Resource-bounded paraconsistent inference *

Pierre Marquis

CRIL-CNRS/Université d'Artois
rue de l'Université - S.P. 16
F-62307 Lens Cedex
FRANCE

Nadège Porquet

CRIL-CNRS/Université d'Artois
rue de l'Université - S.P. 16
F-62307 Lens Cedex
FRANCE

In this paper, a new framework for reasoning from inconsistent propositional belief bases is presented. A family of resource-bounded paraconsistent inference relations is introduced. Such inference relations are based on $S-3$ entailment, an inference relation logically weaker than the classical one and parametrized by a set S of propositional variables. The computational complexity of our relations is identified, and their logical properties are analyzed. Among the strong features of our framework is the fact that tractability is ensured each time $|S|$ is bounded and a limited amount of knowledge is taken into account within the belief base. Furthermore, binary connectives $\wedge$, $\vee$ behave in a classical manner. Finally, our framework is general enough to encompass several paraconsistent multi-valued logics (including $S-3$, J_3 and its restrictions), the standard coherence-based approach to inconsistency handling (based on the selection of consistent subbases) and some signed systems for paraconsistent reasoning as specific cases.

Keywords: Resource-bounded reasoning, inconsistency handling.

1. Introduction

Few would dispute the fact that classical entailment is not suited to commonsense inference. One of its main drawbacks is its inability to deal with inconsistency in a non trivial way. This is often referred to as *ex falso quodlibet sequitur* and classical entailment is said to be *explosive* (i.e., not paraconsistent): any formula is a logical consequence of a contradiction.

* This paper is an extended and revised version of the paper "Resource-bounded inference from inconsistent belief bases" appeared in the proceedings of the 17^{th} International Joint Conference on Artificial Intelligence (IJCAI'01), pages 103–108, 2001.

Many approaches have been proposed so far to address this issue (see [9] for a survey). Some of them, like belief revision, aim at preventing inconsistencies from being generated; other approaches aim at removing inconsistencies once they appeared. Complementary to them, many approaches deal with inconsistencies. Among them are various paraconsistent logics (see e.g., [26] [23] [4] [33] [44] [54] [55] [8] [11] [12] [2]), argumentative logics (see e.g., [28] [29] [13] [1]), belief merging (see e.g., [46] [57] [39] [40]) and the so-called coherence-based approach to inconsistency handling (see e.g., [56] [31] [34] [53] [5]).

In the latter approach, beliefs are represented by a belief base, i.e., a finite set of propositional formulas; some preferred consistent subbases of the belief base are first selected, then inference amounts to classical inference from some of them. When all preferred subbases are considered, the corresponding (skeptical) inference relations exhibit valuable logical properties, since the rules of system P [41] are satisfied [5] [6]; other valuable deduction rules may be satisfied as well, depending on the selection policy into consideration (see [17] for more details).

Unfortunately, such inference relations are intractable, their computational complexity varying from the first level to the second level of the polynomial hierarchy [18]. This is far from being surprising since these relations come down to classical entailment whenever the underlying belief base is consistent. Here appears a second major drawback of classical entailment to model common-sense inference: classical entailment is coNP-complete, hence intractable, i.e., there does not exist a deterministic polytime algorithm to decide classical entailment, unless $\mathsf{P} = \mathsf{NP}$.

Dealing with these two drawbacks simultaneously is very important, both from the theoretical side and from the practical side since actual knowledge bases are often inconsistent and so large that the computational aspects cannot be neglected. Alltogether, they were some of the basic motivations for the development of nonmonotonic logics twenty years ago [43] [35].

In this paper, we show that trivialization in presence of inconsistency and intractability can be handled within a uniform framework. A family of paraconsistent resource-bounded inference relations for reasoning from inconsistent belief bases is presented. Our starting point is the notion of $S-3$ inference [58] and the coherence-based approach to inconsistency handling [53] [5]. In $S-3$ logic, every propositional symbol is interpreted in a weak way, allowing it to be both true and false at the same time, except for the variables of a given set S that are interpreted classically. Tractability is ensured by limiting the size of S. Interestingly, $S-3$ entailment proves sufficient to handle some inconsistencies, those that are not reachable when classical inference is limited to the variables of S.

In order to deal with the remaining inconsistencies while avoiding trivialization, we introduce a family of inference relations which are to $S-3$ entailment what the inference relations of the standard coherence-based approach are to classical entailment. Specifically, consistency is restored by removing variables (from S) instead of removing some explicit beliefs from the belief base. Among the strong features of our framework is the fact that tractability is ensured each time a computational bound $|S|$ is set, and a limited amount of knowledge is taken into account within the belief base. Moreover, binary connectives behave in a classical manner (while it is not the case for many paraconsistent logics). Contrariwise to the standard coherence-based approach to inconsistency handling, there is no need of a special (non truth functional) comma connective in our approach. Compared with $S-3$ logic [58], our framework is based on the more expressive language of paraconsistent logic J_3 [26], allowing integrity constraints to be encoded in a satisfying way. In addition, the explicit handling of preferences over S (which is not offered in $S-3$ logic) enables some forms of graded paraconsistency to be achieved. Finally, our framework is general enough to encompass several paraconsistent multi-valued logics (including J_3 [26] and its restrictions LP and LP_m [54] [55]), $S-3$ logic [58], the standard coherence-based approach to inconsistency handling [53] [5] and some systems for paraconsistent reasoning given in [12] as specific cases.

The rest of this paper is organized as follows. In Section 2, the standard coherence-based approach to inconsistency handling and the notion of $S-3$ inference are briefly presented. In Section 3, we introduce our family of resource-bounded inference relations. We analyze these inference relations, both from the computational and the logical side, and we compare them with the corresponding inference relations of the standard coherence-based approach. Section 4 shows connections with some other related work. Finally, Section 5 concludes the paper. Proofs of propositions are reported in an Appendix.

2. Belief bases, the standard coherence-based approach and $S-3$ inference

2.1. Belief bases and the standard coherence-based approach

Before presenting our contribution, we first need to briefly recall what is the standard coherence-based approach to inconsistency handling since it will serve as a base line. In this approach, beliefs are represented under the form of belief bases, possibly stratified. Each piece of belief is encoded as a propositional formula from the language $PROP_{PS}$ inductively defined in the usual way, from a finite set PS of propositional symbols (also called variables) and the connectives $\neg$, $\wedge$, $\vee$[1]. For every formula or set of formulas ϕ, $Var(\phi)$ denotes the set of variables from PS occurring in ϕ.

Definition 2.1. A *belief base* (BB) is a finite set $B = \{\phi_1, \ldots, \phi_n\}$ of formulas

[1] Other non-genuine connectives could be introduced as well, as syntactic sugars.

from $PROP_{PS}$. A *stratified belief base* (SBB) is a belief base B equipped with a total pre-order $\preceq$ over B. Every subset of B is a *subbase* of B.

Equivalently, B can be defined as a finite sequence $(\Delta_1, \ldots, \Delta_k)$ of subbases of B, where each Δ_i $(i \in 1 \ldots k)$ is the non-empty set that contains all the minimal elements of $B \setminus (\bigcup_{j=1}^{i-1} \Delta_j)$ [2] w.r.t. $\preceq$. Clearly enough, $\{\Delta_1, \ldots, \Delta_k\}$ is a partition of B. Each subset Δ_i $(i \in 1 \ldots k)$ is called a *stratum* of B, and i is the priority level of each formula of Δ_i. The lowest the priority level of a formula the highest its plausibility.

The standard coherence-based approach to inconsistency handling consists in selecting some consistent subbases $B_{\mathcal{P}}$ of B, and then performing classical inference from some of them. Accordingly, it is very close to the syntax-based revision framework proposed in [48] [49] as well as to the framework for supernormal default reasoning with priorities reported in [14]. Given a selection policy $\mathcal{P}$, several entailment principles can be considered. We focus on the most rational one [17], i.e., skeptical reasoning:

Definition 2.2. Let $B = (\Delta_1, \ldots, \Delta_k)$ be an SBB, $\mathcal{P}$ be a policy for the generation of preferred subbases, and γ be a formula from $PROP_{PS}$. γ is a *consequence* of B w.r.t. $\mathcal{P}$, noted $B \mathrel{|\!\!\sim}^{\mathcal{P}} \gamma$, iff $\forall B' \in B_{\mathcal{P}}$, $B' \models \gamma$.

Several policies $\mathcal{P}$ for the selection of consistent subbases can be defined [53] [5] [6]. Among them are:

Definition 2.3. Let $B = (\Delta_1, \ldots, \Delta_k)$ be an SBB.

- The set $B_{\mathcal{PO}}$ of all the preferred consistent subbases of B for the *possibilistic policy* is the singleton $\{\bigcup_{i=1}^{s-1} \Delta_i\}$, where s is the smallest index $(1 \leq s \leq k)$ s.t. $\bigcup_{i=1}^{s} \Delta_i$ is inconsistent if B is inconsistent and $s = k+1$ otherwise.
- The set $B_{\mathcal{LO}}$ of all the preferred consistent subbases of B for the *linear order policy* is the singleton $\{\bigcup_{i=1}^{k} \Delta'_i\}$, where Δ'_i $(i \in 1 \ldots k)$ is defined by $\Delta'_i = \Delta_i$ if $\Delta_i \cup \bigcup_{j=1}^{i-1} \Delta'_j$ is consistent, $\emptyset$ otherwise.
- The set $B_{\mathcal{IP}}$ of all the preferred consistent subbases of B for the *inclusion-preference policy* is $\{B' \subseteq B \mid B'$ is consistent and $\forall B'' \subseteq B$ s.t. B'' is consistent, $\forall i \in 1 \ldots k$ $((\forall j < i \;\; (B' \cap \Delta_j = B'' \cap \Delta_j)) \Rightarrow B' \cap \Delta_i \not\subset B'' \cap \Delta_i)\}$.
- The set $B_{\mathcal{LE}}$ of all the preferred consistent subbases of B for the *lexicographic policy* is $\{B' \subseteq B \mid B'$ is consistent and $\forall B'' \subseteq B$ s.t. B'' is consistent, $\forall i \in 1 \ldots k$ $((\forall j < i \;\; (|B' \cap \Delta_j| = |B'' \cap \Delta_j|)) \Rightarrow |B' \cap \Delta_i| \geq |B'' \cap \Delta_i|)\}$.

These definitions are equivalent to those presented in [5] (the possibilistic policy was called *best-out*) and [18] for the *linear order* policy. In[5], preferred

[2] By convention, $\bigcup_{j=1}^{0} \Delta_j = \emptyset$.

subbases are defined as the consistent subbases of B which are maximal w.r.t. some pre-orders. For the sake of brevity, we refrain from recalling such pre-orders here; the main point is that our definitions characterize the same preferred subbases, without referring explicitly to the pre-orders.

The corresponding inference relations have their own motivations; none of them is better than the other ones when both the computational aspects, the logical ones and cautiousness are considered simultaneously as evaluation criteria [17] [18].

Example 2.1. Let $B = (\{penguin\}, \{(\neg penguin \vee bird), (\neg feathers \vee bird), (\neg bird \vee fly), (feathers \wedge \neg fly)\}, \{eat\_fish, (\neg eat\_fish \vee piscivore)\})$. Using the policies previously defined, we have:

- $B_{\mathcal{PO}} = \{\{penguin\}\}$
- $B_{\mathcal{LO}} = \{\{penguin, eat\_fish, (\neg eat\_fish \vee piscivore)\}\}$
- $B_{\mathcal{IP}} = \{\{penguin, (\neg penguin \vee bird), (\neg feathers \vee bird), (\neg bird \vee fly), eat\_fish, (\neg eat\_fish \vee piscivore)\}, \{penguin, (\neg penguin \vee bird), (\neg feathers \vee bird), (feathers \wedge \neg fly), eat\_fish, (\neg eat\_fish \vee piscivore)\}, \{penguin, (\neg bird \vee fly), (feathers \wedge \neg fly), eat\_fish, (\neg eat\_fish \vee piscivore)\}\}$
- $B_{\mathcal{LE}} = \{\{penguin, (\neg penguin \vee bird), (\neg feathers \vee bird), (\neg bird \vee fly), eat\_fish, (\neg eat\_fish \vee piscivore)\}, \{penguin, (\neg penguin \vee bird), (\neg feathers \vee bird), (feathers \wedge \neg fly), eat\_fish, (\neg eat\_fish \vee piscivore)\}\}$

When skeptical reasoning is applied, we obtain:

- $\forall \mathcal{P} \in \{\mathcal{PO}, \mathcal{LO}, \mathcal{IP}, \mathcal{LE}\}$, $B \mid\!\sim^{\mathcal{P}} penguin$
- $\forall \mathcal{P} \in \{\mathcal{LO}, \mathcal{IP}, \mathcal{LE}\}$, $B \mid\!\sim^{\mathcal{P}} eat\_fish$ and $B \mid\!\sim^{\mathcal{P}} piscivore$
- $B \mid\!\sim^{\mathcal{LE}} bird$

2.2. $S-3$ inference

Our family of inference relations is based on an extension of $S-3$ logic [58], a multi-valued logic which can be viewed as a generalization of Levesque's logic of limited inference [44]. In $S-3$ logic, every propositional variable is interpreted in a weak way, allowing it to be both true and false, except for the variables of a given set S. The language of $S-3$ logic is the subset of $PROP_{PS}$ containing all the formulas in *Negation Normal Form* (NNF for short), where a formula $\phi \in PROP_{PS}$ is in NNF iff the sole subformulas of ϕ that are in the scope of any occurrence of $\neg$ are variables. It is well-known that any formula from $PROP_{PS}$ can be turned into a classically equivalent NNF formula in linear time using De Morgan's laws and involution of $\neg$. Among NNF formulas are CNF formulas, consisting of conjunctions of clauses, where a clause is a disjunction of literals.

