (1) as $\alpha \rightarrow \beta$, where α and β are formulae of $\mathfrak{L}$. Note that "→" is a non-classical arrow, and it should not be confused with material implication. Given a default

- *Syntax independence*: the consequences of a knowledge base should not depend on the syntactical form used to represent the available knowledge. In particular, they should not be sensitive to duplications of rules in the knowledge base; failure to do this is referred to as the problem of *redundancy*.

Since Reiter's paper (1980), many proposals for reasoning with default information have appeared in the literature. Some of them are based on the use of uncertainty models such as probability theory (Adams, 1975; Pearl, 1988), or possibility theory (Dubois & Prade, 1988; Benferhat et al., 1992). Up to now, however, no single system has been reported that fulfils 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; Smets and Kennes, 1994), 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 numerical values, whose origin is often an open question. Finding a representation free from such assessments would somehow avoid the problem of the origin of the numbers. The representation that we develop in this paper addresses this problem. Our starting point is an interpretation of default information that uses a class of infinitesimal belief functions, called ε-belief functions, whose non-null masses are either infinitesimally close to 0 or to 1.

The idea of using extreme values is not new in plausible reasoning. For instance, Adams (1975) uses extreme probabilities to encode default information; De Kleer (1990) and Poole (1993) apply extreme probabilities to diagnostics problems; and Wilson (1993) represents default rules by limits of belief functions. In this paper, we show that Adams' results for reasoning with default information based on probability functions can also be obtained with belief functions. Thanks to the greater expressive power and the greater flexibility of belief functions, we are able to build a powerful framework that can be specialized to capture any one of several existing systems. In particular, this framework encompasses Adams' system, the **P** system, the **Z** system, the possibilistic logic, the penalty logic, the lexicographic approach, Geffner's conditional entailment, and Brewka's preferred sub-theories. Moreover, we propose a new system based on this framework, called LCD, which is different from all these existing systems. LCD correctly addresses the problems of specificity, of irrelevance, of inheritance blocking, of ambiguity, and of redundancy, that are encountered among the other systems.

The rest of this paper is organized as follows. The next section fixes the notation, recalls a few notions of the theory of belief functions, and reviews Adams' ε-semantics. In Section 3, we introduce ε-belief functions, and show how to use them to define a non-monotonic consequence relation that is equivalent to Adams' ε-semantics. This relation turns out to be too cautious given flies. If we later learn that it is a penguin, however, we should withdraw this conclusion. Similar problems are encountered in diagnostic settings (e.g., De Kleer, 1990). 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.

- *Rationality*: the consequence operator used to generate plausible conclusions from a knowledge base should satisfy the rationality postulates proposed by Kraus, Lehmann and Magidor (1990) (see section 2.1).

- *Specificity*: results obtained from more specific classes should override those obtained from more generic ones (Touretzky, 1984). For example, from the default base Δ = {"Birds fly", "Penguins do not fly", "Penguins are birds"}, one should deduce that Tweety, who is both a penguin and a bird, does not fly, since penguins are a subclass of birds.

- *Irrelevance* : if a formula δ is a plausible consequence of α, and if a formula β is "irrelevant" to the problem) then δ should also be a plausible consequence of $\alpha \wedge \beta$. For example, given the default rule "birds fly", we should also deduce that "red birds fly" since being red is irrelevant to the property of flying, meaning that no conditional assertion in the default base deals with red things. This intuitive notion of being "irrelevant" is not precisely defined in the NMR literature — in fact, the very use of the word "irrelevant" here may be criticized. In this paper, we consider a narrow definition of irrelevance, and say that a formula β is irrelevant to a default base Δ if β is composed of propositional symbols that do not appear in Δ.

- *Property inheritance:* a sub-class that is exceptional with respect to some property should still inherit the other properties from super-classes, unless there is a contradiction. For example, from the default base Δ above plus the rule "Birds have legs", one should deduce that penguins have legs too, since having legs is not a conflicting property — the only conflicting property is flying. Failure to perform these deductions is referred to as the problem of *inheritance blocking*.

- *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. The most popular example is the so-called Nixon diamond: knowing that Quakers are pacifists, that Republicans are not pacifists, and that Nixon is both a Quaker and a Republican, one should not deduce that Nixon is a pacifist, nor that he is not.[1]

[1] Note that this is different from the situation in which the same argument supports both a conclusion and its contrary. For example, given the rules "$\alpha \rightarrow \beta$" and "$\alpha \rightarrow \neg\beta$", the argument α supports both β and its contrary $\neg\beta$. In the latter case, we are in the presence of *inconsistency*.

Published in: Artificial Intelligence 122(1-2):1-69, 2000

Belief functions and default reasoning*

S. Benferhat[1] **A. Saffiotti**[2] **P. Smets**[3]

benferha@irit.fr asaffio@aass.oru.se psmets@ulb.ac.be

[1]IRIT, Université Paul Sabatier
118 Route de Narbonne
Toulouse 31062 Cedex, France

[2]AASS, Dept of Technology
Orebro University
SE-70182 Orebro, Sweden

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

Abstract. We present a new approach to deal with default information based on the theory of belief functions. Our semantic structures, inspired by Adams' epsilon semantics, are epsilon-belief assignments, where mass values are either close to 0 or close to 1. In the first part of this paper, we show that these structures can be used to give a uniform semantics to several popular non-monotonic systems, including Kraus, Lehmann and Magidor's system **P**, Pearl's system **Z**, Brewka's preferred sub-theories, Geffner's conditional entailment, Pinkas' penalty logic, possibilistic logic and the lexicographic approach. In the second part, we use epsilon-belief assignments to build a new system, called **LCD**, and show that this system correctly addresses the well-known problems of specificity, irrelevance, blocking of inheritance, ambiguity, and redundancy.

1 Introduction

Default reasoning is a form of non-monotonic reasoning (NMR) that consists in drawing conclusions from (i) a set of general rules which may have exceptions (like "birds fly"), and (ii) a set of facts representing the available information (which is often incomplete). The conclusions so drawn are only plausible, and they can be subsequently revised in the light of the new information. In the canonical penguin example, if we know that Tweety is a bird, then we can conclude that it

---

* A preliminary short version of this paper appeared in the Proceedings of the Conference on Uncertainty in Artificial Intelligence, Montreal, CA, 1995, pp. 19-26.