Definition 2.4. [58] Let S be a subset of PS.

- Let $L_{PS} = PS \cup \{\neg x \mid x \in PS\}$ be the set of all literals built up from PS. An $S-3$ *interpretation* I over PS is a mapping from L_{PS} to $\{0,1\}$ s.t. for every variable x of PS, we never have $I(x) = I(\neg x) = 0$ and for every variable x of S, we never have $I(x) = I(\neg x)$.
- An $S-3$ *model* I of an NNF formula ϕ is an $S-3$ interpretation s.t. $I(\phi) = 1$ (it is assumed that $\vee$ and $\wedge$ are classically interpreted).
- Let B be a conjunctively-interpreted belief base from $PROP_{PS}$. Let γ be an NNF formula from $PROP_{PS}$. γ is an $S-3$ *consequence* of B, noted $B \models^3_S \gamma$, iff every $S-3$ model of B is an $S-3$ model of γ.

A formula is said to be $S-3$ *consistent* when it has an $S-3$ model. A 3-interpretation (resp. 3-model) is just an $\emptyset-3$ interpretation (resp. $\emptyset-3$ model) and a formula is 3-consistent when it has a 3-model. Since the Boolean constant $false$ is not taken into account in the language $PROP_{PS}$, every formula and every belief base is 3-consistent: the 3-interpretation I_{max} over PS s.t. $I(l) = 1$ for every $l \in L_{PS}$ is a 3-model of it.

Some standard inference rules of the proof theory of classical propositional logic do not hold in $S-3$ logic, or only in a restricted way. Specifically, this is the case of the *disjunctive syllogism* [58]:

Example 2.2. fly is a $\{bird\}-3$ consequence of $bird$ and $\neg bird \vee fly$. However, fly is *not* a $\{fly\}-3$ consequence of $bird$ and $\neg bird \vee fly$ since the $\{fly\}-3$ interpretation I s.t. $I(bird) = I(\neg bird) = 1$ and $I(fly) = 0$ is a $\{fly\}-3$ model of $bird$ and $\neg bird \vee fly$, but it does not satisfy fly.

Interestingly, the limited power of $S-3$ inference enables to handle *some* inconsistencies (by avoiding to make them reachable):

Example 2.3. *(2.1 continued)* Let $B = \{penguin, (\neg penguin \vee bird), (\neg feathers \vee bird), (\neg bird \vee fly), (feathers \wedge \neg fly), eat\_fish, (\neg eat\_fish \vee piscivore)\}$. B is classically inconsistent.

- Let $S = \{feathers, fly, eat\_fish\}$. We have $B \models^3_S penguin \wedge feathers \wedge \neg fly \wedge bird \wedge \neg bird \wedge eat\_fish \wedge piscivore$ but $B \not\models^3_S fly$.
- Let $S = \{bird, feathers, fly, eat\_fish\}$. For every (NNF) formula γ from $PROP_{PS}$, we have $B \models^3_S \gamma$. By definition, $B \models^3_S$ is explosive: the image of B by $\models^3_S$ is the whole language $PROP_{PS}$.

For every $S \subseteq PS$, the corresponding inference relation $\models^3_S$ is an approximation of $\models$ by below: if $B \models^3_S \gamma$ then $B \models \gamma$. This is easy to envision in the case where B is classically inconsistent since in this situation, $B \models \gamma$ holds for every formula γ.

Tractability of $S-3$ entailment is ensured by limiting the size of S and the format of the queries. To be more specific, the time complexity of $S-3$ inference is in $\mathcal{O}(2^{|S|} \cdot |B| \cdot |\gamma|)$ when B is a CNF formula and γ is a clause [58]. Accordingly, the complexity of $S-3$ inference depends essentially on the number of variables in S. Combining both Theorem A.3 from [58], Lemma 3.1 and Theorem 4.1 from [15], this result can be easily generalized to the case where B is any NNF formula and γ is a CNF formula; in this situation, $S-3$ inference can still be decided in time linear in the size of the belief base.

3. Resource-bounded paraconsistent inference

3.1. Limitations of $S-3$ inference

Adhering to $S-3$ entailment is a way to prevent some inconsistencies from being harmful. Indeed, as the previous example illustrates, performing $S-3$ inference from an inconsistent belief base does not necessarily result in trivialization. Several inconsistencies are not detected because classical reasoning over some variables outside S would be required to point them out. Especially, if S does not contain all the variables belonging to a (minimally) inconsistent subbase of B, then the corresponding inference relation $B \models^3_S$ is not explosive. In the previous example with $S = \{feathers, fly, eat\_fish\}$, while we have both $B \models^3_S bird$ and $B \models^3_S \neg bird$, we do **not** have $B \models^3_S fly$. The fact that $bird$ is both true and false is not problematic if $bird$ is an irrelevant query, and $bird$ does not participate to the inference of any relevant query. Actually, common-sense reasoning is often like that: only a subset of the current beliefs is typically exploited when a query is to be answered [36] [19] [3]; efficiency is ensured this way and many inconsistencies are handled as well; specifically, all the inconsistencies related to the discarded beliefs are ignored.

Nevertheless, $S-3$ logic in its current turn suffers from some limitations. First of all, as the previous example shows, (1) restricting S to an arbitrary subset of the variables occurring in B is not sufficient to avoid trivialization in every situation. In addition, (2) $S-3$ inference is often too cautious since the disjunctive syllogism does not hold in the general case. Furthermore, (3) the restriction to NNF formulas is unnecessarily restrictive. Finally, (4) the language upon which this logic is based is not expressive enough to make the distinction between the necessary consequences of a belief base and the plausible ones, which means that integrity constraints cannot be handled in a satisfying way. Indeed, suppose that

[3] This is also the case in mathematics. Thus, the existence of paradoxes (inconsistencies) (like Russel's paradox - let $e = \{x \mid x \notin x\}$; is $e \in e$ true?) did not prevent mathematicians from pointing out many interesting results over the past centuries, specifically because these paradoxes were mainly irrelevant to such results.

our belief base B is $\{fly, \neg fly, (fly \vee penguin)\}$. We can reasonably conclude that fly is only possible since there are pros and cons for it. Since $(fly \vee penguin)$ is considered possible as well (but not certain), we cannot conclude anything about $penguin$. However, in the case where $(fly \vee penguin)$ is an integrity constraint (i.e., it *must* be true), $penguin$ is expected to be true. Using $S-3$ inference, this conclusion cannot be drawn: because $fly \notin S$ is required to avoid trivialization, no choice of S enables $penguin$ to be derived from B. As shown by this example, integrity constraints cannot be expressed in $S-3$ logic through a relevant choice of S since a uniform treatment of variables of S is imposed: if $x \in S$, then x is classically interpreted everywhere, not only locally to some formulas of B.

3.2. A family of resource-bounded paraconsistent relations

In order to circumvent the four limitations (1) (2) (3) and (4) of $S-3$ inference discussed above, we introduce a new family of resource-bounded paraconsistent inference relations. This family is based on the propositional language $PROP^3_{PS}$ of paraconsistent logic J_3 [26]. J_3 is a 3-valued logic whose "truth values" are $0, \frac{1}{2}, 1$. While 0 (resp. 1) stands for false (resp. true) (as usual), the "truth value" $\frac{1}{2}$ stands for *both true and false*.

Formally, $PROP^3_{PS}$ is inductively defined in the usual way from a finite set PS of propositional symbols, the Boolean constants $true$ and $false$ and the connectives $\neg$, $\wedge$, $\vee$ and $\Box$. Accordingly, $PROP_{PS}$ is a proper subset of $PROP^3_{PS}$.

The semantics of formulas in J_3 is given by the following truth tables:

A	$\neg A$
1	0
$\frac{1}{2}$	$\frac{1}{2}$
0	1

$\neg$

$A \backslash B$	1	$\frac{1}{2}$	0
1	1	1	1
$\frac{1}{2}$	1	$\frac{1}{2}$	$\frac{1}{2}$
0	1	$\frac{1}{2}$	0

$\vee$

$A \backslash B$	1	$\frac{1}{2}$	0
1	1	$\frac{1}{2}$	0
$\frac{1}{2}$	$\frac{1}{2}$	$\frac{1}{2}$	0
0	0	0	0

$\wedge$

A	$\Box A$
1	1
$\frac{1}{2}$	0
0	0

$\Box$

Clearly enough, $\wedge$ (resp. $\vee$) stands for a conjunction (resp. disjunction) connective, while $\neg$ is a negation. $\Box$ is a necessity operator since $\Box\phi$ is true if ϕ is true and $\Box\phi$ is false otherwise (in particular, when ϕ is both true and false). Let us say that a *square literal* is a formula of the form $\Box x$ or $\Box\neg x$ where $x \in PS$. A *signed square literal* is a square literal $\Box l$ or its negation $\neg\Box l$ and an *unsigned square literal* is just a square literal.

Interestingly, some other connectives can be derived in J_3, for instance:

- possibility: $\diamond\phi =_{def} \neg\Box\neg\phi$.
- material implication: $\phi \supset \psi =_{def} \Box(\neg\phi) \vee \psi$.

We now introduce the notions of belief base, $S-3$ interpretation, $S-3$ model, $S-3$ consistency, $S-3$ consequence related to the language of J_3; we also introduce the corresponding notion of $S-3$ equivalence. As in $S-3$ logic, whenever $S=\emptyset$, we simply refer to the notions of 3-interpretation, 3-model, 3-consistency, 3-consequence, 3-equivalence respectively.

Definition 3.1. Let S be a subset of PS.

- A *belief base* B (from $PROP^3_{PS}$) is a finite and conjunctively-interpreted set of formulas from $PROP^3_{PS}$.
- An $S-3$ *interpretation* I over PS is a mapping from PS to $\{0,1,\frac{1}{2}\}$ s.t. $\forall x \in S$, $I(x) \neq \frac{1}{2}$.
- The semantics $I(\phi)$ of a formula $\phi \in PROP^3_{PS}$ in I is inductively defined by:
  - $I(true)=1$;
  - $I(false)=0$;
  - the truth tables presented above.
- An $S-3$ interpretation I is said to be an $S-3$ *model* of ϕ iff $I(\phi)>0$.
- A formula ϕ from $PROP^3_{PS}$ is said to be $S-3$ *consistent* iff there exists an $S-3$ model of ϕ.
- Let γ be a formula from $PROP^3_{PS}$. γ is an $S-3$ *consequence* of a belief base B from $PROP^3_{PS}$, noted $B \models^3_S \gamma$, iff every $S-3$ model of B is a $(S-3)$ model of γ.
- Let ϕ and ψ be formulas from $PROP^3_{PS}$. ϕ and ψ are said to be $S-3$ *equivalent*, noted $\phi \equiv^3_S \psi$, iff both $\phi \models^3_S \psi$ and $\psi \models^3_S \phi$ hold.

Example 3.1. *(2.1 continued)* Let $B=\{penguin, (\neg penguin \vee bird), (\neg feathers \vee bird), (\neg bird \vee fly), (feathers \wedge \neg fly), eat\_fish, (\neg eat\_fish \vee piscivore)\}$. B is classically inconsistent.

- Let $S=\{penguin, feathers, fly, eat\_fish\}$. There are two $S-3$ models of B, the $S-3$ interpretations I_1 and I_2 such that:
  - $I_1(fly)=I_2(fly)=0$
  - $I_1(penguin)=I_2(penguin)=1$
  - $I_1(feathers)=I_2(feathers)=1$
  - $I_1(eat\_fish)=I_2(eat\_fish)=1$

  As these variables have to be classically interpreted, we also have:
  - $I_1(\neg fly)=I_2(\neg fly)=1$
  - $I_1(\neg penguin)=I_2(\neg penguin)=0$
  - $I_1(\neg feathers)=I_2(\neg feathers)=0$

* $I_1(\neg eat\_fish) = I_2(\neg eat\_fish) = 0$

As we have $I_1(\neg penguin) = I_2(\neg penguin) = 0$ and $I_1(fly) = I_2(fly) = 0$, we must have:

* $I_1(bird) = I_1(\neg bird) = I_2(bird) = I_2(\neg bird) = \frac{1}{2}$.

As we have $I_1(\neg eat\_fish) = I_2(\neg eat\_fish) = 0$, we must have:

* $I_1(piscivore) = \frac{1}{2}$ and $I_2(piscivore) = 1$.

Given these two $S-3$ models of B, we have $B \models^3_S penguin \wedge feathers \wedge \neg fly \wedge bird \wedge \neg bird \wedge eat\_fish \wedge piscivore$ but $B \not\models^3_S fly$.

- Let $S = \{bird, feathers, fly, eat\_fish\}$. We have $B \models^3_S false$. Equivalently, for every formula γ from $PROP^3_{PS}$, we have $B \models^3_S \gamma$ (or stated in a third way, B is not $S-3$ consistent).

These definitions show that our framework includes J_3 logic as a specific case, which is recovered by setting S to the empty set. Clearly enough, we slightly abuse words here since we use the terminology and notations already at work for $S-3$ logic (Section 2.2). However, this is harmless because the notions above are extensions to the language $PROP^3_{PS}$ of the corresponding notions for $PROP_{PS}$, i.e., they pairwise coincide whenever NNF formulas from $PROP_{PS}$ are considered, only. This is stated formally by the next proposition:

Proposition 3.1. Let S be any subset of PS. To every mapping I from L_{PS} to $\{0,1\}$ s.t. $\forall x \in PS$, we never have $I(x) = I(\neg x) = 0$ and $\forall x \in S$, we never have $I(x) = I(\neg x)$, we can associate in a bijective way the mapping $\hat{I}$ from PS to $\{0, \frac{1}{2}, 1\}$ s.t. for every $x \in S$, we never have $\hat{I}(x) = \frac{1}{2}$, where $\hat{I}$ is defined by:

- $\hat{I}(x) = 1$ iff $I(x) = 1$ and $I(\neg x) = 0$,
- $\hat{I}(x) = 0$ iff $I(x) = 0$ and $I(\neg x) = 1$, and
- $\hat{I}(x) = \frac{1}{2}$ iff $I(x) = 1$ and $I(\neg x) = 1$.

In addition, for every NNF formula ϕ from $PROP_{PS}$, we have $I(\phi) = 1$ iff $\hat{I}(\phi) > 0$.

Specifically, this proposition shows that $S-3$ logic can be recovered as a specific case in our framework (just restrict the language to the NNF formulas of $PROP_{PS}$). Interestingly, unlike $S-3$ logic, J_3 is fully truth functional; negation does not play a special role, so that there is no need to restrict formulas to NNF ones here. Furthermore, considering the $\Box$ connective actually leads to enrich the language, in the sense that it cannot be encoded using $\neg$, $\wedge$, $\vee$, only (see [30]). This enables new inferences to be drawn[4] (e.g., *modus ponens* - from $bird$ and $bird \supset fly$, conclude fly - holds, as well as disjunctive syllogism whenever a square literal is resolved upon; thus, from $\neg bird$ and $(\Box bird \vee fly)$, it is possible to conclude fly). The necessity operator $\Box$ also allows us to make a distinction between credible pieces of belief (hypotheses) and certain ones (pieces of knowledge), including integrity constraints. If we step back to a previous example, the integrity constraint $(fly \vee penguin)$ can be expressed by the formula $\Box(fly \vee penguin)$ in $PROP^3_{PS}$. The whole example can be handled in an adequate way by considering the belief base $B = \{fly, \neg fly, \Box(fly \vee penguin)\}$ from which we can conclude $penguin$. Accordingly, while $S-3$ logic is unable to handle integrity constraints, they can be taken into account in our framework in a satisfying way.

Hitherto, the three limitations (2) (3) and (4) of $S-3$ logic over the four reported in Section 3.1 have been circumvented by extending the framework to J_3. The remaining one (1) lies in the fact that trivialization is not avoided in every situation. Actually, we will not ensure it for every belief base from $PROP^3_{PS}$ but consider in what follows 3-*consistent belief bases, only.* This amounts to assume that there is no contradiction between the pieces of knowledge (certain beliefs) and the remaining ones within a belief base. Especially, if belief bases generated from the fragment $PROP_{PS}$ are considered (as in $S-3$ logic), then this assumption is always satisfied (as shown before).

In order to deal with the inconsistencies that may be revealed when using $S-3$ inference in $PROP^3_{PS}$, we suggest to focus on some subsets S' of S, those for which the corresponding inference relations are not explosive. The approach is similar to the standard coherence-based approach to inconsistency handling, except that the inference relation is weakened by removing variables from S instead of removing explicit beliefs from B. Formally, we first need a notion playing a role similar to the notion of consistent subbase in the standard coherence-based approach.

Definition 3.2. Let B be a belief base from $PROP^3_{PS}$. Let $S \subseteq PS$ and $S_0 \subseteq S$ s.t. $B \not\models^3_{S_0} false$. A *consistent subset* S' of S w.r.t. B and S_0 is a subset of S containing S_0 and s.t. $B \not\models^3_{S'} false$. $\mathcal{S}(B, S, S_0)$ denotes the set of all consistent subsets of S w.r.t. B and S_0.

In this definition, S_0 is a given set of variables which *must* be interpreted classically. The condition $B \not\models^3_{S_0} false$ precisely means that $B \models^3_{S_0}$ is not *explosive*, i.e., the set of its consequences is not the whole language $PROP^3_{PS}$. Formally, we could avoid considering such a set S_0 by replacing in B every oc-

[4] As we will see in the following, the price to be paid is an increase in computational complexity compared with $S-3$ logic in some situations, especially when S is empty and clausal queries are considered.

currence of x by $\Box x$ and every occurrence of $\neg x$ by $\Box\neg x$, for each $x \in S_0$. We keep it nevertheless for an ease of presentation. Note that the existence of such a set S_0 (possibly empty) is ensured provided that B is 3-consistent.

Example 3.2. *(2.1 continued)* Let $B = \{penguin, (\neg penguin \vee bird), (\neg feathers \vee bird), (\neg bird \vee fly), (feathers \wedge \neg fly), eat\_fish, (\neg eat\_fish \vee piscivore)\}$.
Let $S = \{penguin, bird, feathers, fly, eat\_fish, piscivore\}$ and $S_0 = \{penguin, feathers, eat\_fish\}$.
We have $\mathcal{S}(B, S, S_0) = \{\{penguin, feathers, eat\_fish\}, \{penguin, feathers, eat\_fish, bird\}, \{penguin, feathers, eat\_fish, fly\}, \{penguin, feathers, eat\_fish, piscivore\}, \{penguin, feathers, eat\_fish, bird, piscivore\}, \{penguin, feathers, eat\_fish, fly, piscivore\}\}$.

Among the consistent subsets, we are interested in those from which the maximum amount of information is kept but trivialization is avoided. This calls for a selection policy $\mathcal{P}$ which aims at pointing out a preferred subset of $\mathcal{S}(B, S, S_0)$.

In the following, we adhere to a skeptical approach: any piece of belief is considered as a consequence of B given S and S_0 iff it is $S' - 3$ entailed by B for every preferred consistent subset S' of S w.r.t. B and S_0.

Definition 3.3. Let B be a belief base from $PROP^3_{PS}$ and S_0 be a subset of PS s.t. $B \not\models^3_{S_0} false$. Let $\mathcal{P}$ be a selection policy s.t. $\mathcal{S}_{\mathcal{P}}(B, S, S_0)$ is a subset of $\mathcal{S}(B, S, S_0)$ and let γ be a formula from $PROP_{PS}$. γ is a *consequence* of B w.r.t. $\mathcal{P}$ and S_0, noted $B \approx^{\mathcal{P},S_0}_{S} \gamma$, iff $\forall S' \in \mathcal{S}_{\mathcal{P}}(B, S, S_0)$, $B \models^3_{S'} \gamma$.

As in Belnap's four valued logic [4], four different epistemic status for every formula γ w.r.t. a belief base are reached: *accepted, rejected, over-determined, undetermined*:

- γ is *over-determined* if $B \approx^{\mathcal{P},S_0}_{S} \gamma$ and $B \approx^{\mathcal{P},S_0}_{S} \neg\gamma$.
- γ is *accepted* if $B \approx^{\mathcal{P},S_0}_{S} \gamma$ and $B \not\approx^{\mathcal{P},S_0}_{S} \neg\gamma$.
- γ is *rejected* if $B \not\approx^{\mathcal{P},S_0}_{S} \gamma$ and $B \approx^{\mathcal{P},S_0}_{S} \neg\gamma$.
- γ is *undetermined* if $B \not\approx^{\mathcal{P},S_0}_{S} \gamma$ and $B \not\approx^{\mathcal{P},S_0}_{S} \neg\gamma$.

Such a semantics shows that our approach has a flavor of argumentative logic [29]. Indeed, a formula is *accepted* if it is derivable in every preferred, computationally accessible, context[5] but this is not the case for its negation (which can be derivable in some of them nevertheless); roughly speaking, there is

[5] Clearly enough, it would be easy to avoid the over-determined epistemic status by slightly modifying the notion of consequence and requiring that $B \approx^{\mathcal{P},S_0}_{S} \gamma$ and $B \not\approx^{\mathcal{P},S_0}_{S} \neg\gamma$ in order to consider γ as a consequence of B w.r.t. $\mathcal{P}$ and S_0. However, this trick would lead to make no distinction between indetermination and over-determination, which is not always desirable.

always (i.e., in every context) an argument to believe in the truth of an accepted formula but not in the truth of its negation.

It is worth noting that queries are restricted to formulas from $PROP_{PS}$. While enabling queries containing some occurrences of the $\Box$ connective would not raise many difficulties, it would have the undesirable effect of "amplifying the truth" of some consequences, leading to confuse some certain pieces of beliefs with credible ones. For instance, we would have $B = \{fly\} \approx_{PS}^{\mathcal{P},S_0} \Box fly$ (for every $\mathcal{P} \in \{\mathcal{PO}, \mathcal{LO}, \mathcal{IP}, \mathcal{LE}\}$ and every S_0); would such a query be admissible? This is not an expected conclusion when fly is not certain – and this is the case here. Especially, it can be the case that the new piece of evidence $\neg fly$ is added to B without questioning the 3-consistency of it.

Now, we can define selection policies for consistent subsets that are similar to the ones defined for consistent subbases in the standard coherence-based approach [53] [5] [6]. Slightly abusing words, we use the same names here since no ambiguity is possible.

Definition 3.4. Let B be a belief base from $PROP^3_{PS}$. Let $S = (S_0, \ldots, S_j)$ be a stratification, i.e., a totally ordered partition of a given subset of PS where S_0 is s.t. $B \not\models^3_{S_0} false$.

- The set $\mathcal{S}_{\mathcal{PO}}(B, S, S_0)$ of all preferred consistent subsets of S w.r.t. B and S_0 for the *possibilistic policy* is the singleton $\{\bigcup_{i=0}^{s-1} S_i\}$, where s is either the smallest index $(1 \leq s \leq j)$ s.t. $\bigcup_{i=0}^{s} S_i$ is an inconsistent subset of S w.r.t. B and S_0 or $j+1$ if S is consistent w.r.t. B.
- The set $\mathcal{S}_{\mathcal{LO}}(B, S, S_0)$ of all preferred consistent subsets of S w.r.t. B and S_0 for the *linear order policy* is the singleton $\{\bigcup_{i=0}^{j} S'_i\}$, where S'_i $(i \in 0 \ldots j)$ is defined by $S'_i = S_i$ if $S_i \cup \bigcup_{l=0}^{i-1} S'_l$ is a consistent subset of S w.r.t. B and S_0, $\emptyset$ otherwise.
- The set $\mathcal{S}_{\mathcal{IP}}(B, S, S_0)$ of all preferred consistent subsets of S w.r.t. B and S_0 for the *inclusion-preference policy* is $\{S' \mid S'$ is a consistent subset of S w.r.t. B and S_0 and $\forall S'' \neq S'$ s.t. $S'' \in \mathcal{S}(B, S, S_0)$, $\forall i \in 0 \ldots j$ $((\forall l < i(S' \cap S_l = S'' \cap S_l)) \Rightarrow S' \cap S_i \not\subset S'' \cap S_i)\}$.
- The set $\mathcal{S}_{\mathcal{LE}}(B, S, S_0)$ of all preferred consistent subsets of S w.r.t. B and S_0 for the *lexicographic policy* is $\{S' \mid S'$ is a consistent subset of S w.r.t. B and S_0 and $\forall S'' \neq S'$ s.t. $S'' \in \mathcal{S}(B, S, S_0)$, $\forall i \in 0 \ldots j$ $((\forall l < i(|S' \cap S_l| = |S'' \cap S_l|)) \Rightarrow |S' \cap S_i| \geq |S'' \cap S_i|)\}$.

These definitions assume that a stratification of the given set S of variables is available. Such a stratification is the formal counterpart of the preferential information used to discriminate among consistent subsets in the standard

We prefer to keep separate the two notions and let the knowledge engineer free to handle undetermined and over-determined formulas in the same way or in different ways.

coherence-based approach. It enables to achieve some graded forms of paraconsistent reasoning (which is not so frequent, see [10]) by giving more credit to some variables than to others; as we will see in the following, we can easily take advantage of such a stratification of variables to encode a stratification of beliefs (see Proposition 3.10).

Here is an example showing the variety of conclusions that can be drawn in our framework, depending on the policy $\mathcal{P}$ and the stratification of S:

Example 3.3. *(2.1 continued)* Let $B = \{penguin, (\neg penguin \vee bird), (\neg feathers \vee bird), (\neg bird \vee fly), (feathers \wedge \neg fly), eat\_fish, (\neg eat\_fish \vee piscivore)\}$. With $S = (\emptyset, \{penguin, bird, feathers, fly, eat\_fish\})$, we have:

- $\mathcal{S}_{\mathcal{PO}}(B, S, S_0) = \{S_0\} = \{\emptyset\}$
- $\mathcal{S}_{\mathcal{LO}}(B, S, S_0) = \{S_0\} = \{\emptyset\}$
- $\mathcal{S}_{\mathcal{IP}}(B, S, S_0) = \{\{penguin, bird, feathers, eat\_fish\}, \{penguin, feathers, fly, eat\_fish\}, \{bird, fly, eat\_fish\}\}$
- $\mathcal{S}_{\mathcal{LE}}(B, S, S_0) = \{\{penguin, bird, feathers, eat\_fish\}, \{penguin, feathers, fly, eat\_fish\}\}$

With $S' = (\emptyset, \{penguin\}, \{bird, feathers, fly\}, \{eat\_fish\})$, we have:

- $\mathcal{S}_{\mathcal{PO}}(B, S', S'_0) = \{\{penguin\}\}$
- $\mathcal{S}_{\mathcal{LO}}(B, S', S'_0) = \{\{penguin, eat\_fish\}\}$
- $\mathcal{S}_{\mathcal{IP}}(B, S', S'_0) = \{\{penguin, bird, feathers, eat\_fish\}, \{penguin, feathers, fly, eat\_fish\}\}$
- $\mathcal{S}_{\mathcal{LE}}(B, S', S'_0) = \{\{penguin, bird, feathers, eat\_fish\}, \{penguin, feathers, fly, eat\_fish\}\}$

With S or S', $penguin$, $feathers$, $\neg fly$ and eat_fish are consequences of B (but their negations are not) whatever $\mathcal{P}$ among the four policies considered above. Moreover:

- With S:
  * w.r.t. $\mathcal{PO}$ (or $\mathcal{LO}$), none of $bird$, $\neg bird$, $piscivore$ and $\neg piscivore$ are consequences of B.
  * w.r.t. $\mathcal{IP}$, $piscivore$ is a consequence of B, but $\neg piscivore$, $bird$ and $\neg bird$ are not.
  * w.r.t. $\mathcal{LE}$, $bird$ and $piscivore$ are consequences of B, but their negations are not.
- With S':
  * w.r.t. $\mathcal{PO}$, $bird$ is a consequence of B, but $\neg bird$, $piscivore$ and $\neg piscivore$ are not.

* w.r.t. $\mathcal{LO}$, $\mathcal{IP}$ or $\mathcal{LE}$, *bird* and *piscivore* are consequences of B, but their negations are not.

Clearly enough, in the case where $B \not\models^3_S false$, all our inference relations $B \mathrel{|\!\!\approx}^{\mathcal{P},S_0}_S$ coincide with $B \models^3_S$. Accordingly, both J_3-inference and $S-3$ inference can be recovered as specific cases in our framework in this situation (just set S to $\emptyset$ in the first case, and consider formulas from $PROP_{PS}$ in the second case).

The next proposition states that one of our basic objectives (avoiding trivialization) is always reached.

Proposition 3.2. Let B be a belief base from $PROP^3_{PS}$. Let $S = (S_0, \ldots, S_j)$ be a stratification of a given subset of PS where S_0 is s.t. $B \not\models^3_{S_0} false$. None of the relations $B \mathrel{|\!\!\approx}^{\mathcal{P},S_0}_S$ with $\mathcal{P}$ among $\mathcal{PO}$, $\mathcal{LO}$, $\mathcal{IP}$, $\mathcal{LE}$ is explosive.

The inclusion-preference policy actually focuses on the consistent subsets from which the maximum amount of information is preserved while trivialization is avoided: $\mathcal{S}_{\mathcal{IP}}(B, S, S_0)$ contains exactly the consistent subsets S' of S that are minimal (i.e. preferred) w.r.t. the pre-ordering defined by $S' \leq_{\mathcal{IN}} S''$ iff $\forall \gamma \in PROP_{PS}$, if $B \models^3_{S''} \gamma$ holds, then $B \models^3_{S'} \gamma$ holds as well. Noting $B \mathrel{|\!\!\approx}^{\mathcal{IN},S_0}_S \gamma$ when $B \models^3_{S'} \gamma$ for each consistent subset S' of S that is minimal w.r.t. $\leq_{\mathcal{IN}}$, we have:

Proposition 3.3. Let B be a belief base from $PROP^3_{PS}$. Let $S = (S_0, \ldots, S_j)$ be a stratification of a given subset of PS where S_0 is s.t. $B \not\models^3_{S_0} false$. For every formula γ of $PROP_{PS}$, we have $B \mathrel{|\!\!\approx}^{\mathcal{IN},S_0}_S \gamma$ iff $B \mathrel{|\!\!\approx}^{\mathcal{IP},S_0}_S \gamma$.

A similar conclusion can be reached in the standard coherence-based approach (where the $\mathcal{IN}$ policy leads to the same skeptical inference relation as the $\mathcal{IP}$ policy).

3.3. Instantiating our framework

Up to now, we have assumed that both a stratified set S and a set S_0 are available. In the following, we suggest several ways to define them from more basic data.

Query-directedness of inference.

A first question to be considered concerns the query-directedness of inference. Indeed, is it a good point for an inference relation to be dependent on the belief base alone, or may it also vary with the queries? The traditional view is that inference is fully characterized from the agent beliefs and is query independent. This assumption results in inference relations with a much more predictable behaviour than if the other choice were made. However, the traditional view also ignores the computational dimension of inference; while it is perfectly reasonable not to take queries into account when every piece of belief is exploited, this is not so reasonable when computational limits prevent from a full exploitation of beliefs. In this situation, queries can be used to characterize (at least partly) the pieces of beliefs we should focus on, and the computational effort we are ready to spend. Tuning the computational resources depending on the importance of getting the right answers seems to be an important feature of common-sense reasoning. Several approaches to common-sense inference (e.g., [36], [19]) adhere to this point of view.

Our purpose is not to address such an issue in this paper, but to note that our framework is general enough to enable both points of view. Specifically, a very simple way to take queries γ into account consists in systematically adding the variables of γ into S (as long as their number remains small enough, which is a reasonable assumption). Of course, much more sophisticated policies can be defined, see [61] [32]. From now on, we will only consider belief-driven inference relations and let their extensions to query-directed inference for further research.

A policy to define a stratified S from an SBB.

There are many ways to define a stratified set S from a stratified belief base (SBB) B from $PROP_{PS}$[6]. First, S_0 can be defined as the set of all variables from the certain beliefs (i.e., pieces of knowledge) of B when available, provided that such a set is small enough and satisfies $B \not\models^3_{S_0} false$. Otherwise, S_0 is defined as $\emptyset$.

Defining S is known as a difficult problem in the general case [58]. Interestingly, when B is stratified, it is possible to exploit the given preferential information so as to define S. A very natural way consists in considering in priority the variables of the most plausible beliefs until a preset bound b on the size of S is reached. Formally, we consider the following set of variables S^b:

Definition 3.5. Let $B = (\Delta_1, \ldots, \Delta_k)$ be an SBB from $PROP_{PS}$. Let b be a positive integer. Let S_0 be a subset of PS s.t. $B \not\models^3_{S_0} false$ and $|S_0| \leq b$. The *inference level* of B given b and S_0, noted i^b_B is either k if $|Var(B)| \leq b$ or the lowest i ($i \in 0 \ldots k-1$) s.t. $S^b_i = S^b_k$, where the sequence (S^b_i) is inductively defined by $S^b_0 = S_0$ and $(\forall i \in 1 \ldots k)$ $S^b_i = S^b_{i-1} \cup Var(\Delta_i)$ if $|S^b_{i-1} \cup Var(\Delta_i)| \leq b$, $S^b_i = S^b_{i-1}$ otherwise. We define S^b as $S^b_{i^b_B}$.

Example 3.4. Let $B = (\Delta_1, \Delta_2, \Delta_3, \Delta_4)$ with $\Delta_1 = \{penguin, (\neg penguin \vee bird\}$, $\Delta_2 = \{(\neg feathers \vee bird), (feathers \wedge \neg fly)\}$, $\Delta_3 = \{(\neg bird \vee fly), eat\_fish\}$ and $\Delta_4 = \{(\neg eat\_fish \vee piscivore\}$.

[6] Upstream to our work are several techniques that aim at deriving an SBB from a "flat" belief base, see [7] for a survey.

- Let $b = 3$. We have $S^b = \{penguin, bird\}$.
- Let $b = 4$. We have $S^b = \{penguin, bird, feathers, fly\}$.

At a first glance, the computation of the inference level of a belief base looks like the computation of the inconsistency level for possibilistic belief bases. However, the former is guided by computational motivations, only; especially, it is not the case that $B \not\models^3_{S^b} false$ in the general case.

Finally, once both S_0 and a flat set S are available, a stratification of S can be easily derived from the given stratification of B. Basically, once S_0 has been taken into account, it is reasonable to give more importance to the variables of S that belong to the more plausible beliefs. Loosely speaking, we want to reason as much classically as possible from the most plausible beliefs in order to take into account their, possibly conflicting, classical consequences: inconsistencies must not be ignored when supported by plausible beliefs.

Definition 3.6. Let $B = (\Delta_1, \ldots, \Delta_k)$ be an SBB from $PROP_{PS}$, S a subset of PS and S_0 a subset of S. The stratification of S induced by B given S_0 is $\mathcal{S} = (S_0, \ldots, S_{k+1})$ s.t. $S_i = (S \cap Var(\Delta_i)) \setminus \bigcup_{l=0}^{i-1} S_l$ for every $i \in 1 \ldots k$, and $S_{k+1} = S \setminus \bigcup_{l=0}^{k} S_l$.

Example 3.5. Let $B = (\Delta_1, \Delta_2, \Delta_3, \Delta_4)$ with $\Delta_1 = \{penguin, (\neg penguin \vee bird\}$, $\Delta_2 = \{(\neg feathers \vee bird), (feathers \wedge \neg fly)\}$, $\Delta_3 = \{(\neg bird \vee fly), eat\_fish\}$ and $\Delta_4 = \{(\neg eat\_fish \vee piscivore\}$. Let $S_0 = \{penguin\}$ and $S = \{penguin, bird, fly, feathers, eat\_fish, piscivore, swim\}$. The stratification of S induced by B given S_0 is $\mathcal{S} = (S_0, S_1, S_2, S_3, S_4, S_5)$ with $S_0 = \{penguin\}$, $S_1 = \{bird\}$, $S_2 = \{feathers, fly\}$, $S_3 = \{eat\_fish\}$, $S_4 = \{piscivore\}$, $S_5 = \{swim\}$.

There are many other ways to derive a stratification of S. For instance, in the situation where we are ready to give more credit to facts than to rules, we can take advantage of the notion of degree of definiteness of a variable (as defined in [12]) to achieve this goal from a CNF belief base B. Roughly, every variable x of S occurring in B can be associated with the size of one of the smallest clauses of B containing x; a stratification of S is obtained by considering x before y whenever the score attached to x is lower than the score attached to y.

3.4. Semantics, logical properties and cautiousness

Designing a tractable paraconsistent inference relation is really simple. For instance, restricting $\models$ so as to consider as consequences of a belief base B only valid formulas containing, say, ten literals s.t. exactly five of them occurs in B, achieves this goal. Is such an inference relation for all that valuable? Obviously, the answer is no because it is very difficult to envision the effects of any change on

B on the set of its consequences and because only few expected consequences of B are reached this way. Indeed, on the one hand, it is important to give a precise semantics to the representation language on which the inference relation is built and to identify the main inferential properties of this relation. Without them, the relation is necessarily of a limited practical use since a knowledge engineer must have a good understanding of the meaning and the expected consequences of the pieces of beliefs she deals with. On the other hand, over-cautious inference relations must be avoided.

Semantics and logical properties.

Our framework offers a wide spectrum of inference relations, parametrized by the set S of variables that are preferred interpreted classically. As in the standard coherence-based approach, it is easy to give a preferential models semantics to each of our relations $\approx_S^{\mathcal{P},S_0}$.

Proposition 3.4. Let B be a belief base from $PROP_{PS}^3$. Let $S = (S_0, \ldots, S_j)$ be a stratification of a given subset of PS where S_0 is s.t. $B \not\models_{S_0}^3 false$. Let γ be a formula from $PROP_{PS}$. For every selection policy $\mathcal{P}$ among $\mathcal{PO}$, $\mathcal{LO}$, $\mathcal{IP}$ and $\mathcal{LE}$, we have $B \approx_S^{\mathcal{P},S_0} \gamma$ iff every preferred 3-model of B is a 3-model of γ, where the preferred 3-models of B are exactly the $S' - 3$-models of B with $S' \in \mathcal{S}_{\mathcal{P}}(B, S, S_0)$.

Thus, when $S = (\emptyset, PS)$, the preferred 3-models of B w.r.t. $\mathcal{IP}$ are the 3-models of it encoding worlds that are "as normal as possible" (i.e., as close as possible to classical interpretations w.r.t. $\subseteq$).

Another strong point of our approach is that our inference relations exhibit a limited form of syntax-sensitivity. Indeed, the binary connectives $\wedge$ and $\vee$ of our language behave in a classical manner: the set of $S - 3$ models of $\phi \wedge \psi$ (resp. $\phi \vee \psi$) is the intersection (resp. union) of the set of $S - 3$ models of ϕ and the set of $S - 3$ models of ψ. Furthermore, given a stratification $S = (S_0, \ldots, S_j)$, every formula from B (and B itself!) can be put into *normal forms* without modifying the corresponding inference relations $\approx_S^{\mathcal{P},S_0}$.

Definition 3.7. Let ϕ be a formula from $PROP_{PS}^3$.

- ϕ is said to be in $\Box NF$ iff only literals are in the scope of any occurrence of $\Box$ in ϕ.
- ϕ is said to be in $\Box NNF$ iff among the subformulas of ϕ, only literals are in the scope of any occurrence of $\Box$ and only positive literals and square literals are in the scope of any occurrence of $\neg$.
- ϕ is said to be in $\Box CNF$ iff it is a conjunction of disjunctions of literals or (signed or unsigned) square literals.

Interestingly, every formula from $PROP^3_{PS}$ can be turned efficiently into a *strongly equivalent* $\Box NNF$ one, where two formulas ϕ and ψ from $PROP^3_{PS}$ are said to be strongly equivalent, noted $\phi \equiv^3_{strong} \psi$, iff for every 3-interpretation I over PS, we have $I(\phi) = I(\psi)$. Note that this relation is more restrictive than the notion of 3-equivalence associated to the notion of 3-consequence. For instance, while $p \vee \neg p$ and $q \vee \neg q$ are 3-equivalent (in the sense that $p \vee \neg p \models^3 q \vee \neg q$ and $q \vee \neg q \models^3 p \vee \neg p$), they are not strongly equivalent since for every 3-interpretation I over PS s.t. $I(p) = 1$ and $I(q) = \frac{1}{2}$, we have $I(p \vee \neg p) = 1 \neq I(q \vee \neg q) = \frac{1}{2}$. Actually, there is no replacement theorem in J_3 for 3-equivalence, while a restricted replacement theorem obviously holds when strong equivalence is considered (because every connective is truth functional).

Proposition 3.5. For every $\phi \in PROP^3_{PS}$, we can derive in time linear in $|\phi|$ a $\Box NNF$ formula $\phi\downarrow$ that is strongly equivalent to ϕ.

The normalization process is based on the following strong equivalences (viewed as rewrite rules, left-to-right oriented):

- $\Box(\phi_1 \wedge \phi_2) \equiv^3_{strong} (\Box\phi_1 \wedge \Box\phi_2)$;
- $\Box(\phi_1 \vee \phi_2) \equiv^3_{strong} (\Box\phi_1 \vee \Box\phi_2)$;
- $\Box\Box\phi \equiv^3_{strong} \Box\phi$;
- $\neg(\phi_1 \wedge \phi_2) \equiv^3_{strong} (\neg\phi_1 \vee \neg\phi_2)$;
- $\neg(\phi_1 \vee \phi_2) \equiv^3_{strong} (\neg\phi_1 \wedge \neg\phi_2)$;
- $\neg\neg\phi \equiv^3_{strong} \phi$;
- $\neg true \equiv^3_{strong} false$;
- $\neg false \equiv^3_{strong} true$.

Let us illustrate it on the formula $\phi = \neg\Box(\neg(p \wedge q) \vee \Box(\neg(q \vee \neg p)))$. This is a 3-step process:

- First, the last five rewrite rules are systematically fired to push down every occurrence of $\neg$: $\phi = \neg\Box(\neg p \vee \neg q \vee \Box(\neg q \wedge p))$.
- Then, the square connectives are pushed down taking advantage of the first three rewrite rules. In the resulting formula, only literals are in the scope of any occurrence of $\Box$ (in other words, the formula is in $\Box NF$): $\phi = \neg(\Box\neg p \vee \Box\neg q \vee (\Box\neg q \wedge \Box p))$.
- Finally, using the last five rules again, the formula obtained at the end of the previous step is turned into a (strongly equivalent) $\Box NNF$: $\phi = \neg\Box\neg p \wedge \neg\Box\neg q \wedge (\neg\Box\neg q \vee \neg\Box p)$.

On this ground, it is easy to show that every formula ϕ from $PROP^3_{PS}$ can be turned into a strongly equivalent $\Box CNF$ formula. The reason is that for every triple of formulas ϕ_1, ϕ_2 and ϕ_3 from $PROP_{PS}^3$, the formulas $\phi_1 \vee (\phi_2 \wedge \phi_3)$ and $(\phi_1 \vee \phi_2) \wedge (\phi_1 \vee \phi_3)$ are strongly equivalent. The possibility of normalizing the belief base under $\Box CNF$ is particularly important for readibility reasons and the practical computing of the inference relations $\approx_S^{\mathcal{P},S_0}$. Moreover, this valuable property is usually not satisfied by paraconsistent logics (e.g., the one given in [23], the one given in [45], or the standard coherence-based approach).

It is well-known that $S-3$ logic satisfies the monotonicity property (Theorem 4.1 from [58]) stating that if $S \subseteq S'$ and $B \models_S^3 \gamma$, then $B \models_{S'}^3 \gamma$. Roughly speaking, increasing S never requires to question the consequences obtained so far. It is not difficult to show that this property still holds when $S-3$ inference is extended to the language $PROP_{PS}^3$, as this is the case in this paper. While monotonicity is a highly desirable property when B is classically consistent, that's not the case otherwise. Indeed, what is required is that increasing the computational power leads to increase the set of *expected* consequences, i.e., those required by common sense. It is natural that the epistemic status of beliefs changes when more evidence is got through the incorporation of additional explicit beliefs or through an improvement of inferential capabilities. Specifically, increasing the computational effort can lead to discover inconsistencies that were hidden before. Accordingly, monotonicity is not satisfied by our inference relations $\approx_S^{\mathcal{P},S_0}$ in the general case.

Example 3.6. Let $B = \{bird, (\neg bird \vee fly), \neg fly, (\neg bird \vee feathers)\}$. With $S = (\emptyset, \{bird\})$:

- $\mathcal{S}_{\mathcal{PO}}(B, S, S_0) = \{\{bird\}\}$
- $\mathcal{S}_{\mathcal{LO}}(B, S, S_0) = \{\{bird\}\}$
- $\mathcal{S}_{\mathcal{IP}}(B, S, S_0) = \{\{bird\}\}$
- $\mathcal{S}_{\mathcal{LE}}(B, S, S_0) = \{\{bird\}\}$

With $S' = (\emptyset, \{bird, fly\})$:

- $\mathcal{S}_{\mathcal{PO}}(B, S', S'_0) = \{\emptyset\}$
- $\mathcal{S}_{\mathcal{LO}}(B, S', S'_0) = \{\emptyset\}$
- $\mathcal{S}_{\mathcal{IP}}(B, S', S'_0) = \{\{bird\}, \{fly\}\}$
- $\mathcal{S}_{\mathcal{LE}}(B, S', S'_0) = \{\{bird\}, \{fly\}\}$

We have $B \approx_S^{\mathcal{IP},S_0} feathers$ but $B \not\approx_{S'}^{\mathcal{IP},S_0} feathers$. When S is considered, no argument against $feathers$ can be found because fly can be interpreted as $\frac{1}{2}$. Adding fly to S leads to question this attitude. This example also shows that the inference relations $\approx_S^{\mathcal{LE},S_0}$ do not satisfy the monotonicity property as well in the general case. Let $B = \{bird, fly, \neg fly, (\neg bird \vee feathers)\}$. With $S = (\emptyset, \{bird\})$:

- $\mathcal{S}_{\mathcal{PO}}(B, S, S_0) = \{\{bird\}\}$

- $\mathcal{S}_{\mathcal{LO}}(B, S, S_0) = \{\{bird\}\}$
- $\mathcal{S}_{\mathcal{IP}}(B, S, S_0) = \{\{bird\}\}$
- $\mathcal{S}_{\mathcal{LE}}(B, S, S_0) = \{\{bird\}\}$

With $S' = (\emptyset, \{bird, fly\})$:

- $\mathcal{S}_{\mathcal{PO}}(B, S', S'_0) = \{\emptyset\}$
- $\mathcal{S}_{\mathcal{LO}}(B, S', S'_0) = \{\emptyset\}$
- $\mathcal{S}_{\mathcal{IP}}(B, S', S'_0) = \{\{bird\}\}$
- $\mathcal{S}_{\mathcal{LE}}(B, S', S'_0) = \{\{bird\}\}$

With $\gamma = feathers$, we can also easily show that the inference relations $\approx_S^{\mathcal{PO},S_0}$ and $\approx_S^{\mathcal{LO},S_0}$ do not satisfy the monotonicity property.

Contrariwise to what this example illustrates, our inference relations do not satisfy the antimonotonicity property (if $S \subseteq S'$ and $B \not\models_S^3 \gamma$, then $B \not\models_{S'}^3 \gamma$) in the general case.

Example 3.7. Let $B = \{fly, \neg fly, bird, (\neg bird \vee feathers)\}$, $S = (\emptyset, \{fly\})$ and $S' = (\emptyset, \{bird\}, \{fly\})$. We have $\mathcal{S}_{\mathcal{IP}}(B, S, S_0) = \{\emptyset\}$ and $\mathcal{S}_{\mathcal{IP}}(B, S', S'_0) = \{\{bird\}\}$, so $B \not\approx_S^{\mathcal{IP},S_0} feathers$ but $B \approx_{S'}^{\mathcal{IP},S_0} feathers$.

Nevertheless, we have a "stratumwise" monotonicity property. This property is helpful to check whether $B \approx_S^{\mathcal{P},S_0} \gamma$ in an incremental fashion.

Proposition 3.6. Let B be a belief base from $PROP_{PS}^3$. Let $S = (S_0, \ldots, S_j)$ be a stratification of a given subset of PS where S_0 is s.t. $B \not\models_{S_0}^3 false$. Let γ be a formula from $PROP_{PS}$. For $\mathcal{P} \in \{\mathcal{PO}, \mathcal{LO}, \mathcal{IP}, \mathcal{LE}\}$, for every $i \in 0 \ldots j-1$, if $B \approx_{(S_0,\ldots,S_i)}^{\mathcal{P},S_0} \gamma$, then $B \approx_{(S_0,\ldots,S_{i+1})}^{\mathcal{P},S_0} \gamma$.

Finally, like any paraconsistent inference relation, our relations do not satisfy left logical equivalence. For the same reason, they do not satisfy right weakening (since they satisfy reflexivity).

Cautiousness.

In the spectrum of inference relations in our framework, one extreme bound corresponds to the case $S = \emptyset$, while at the other extreme bound, we have $S = PS$. None of these two extreme cases can be considered as a "reasonable" inference relation for common-sense reasoning. On the one hand, the latter case corresponds to an ideal agent, with unlimited inferential capabilities: $\models_{PS}^3$ coincides with $\models$. Accordingly, if B (from $PROP_{PS}$) is classically consistent, all our inference relations $B \approx_{PS}^{\mathcal{P},S_0}$ coincide with $B \models$, whatever $\mathcal{P} \in \{\mathcal{PO}, \mathcal{LO}, \mathcal{IP}, \mathcal{LE}\}$ (admittedly, this contrasts with many paraconsistent inference relations). On the other hand, the former case corresponds to an agent with very limited reasoning capabilities. Indeed, if $S = \emptyset$ and B is a belief base from $PROP_{PS}$, all our inference relations correspond to Levesque's 3-entailment [44]. This inference relation is tractable when CNF queries are considered and trivialization is avoided. However, with this inference relation, disjunctive syllogism can never be applied: it is not possible to conclude fly from $bird \wedge (\neg bird \vee fly)$. This suggests that cautiousness must be taken into account as an important criterion in the choice of a relation. As usual, a first inference relation is said to be at least as cautious as a second one whenever the set of consequences of the former is included into the set of consequences of the latter.

In the case where S is a consistent subset, all our inference relations coincide with $\models^3_S$, so they are just as cautious as $\models^3_S$, and the size of S can be viewed as a degree of cautiousness.

In the remaining case, our inference relations are more cautious relations than $\models^3_S$ but this is what is expected since trivialization must be avoided. In this situation, our inference relations do not coincide in the general case, and results similar to the corresponding ones in the standard coherence-based approach can be obtained; especially, the less cautious inference relations correspond to the $\mathcal{LO}$ and $\mathcal{LE}$ policies.

Proposition 3.7.

- $\approx^{\mathcal{PO},S_0}_S$ is at least as cautious as $\approx^{\mathcal{LO},S_0}_S$ and $\approx^{\mathcal{IP},S_0}_S$, but the converse does not always hold.
- $\approx^{\mathcal{IP},S_0}_S$ is at least as cautious as $\approx^{\mathcal{LE},S_0}_S$ but the converse does not always hold.
- $\approx^{\mathcal{LO},S_0}_S$ cannot be compared with $\approx^{\mathcal{IP},S_0}_S$ and $\approx^{\mathcal{LE},S_0}_S$ w.r.t. cautiousness in the general case.

As to cautiousness, it is quite difficult to compare our approach with the standard one because our inference relations have some additional parameters that may have a great influence w.r.t. this issue. Thus, limiting S to a bounded subset of PS ensures tractability for clausal queries (provided that B is from $PROP_{PS}$) but it can give rise to very cautious inference relations compared with the standard approach. For instance, if we consider the one-stratum belief base $B = (\{summer, \neg summer, summer \vee conifers\_have\_needles, \neg summer \vee conifers\_have\_needles\})$, then, using the standard approach, we can infer $conifers\_have\_needles$; but we cannot infer $conifers\_have\_needles$ in our approach[7] since all the formulas of B can be satisfied by interpreting $summer$ to $\frac{1}{2}$. This does not mean that our inference relations are always more cautious than the corresponding ones in the standard approach, even in the case where

[7] Nevertheless, we have $\{summer, \neg summer, \Box(summer \vee conifers\_have\_needles), \Box(\neg summer \vee conifers\_have\_needles)\} \approx^{\emptyset,\mathcal{IP}}_{PS} conifers\_have\_needles$.

$S \neq PS$. For instance, let $B = (\{summer \wedge \neg summer \wedge conifers\_have\_needles\})$ be a one-stratum SBB. As it is the case for any belief base containing only one (classically) inconsistent formula, only tautologies are consequences of B w.r.t. any of the four inference relations in the standard approach. Contrastingly, $conifers\_have\_needles$ is a consequence of B w.r.t. any of the four inference relations in our approach for every choice of S such that $summer \notin S_0$.

3.5. A general framework

Our framework is also general. As evoked previously, it includes both J_3 inference and $S-3$ inference as specific cases. As a consequence, it also includes the inference relation $\vdash_{LP}$ of the 3-valued logic LP [54], the restriction of J_3 on the $\{\neg, \wedge, \vee\}$ fragment. More interestingly, let $\leq$ be the preference ordering over 3-interpretations defined by $I \leq J$ iff $\{x \in PS \mid I(x) = \frac{1}{2}\} \subseteq \{x \in PS \mid J(x) = \frac{1}{2}\}$; then, the relation $\vdash_{LP_m}$ [55] defined by $\Sigma \vdash_{LP_m} \gamma$ iff every 3-model of Σ that is minimal in its set w.r.t. $\leq$ is a (3-)model of γ, can be recovered as a specific case of our $\approx_S^{\mathcal{IP},S_0}$ relation:

Proposition 3.8. Let Σ and γ be two formulas from $PROP_{PS}$.
We have $\Sigma \vdash_{LP_m} \gamma$ iff $\{\Sigma\} \approx_{PS}^{\mathcal{IP},\emptyset} \gamma$.

Our work is also closely related to the approach to paraconsistent reasoning given in [12]. In this work, every NNF formula Σ of $PROP_{PS}$ is associated to a default theory $\langle \Sigma^{\pm}, D_{PS} \rangle$ where:

- $\Sigma^{\pm}$ is a formula in the language $PROP_{PS^{\pm}}$ where $PS^{\pm} = \{x^+ \mid x \in PS\} \cup \{x^- \mid x \in PS\}$; $\Sigma^{\pm}$ is obtained by replacing in Σ every occurrence of a positive literal x by the positive literal x^+ and every occurrence of a negative literal $\neg x$ by the positive literal x^-.
- $D_{PS} = \{\delta_x \mid x \in PS\}$ is a set of default rules

$$\delta_x = \frac{: x^+ \Leftrightarrow \neg x^-}{(x \Leftrightarrow x^+) \wedge (\neg x \Leftrightarrow x^-)}$$

Among the inference relations that can be defined from $\langle \Sigma^{\pm}, D_{PS} \rangle$ is skeptical signed inference; a formula γ of $PROP_{PS}$ is a skeptical signed consequence of Σ, noted $\Sigma \vdash_s^{\pm} \gamma$ iff $\gamma^{\pm}$ belongs to every extension of $\langle \Sigma^{\pm}, D_{PS} \rangle$.

Interestingly, the relation $\vdash_s^{\pm}$ can be recovered as a specific case of our $\approx_S^{\mathcal{IP},S_0}$ relation.

Proposition 3.9. Let Σ and γ be two formulas from $PROP_{PS}$. We have $\Sigma \vdash_s^{\pm} \gamma$ iff $\{\Sigma\} \approx_{PS}^{\mathcal{IP},\emptyset} \gamma$.

Finally, our approach encompasses as well the skeptical inference relations $\mid\!\sim^{\mathcal{P}}$ of the standard coherence-based approach as defined in [53], [5] and [6].

Proposition 3.10. Let $B = (\Delta_1, \ldots, \Delta_k)$ be an SBB from $PROP_{PS}$. We associate to B in polynomial time the belief base $B_{new} = \{new_{i,j}, \neg new_{i,j} \vee \phi_{i,j} \mid \phi_{i,j} \in \Delta_i\}$ and the stratification $S = (Var(B), \{new_{1,j} \mid \phi_{1,j} \in \Delta_1\}, \ldots, \{new_{k,j} \mid \phi_{k,j} \in \Delta_k\})$ where all the $new_{i,j}$ are new variables (not occurring in B). For every selection policy $\mathcal{P}$ among $\mathcal{PO}$, $\mathcal{LO}$, $\mathcal{IP}$, $\mathcal{LE}$ and every formula γ from $PROP_{PS}$ not containing a new variable, we have $B \mid\!\sim^{\mathcal{P}} \gamma$ iff $B_{new} \mid\!\approx_S^{\mathcal{P}, Var(B)} \gamma$.

Focusing on variables instead of beliefs gives less syntax-sensitivity and some flexibility that is hardly achieved by the standard coherence-based approach. For instance, let us assume that B gathers beliefs stemming from two different sources of information. The first source states $fly \wedge penguin$ while the second source states $\neg fly \wedge \neg penguin$. The two sources are equally reliable; however, the first source is more reliable to what concerns *penguin* while the second source is more reliable to what concerns fly. With $\mathcal{P} \in \{\mathcal{IP}, \mathcal{LE}\}$, the standard coherence-based approach gives $(\{fly \wedge penguin, \neg fly \wedge \neg penguin\}) \mid\!\sim^{\mathcal{P}} fly \vee \neg penguin$, but the expected conclusion is just the *negation* of $fly \vee \neg penguin$! Indexing variables with sources (e.g. fly_1, fly_2, $\ldots$), we can express in our framework that *penguin* from source 1. is more reliable than fly from source 1. and the opposite holds when source 2. is considered. Indeed, with $B = \{(\neg fly_1 \vee fly) \wedge (\neg fly \vee fly_1) \wedge (\neg fly_2 \vee fly) \wedge (\neg fly \vee fly_2) \wedge (\neg penguin_1 \vee penguin) \wedge (\neg penguin \vee penguin_1) \wedge (\neg penguin_2 \vee penguin) \wedge (\neg penguin \vee penguin_2) \wedge fly_1 \wedge penguin_1 \wedge \neg fly_2 \wedge \neg penguin_2\}$, $\mathcal{P} \in \{\mathcal{PO}, \mathcal{LO}, \mathcal{IP}, \mathcal{LE}\}$, and $S = (\{fly, penguin\}, \{penguin_1, fly_2\}, \{fly_1, penguin_2\})$, we have $B \mid\!\approx_S^{\mathcal{P}, S_0} \neg fly \wedge penguin$ and we do **not** have $B \mid\!\approx_S^{\mathcal{P}, S_0} fly \vee \neg penguin$.

Conversely, all our inference relations can be recovered as skeptical inference relations from the standard coherence-based approach. We have:

Proposition 3.11. Let B be a belief base from $PROP_{PS}$ and γ a formula from $PROP_{PS}$. Let $S = (S_0, \ldots, S_j)$ be a stratification of a given subset of PS where S_0 is s.t. $B \not\models^3_{S_0} false$. Let $\mathcal{P}$ be any selection policy among $\mathcal{PO}, \mathcal{LO}, \mathcal{IP}$ and $\mathcal{LE}$. We have $B \mid\!\approx_S^{\mathcal{P}, S_0} \gamma$ iff $(B^{\pm} \cup \bigcup_{x \in PS} \{x^+ \vee x^-\} \cup \bigcup_{x \in S_0} \{\neg x^+ \vee \neg x^-\}, \{\neg x^+ \vee \neg x^- \mid x \in S_1\}, \ldots, \{\neg x^+ \vee \neg x^- \mid x \in S_j\}) \mid\!\sim^{\mathcal{P}} \gamma^{\pm}$.

This reduction can be used to implement our inference relations, taking advantage of existing algorithms for the standard coherence-based approach (see e.g. [18]). Both propositions 3.10 and 3.11 show our approach and the standard coherence-based one computationally equivalent. Accordingly, the previous example can also be handled by the standard coherence-based approach but the required encoding is much more sophisticated (as the translation based on Proposition 3.11 illustrates). A strong point of our approach is that it does not need the (non truth functional) comma connective at work in the standard coherence-based approach.

3.6. Dealing with intractability

We assume the reader familiar with some basic notions of computational complexity, especially the complexity classes P, NP, and coNP, and the classes Δ_k^p, Σ_k^p and Π_k^p of the polynomial hierarchy PH (see [52] for details). $\Delta_2^p[\mathcal{O}(log\ n)]$ is the class of problems that can be decided using only logarithmically many calls to an NP oracle. Let us recall that a decision problem is said to be at the k^{th} level of PH iff it belongs to Δ_{k+1}^p, and is either Σ_k^p-hard or Π_k^p-hard. It is strongly believed that PH does not collapse (at any level), i.e., is a truly infinite hierarchy (for every integer k, $\mathsf{PH} \neq \Sigma_k^p$).

The complexity of $S-3$ logic has been investigated in depth in [44] [58] [15]. In order to circumvent the four limitations of $S-3$ inference (as reported in Section 3.1), we had to extend the framework. Let us first investigate how the complexity of inference is altered by considering a more expressive language than the one of $S-3$ logic, as this is done in this paper.

Definition 3.8. CLAUSE $\models_S^3$ is the following decision problem:

- **Input:** A belief base B from $PROP_{PS}^3$, a subset S of PS, and a CNF formula γ from $PROP_{PS}$.
- **Query:** Does $B \models_S^3 \gamma$ hold?

LITERAL $\models_S^3$ is the restriction of this problem to the case where γ is a conjunction of literals.

Contrariwise to what happens when B is from $PROP_{PS}$ (i.e., when Schaerf and Cadoli's $S-3$ logic is considered), $\models_S^3$ is intractable in our framework, even in the case where queries are literals:

Proposition 3.12. CLAUSE $\models_S^3$ and LITERAL $\models_S^3$ are coNP-complete, even in the case where $S = \emptyset$.

The complexity shift is not surprising due to the well-known expressivity/efficiency tradeoff in knowledge representation (often attributed to Brachman and Levesque). Especially, Proposition 3.12 is an easy consequence of the fact that classical propositional logic (from the fragment $\{\neg, \wedge, \vee\}$) can be translated into J_3, thanks to the $\Box$ connective.

Fortunately, tractability can be achieved whenever $|S|$ is bounded and a limited amount of knowledge is taken into account. This can be measured by the cardinality of the following set $Nec(B)$.

Definition 3.9. Let $\phi\downarrow$ be the $\Box NNF$ formula associated with ϕ as given in Proposition 3.5. Let B be a (conjunctively-interpreted) belief base from $PROP^3_{PS}$. We note $Nec(B) = \{x \in PS \mid \Box x$ or $\Box\neg x$ is a subformula of $\phi\downarrow$ where $\phi \in B\}$.

We have obtained the following proposition, that partly extends Theorem 4.3 from [58]:

Proposition 3.13. Let B be a (conjunctively-interpreted) belief base from $PROP^3_{PS}$. Let S be a subset of PS and γ be a CNF formula from $PROP_{PS}$. CLAUSE $\models^3_S$ can be decided in time $\mathcal{O}(2^{2\cdot|Nec(B)|+|S|}\cdot|B|\cdot|\gamma|)$.

The translation to classical logic used to prove this proposition (see Appendix) is valuable for the implementation issue, especially in the case where B is a $\Box CNF$ formula; indeed in this situation, every instance of CLAUSE $\models^3_S$ can be turned in linear time into an instance of the famous SATisfiability problem for which many "practically efficient" implementations (especially those based on the Davis/Logemann/Loveland procedure [25], like $C-SAT$ [27]) have been pointed out for the past ten years.

Let us finally consider the more sophisticated inference relations introduced in the paper. Despite the generality of our framework (cf. Section 3.5), our inference relations are just as hard as the corresponding ones in the standard coherence-based approach in the general case:

Definition 3.10. For every selection policy $\mathcal{P}$, CLAUSE $\approx^{\mathcal{P},S_0}_S$ is the following decision problem:

- **Input:** A belief base B from $PROP^3_{PS}$, a stratification $S = (S_0,\ldots,S_j)$ of PS, a subset S_0 of S s.t. $B \not\models^3_{S_0} false$ and a CNF formula γ from $PROP_{PS}$.
- **Query:** Does $B \approx^{\mathcal{P},S_0}_S \gamma$ hold?

LITERAL $\approx^{\mathcal{P},S_0}_S$ is the restriction of this problem to the case where γ is a conjunction of literals.

Proposition 3.14. The complexity of CLAUSE $\approx^{\mathcal{P},S_0}_S$ and LITERAL $\approx^{\mathcal{P},S_0}_S$ from a belief base B given a stratification $S = (S_0,\ldots,S_j)$ of a set of propositional variables, and its restrictions to literal inference for $\mathcal{P} \in \{\mathcal{PO},\mathcal{LO},\mathcal{IP},\mathcal{LE}\}$ is reported in Table 1.

$\mathcal{P}$	CLAUSE / LITERAL $\approx_S^{\mathcal{P},S_0}$
$\mathcal{PO}$	$\Delta_2^p[\mathcal{O}(log\ n)]$-complete
$\mathcal{LO}$	Δ_2^p-complete
$\mathcal{IP}$	Π_2^p-complete
$\mathcal{LE}$	Δ_2^p-complete

Table 1

This is an easy consequence of the polynomial reductions whose existence is stated by Propositions 3.10 and 3.11. Now, every "reasonable" inference relation, corresponding to an agent with limited computational resources corresponds to an S of bounded size, and a limited amount of knowledge, corresponding to a $Nec(B)$ of bounded size. Interestingly, limiting both the size of S and the size of $Nec(B)$ is sufficient to ensure tractability. The next proposition shows that our second objective is reached in this case:

Proposition 3.15. $\forall \mathcal{P} \in \{\mathcal{PO}, \mathcal{LO}, \mathcal{IP}, \mathcal{LE}\}$, the complexity of CLAUSE $\approx_S^{\mathcal{P},S_0}$ from a belief base B from $PROP_{PS}^3$ s.t. $|Nec(B)|$ is **bounded**, given a stratification $S = (S_0, \ldots, S_j)$ of a **bounded** set of propositional variables is in P.

From the implementation side, we can again take advantage of the translation to classical logic shown in the proof of Proposition 3.13; thus, algorithms for the WEIGHTED-MAX-SAT problem can be exploited with profit to solve in practice the CLAUSE $\approx_S^{\mathcal{LE},S_0}$ problem. Proposition 3.15 also shows that checking whether $B \approx_{(S_0,\ldots,S_j)}^{\mathcal{P},S_0} \gamma$ in a stratumwise fashion (as stated by Proposition 3.6) is a valuable strategy; indeed, the worst case situation for such a stratumwise computing of CLAUSE $\approx_S^{\mathcal{P},S_0}$ is reached whenever each S_i ($i \in 0 \ldots j$) is a singleton (i.e., S contains j variables and it is stratified into j classes) and for each $i \in 0 \ldots j-1$, we have $B \not\approx_{(S_0,\ldots,S_i)}^{\mathcal{P},S_0} \gamma$: in this situation, the computational effort spent performing these j tests is wasted. The point is that this computational effort can be neglected, compared with the one required to answer CLAUSE $\approx_S^{\mathcal{P},S_0}$ in a direct way (even if a naive algorithm consisting in computing $\mathcal{S}(B, \bigcup_{k=0}^{i} S_k, S_0)$ from the scratch at each step i is used). Especially, the difference of the dominant factors as to time complexity is $2^{2.j} - \sum_{i=0}^{j-1} 2^{2.i}$, which is in $\Omega(4^j)$.

As the previous proposition shows it, a major advantage of our approach over other techniques is that it enables *resource-bounded reasoning*: the computational resources $|S|$ can be tuned to give rise to realistic, computationally viable, inference relations, each time a limited amount of knowledge must be dealt with. The main reason is that when $|S| \leq b$, there is only 2^b subsets of S, which is a constant when b is a constant. Accordingly, the naive algorithm consisting in generating $\mathcal{S}_{\mathcal{P}}(B, S, S_0)$ by filtering out the preferred consistent elements of 2^S, and testing for every of the resulting elements S' whether or not $B \models_{S'} \gamma$ holds runs in time $\mathcal{O}(2^{2\cdot(b+|Nec(B)|)} \cdot |B| \cdot |\gamma|)$, hence in $\mathcal{O}(|B| \cdot |\gamma|)$ for a constant b and a bounded $|Nec(B)|$, for every B and a CNF query γ. Of course, this is a "trick" in some sense but we can hardly do better: unless $\mathsf{P} = \mathsf{NP}$, there is no tractable inference relation that is general enough to include $\models$ as a specific case. In addition, we claim that focusing on $B \models^3_S$ as a tractable approximation of $B \models$ in order to design a paraconsistent inference relation is not so bad, especially when compared with other possible choices. Indeed, let $B \models\!\simeq$ be a tractable approximation by below of $B \models$ upon which we want to build a tractable paraconsistent inference relation. Requiring that such a tractable paraconsistent relation coincides with $B \models\!\simeq$ whenever $B \models\!\simeq$ is not explosive and is obtained by weakening $B \models\!\simeq$ through the removal of some beliefs of B otherwise limits the spectrum of interesting relations $\models\!\simeq$. For instance, considering the restriction $B \vdash_{unit}$ of $B \models$ where only unit-refutation derivable consequences of B are considered would lead to a complex semantics with respect to which $\wedge$ and $\vee$ do not behave classically [21]; considering the restriction $B_{Horn} \models$ of $B \models$ where B_{Horn} is the subset of Horn clauses of B would not lead to a tractable relation [18] [51].

4. Other related work

In this paper, we are mainly concerned with the design of inference relations that are both paraconsistent and tractable. A close look at the AI literature shows that both issues can be handled by considering *infraclassical* (or subclassical) inference relations, i.e., *approximations by default* of $B \models$. Such inference relations can be achieved by restricting the proof theory, or by focusing on some subsets of the belief base. What is quite surprising is the fact that the two issues (trivialization and intractability) have usually **not** been addressed alltogether!

On the one hand, among the various formalisms for dealing with inconsistency (see [9] for a survey), the notion of approximation plays a key role. In paraconsistent logics, the proof theory is limited to a strict subset of the proof theory of classical logic with the objective of avoiding trivialization. In the standard coherence-based approach, the focus is laid on some preferred consistent subbases of B, which can be considered as approximations of B by default. Such preferred subbases can also be defined by splitting the belief base B into several micro-theories characterized by their sets of variables, and selected in a dynamic way by looking at the variables occurring in the queries [19]. Unfortunately, for all these approaches, the corresponding inference relations are typically at least as hard as classical inference in the general case[8] [20]. For instance, quasi-classical entailment [8] [37] is a paraconsistent infraclassical relation but it is not tractable in the general case [47].

[8] Some tractable cases can be achieved sometimes by restricting B to some formulas for which consistency checks can be achieved efficiently, and/or by limiting both the number and the size of the preferred subbases.

On the other hand, many approximation techniques have been developed so far to deal with the intractability issue (e.g., [42] [33] [22] [58] [59] [24] [21]); in such approaches, the inference relation or the knowledge base itself are approximated into computationally easier one(s). However, the trivialization issue is typically not addressed. For instance, the inference relation defined in [59] always is explosive when B is inconsistent, while the ones given in [58] [21] are explosive in the general case when B is inconsistent.

Our approach is clearly related to all these works. Specifically, it is directly inspired by two approximation techniques, $S-3$ logic and the standard coherence-based approach, each of them being developed for a specific purpose (dealing with intractability versus avoiding trivialization). Actually, these two approaches can be recovered as specific cases in our framework. By ensuring that all our infraclassical inference relations are paraconsistent and tractable when the size of S is bounded and a limited amount of knowledge is present in the belief base, we just succeed in taking the best of these two approaches. Since we have intensively presented the relationships and differences between our approach, J_3 logic, $S-3$ logic and the standard coherence-based approach in the course of the previous sections, we refrain from repeating these here.

5. Conclusion

The main contribution of this paper is a family of resource-bounded inference relations for reasoning from inconsistent belief bases, and an analysis of them both from the logical side and the computational side. In contrast to many approaches to inconsistency handling, the computational issue is very central in our framework. Especially, the computational resources can be tuned to give rise to realistic, computationally viable, inference relations. Nevertheless, our framework is general enough to encompass several multi-valued paraconsistent logics, the standard coherence-based approach to inconsistency handling, and some signed systems for paraconsistent reasoning. Some of our resource-bounded inference relations have been implemented in JAVA. The code is available at URL `http://www.cril.univ-artois.fr/~porquet/projet.html`. A next step w.r.t. the implementation issue is to make some experiments in order to compare the cautiousness of policies on practical instances, depending on several parameters (especially the size of S).

Several additional families of inference relations could be defined in our framework. Specifically, other inference principles (including credulous inference, argumentative inference and prudent inference) could be considered. Some additional inference relations could also be obtained by considering a formula as derivable iff it is true (and not false) given B for every preferred consistent subset. Additionally, several policies for the incorporation of variables from queries into S are also to be considered. Investigating the computational and logical properties of the corresponding inference relations is a topic for further research. Additional work will also encompass the methodological issue. Our analysis of both the cautiousness and the computational complexity of inference relations offers some basic criteria for chosing a selection policy. We plan also to expand on [38], where heuristics for the choice of S are presented.

Appendix

Proof of Proposition 3.1:
It is obvious that the function that maps I to $\hat{I}$ is bijective. Now, the remaining proof is easy by structural induction on ϕ. ■

Proof of Proposition 3.2:
By construction, each of the sets $\mathcal{S}_{\mathcal{PO}}(B,S,S_0)$ and $\mathcal{S}_{\mathcal{LO}}(B,S,S_0)$ contains a unique element. Because B is 3-consistent (this is a consequence of the existence of S_0), the set $\mathcal{S}(B,S,S_0)$ is not empty. Moreover, it is finite since PS is finite. Since a finite non-empty set always has a minimal element, we obtain that $\mathcal{S}_{\mathcal{IP}}(B,S,S_0)$ and $\mathcal{S}_{\mathcal{LE}}(B,S,S_0)$ are not empty. Because every element S' of $\mathcal{S}_{\mathcal{P}}(B,S,S_0)$ satisfies $B \not\models^3_{S'} false$ (by definition), we also have $B \not\approx^{\mathcal{P},S_0}_S false$ for every $\mathcal{P}$ in $\mathcal{PO}$, $\mathcal{LO}$, $\mathcal{IP}$, $\mathcal{LE}$. ■

Proof of Proposition 3.3:

- Let us first show that if $B \approx^{\mathcal{IP},S_0}_S \gamma$, then $B \approx^{\mathcal{IN},S_0}_S \gamma$. The monotonicity property for $\models^3_S$ states that if $S \subseteq S'$ and $B \models^3_S \gamma$, then $B \models^3_{S'} \gamma$ (for every formula γ from $PROP_{PS}$). This shows that the pre-ordering corresponding to the $\mathcal{IN}$ policy refines the pre-ordering $\supseteq$ corresponding to the $\mathcal{IP}$ policy; specifically, for every S' in $\mathcal{S}_{\mathcal{IN}}(B,S,S_0)$, there exists a S'' in $\mathcal{S}_{\mathcal{IP}}(B,S,S_0)$ s.t. S' and S'' are equally preferred w.r.t. the $\mathcal{IN}$ policy, i.e., B has the same consequences w.r.t. $\mathcal{IN}$, S' and S_0 or $\mathcal{IP}$, S'' and S_0. Hence, the conclusion follows.
- Conversely, let us show that if $B \approx^{\mathcal{IN},S_0}_S \gamma$, then $B \approx^{\mathcal{IP},S_0}_S \gamma$. To achieve this goal, it is sufficient to show that any pair of distinct elements S_1, S_2 of $\mathcal{S}_{\mathcal{IP}}(B,S,S_0)$ are incomparable w.r.t. the pre-ordering corresponding to the $\mathcal{IN}$ policy. So let S_1, S_2 be two distinct elements of $\mathcal{S}_{\mathcal{IP}}(B,S,S_0)$. By definition, there exist $x_1 \in S_1 \setminus S_2$ and $x_2 \in S_2 \setminus S_1$. Let $\gamma_1 = x_1 \wedge \neg x_1$ and $\gamma_2 = x_2 \wedge \neg x_2$. By construction, there is no $S_1 - 3$ model of γ_1 and no $S_2 - 3$ model of γ_2. However, the existence of S_1 and S_2 of $\mathcal{S}_{\mathcal{IP}}(B,S,S_0)$ shows that B is both $S_1 - 3$ consistent and $S_2 - 3$ consistent. Accordingly, we have $B \not\models^3_{S_1} \gamma_1$ and $B \not\models^3_{S_2} \gamma_2$. It remains to show that $B \models^3_{S_2} \gamma_1$ and $B \models^3_{S_1} \gamma_2$. Let us show that $B \models^3_{S_2} \gamma_1$ (the remaining case is similar). It is sufficient to prove that every $S_2 - 3$ model I of B satisfies $I(x_1) = \frac{1}{2}$. Assume that this is not the case. Then there exists an $S_2 - 3$ model I of B s.t. $I(x_1) = 1$ or

$I(x_1) = 0$. Accordingly, B is $S_2 \cup \{x_1\} - 3$ consistent. This would contradict the fact that S_2 belongs to $\mathcal{S}_{\mathcal{IP}}(B, S, S_0)$.

■

Proof of Proposition 3.4:
By definition, for each selection policy $\mathcal{P}$ under consideration, we have $B{\approx}_S^{\mathcal{P},S_0}\gamma$ iff $B \models^3_{S'} \gamma$ for every $S' \in \mathcal{S}_{\mathcal{P}}(B, S, S_0)$. Equivalently, we have $B{\approx}_S^{\mathcal{P},S_0}\gamma$ iff every $S' - 3$ model of B is a 3-model of γ for $S' \in \mathcal{S}_{\mathcal{P}}(B, S, S_0)$ iff every preferred 3-model of B is a 3-model of γ, where a preferred 3-model of B is a $S' - 3$ model of B with $S' \in \mathcal{S}_{\mathcal{P}}(B, S, S_0)$. ■

Proof of Proposition 3.5:
Easy by structural induction on ϕ, taking advantage of the following equivalences as rewrite rules (oriented from the left-hand side to the right-hand side):

- $\Box(\phi_1 \wedge \phi_2) \equiv^3_{strong} (\Box\phi_1 \wedge \Box\phi_2)$;
- $\Box(\phi_1 \vee \phi_2) \equiv^3_{strong} (\Box\phi_1 \vee \Box\phi_2)$;
- $\Box\Box\phi \equiv^3_{strong} \Box\phi$;
- $\neg(\phi_1 \wedge \phi_2) \equiv^3_{strong} (\neg\phi_1 \vee \neg\phi_2)$;
- $\neg(\phi_1 \vee \phi_2) \equiv^3_{strong} (\neg\phi_1 \wedge \neg\phi_2)$;
- $\neg\neg\phi \equiv^3_{strong} \phi$;
- $\neg true \equiv^3_{strong} false$;
- $\neg false \equiv^3_{strong} true$.

■

Proof of Proposition 3.6:
Let $S = (S_0, \ldots, S_j)$ be a stratification of the set S. For each selection policy $\mathcal{P} \in \{\mathcal{PO}, \mathcal{LO}, \mathcal{IP}, \mathcal{LE}\}$ and for each $i \in 0 \ldots j-1$, we have (by definition):

$$B \approx^{\mathcal{P},S_0}_{(S_0,\ldots,S_i)}\gamma \text{ iff } \forall S' \in \mathcal{S}_{\mathcal{P}}(B, \bigcup_{k=0}^{i} S_k, S_0), B \models^3_{S'} \gamma$$

By construction, for every $S' \in \mathcal{S}_{\mathcal{P}}(B, \bigcup_{k=0}^{i} S_k, S_0)$, there exists $S'' \in \mathcal{S}_{\mathcal{P}}(B, \bigcup_{k=0}^{i+1} S_k, S_0)$ s.t. $S' \subseteq S''$. Subsequently, the preferred models of B obtained at step $i+1$ are among the preferred models of B considered at step i. The conclusion follows immediately. ■

Proof of Proposition 3.7:
By construction, it is obvious that the unique preferred consistent subset of S w.r.t. B and S_0 for the possibilistic policy is a subset of the unique preferred

consistent subset of S w.r.t. B and S_0 for the linear order policy, and a subset of every preferred consistent subset of S w.r.t. B and S_0 for the inclusion-preference policy. It is also obvious that every preferred consistent subset of S w.r.t. B and S_0 for the lexicographic policy is a preferred consistent subset of S w.r.t. B and S_0 for the inclusion-preference policy. Because $S-3$ inference is monotonic in the sense that the set of consequences increases when S increases, and because we are concerned with skeptical inference, we immediately obtain that $\approx_S^{\mathcal{PO},S_0}$ is at least as cautious as $\approx_S^{\mathcal{LO},S_0}$ and $\approx_S^{\mathcal{IP},S_0}$, and $\approx_S^{\mathcal{IP},S_0}$ is at least as cautious as $\approx_S^{\mathcal{LE},S_0}$. It remains to show that:

- $\approx_S^{\mathcal{LO},S_0}$ is not at least as cautious as $\approx_S^{\mathcal{PO},S_0}$. With $B=\{a,\neg a,b,\neg b\vee c\}$ and $S=(\emptyset,\{a\},\{b\})$, we have $B\approx_S^{\mathcal{LO},S_0}c$ but $B\not\approx_S^{\mathcal{PO},S_0}c$.
- $\approx_S^{\mathcal{LE},S_0}$ is not at least as cautious as $\approx_S^{\mathcal{IP},S_0}$. With $B=\{a,\neg a\vee b,\neg b,\neg a\vee c,\neg c,b\vee c\vee d\}$ and $S=(\emptyset,\{a,b,c\})$, we have $B\approx_S^{\mathcal{LE},S_0}d$ but $B\not\approx_S^{\mathcal{IP},S_0}d$.
- $\approx_S^{\mathcal{IP},S_0}$ is not at least as cautious as $\approx_S^{\mathcal{LO},S_0}$. With $B=\{a,\neg a,b,\neg b\vee c\}$ and $S=(\emptyset,\{a,b\})$, we have $B\approx_S^{\mathcal{IP},S_0}d$ but $B\not\approx_S^{\mathcal{LO},S_0}d$. Since we also have $B\approx_S^{\mathcal{LE},S_0}d$ but $B\not\approx_S^{\mathcal{PO},S_0}d$, this example shows that $\approx_S^{\mathcal{IP},S_0}$ is not at least as cautious as $\approx_S^{\mathcal{PO},S_0}$, and $\approx_S^{\mathcal{LE},S_0}$ is not at least as cautious as both $\approx_S^{\mathcal{LO},S_0}$ and $\approx_S^{\mathcal{PO},S_0}$.
- $\approx_S^{\mathcal{LO},S_0}$ is not at least as cautious as $\approx_S^{\mathcal{IP},S_0}$. With $B=\{\neg a,\neg a\vee b,\neg b,c,\neg c\vee d,\neg d,\neg a\vee\neg d\vee e\}$ and $S=(\emptyset,\{a\},\{b,c\},\{d\})$, we have $B\approx_S^{\mathcal{LO},S_0}d$ but $B\not\approx_S^{\mathcal{IP},S_0}d$. Since we also have $B\not\approx_S^{\mathcal{LE},S_0}d$, this example shows that $\approx_S^{\mathcal{LO},S_0}$ is not at least as cautious as $\approx_S^{\mathcal{LE},S_0}$.

■

Proof of Proposition 3.9:
By definition, γ is a skeptical signed consequence of Σ iff $\gamma^{\pm}$ belongs to every extension of $\langle\Sigma^{\pm},D_\Sigma\rangle$. From Theorems A.1 and A.2 of [12], this is equivalent to state that $\gamma^{\pm}$ belongs to every extension of the supernormal default theory $\langle\Sigma^{\pm},\{x^+\Leftrightarrow\neg x^-\mid x\in PS\}\rangle$. This is also equivalent to state that $(\{\Sigma^{\pm}\},\{x^+\Leftrightarrow\neg x^-\mid x\in PS\})\mid\!\sim^{\mathcal{IP}}\gamma^{\pm}$. Now, from the well-known preferred models characterization of skeptical inference from a SBB w.r.t. the $\mathcal{IP}$ policy, this is still equivalent to state that every preferred classical model of $\Sigma^{\pm}$ is a model of $\gamma^{\pm}$, where the preferred models of $\Sigma^{\pm}$ are the minimal models of $\Sigma^{\pm}$ w.r.t. the partial ordering $\leq$ over $BOOL^{PS^{\pm}}$ defined by $I\leq J$ iff $\{x\in PS\mid I(x^+)=I(x^-)\}\subseteq\{x\in PS\mid J(x^+)=J(x^-)\}$ (indeed, I satisfies $x^+\Leftrightarrow\neg x^-$ iff $I(x^+)\neq I(x^-)$).

Now, $\Sigma^{\pm}$ is a monotone (positive) formula from $PROP_{PS^{\pm}}$. As a consequence, if I is a model of $\Sigma^{\pm}$ s.t. $I(x^+)=I(x^-)=0$ for some $x\in PS$, then every interpretation J that coincides with I for every variable from $PS^{\pm}$ (except possibly on x^+ and x^-) is still a model of $\Sigma^\pm$. This is specifically the case for the interpretation J s.t. $J(x^+) \neq J(x^-)$. Since $J < I$ holds (where $<$ is the strict ordering associated with $\leq$), every preferred model M of $\Sigma^\pm$ is s.t. for every variable $x \in PS$, we do **not** have $M(x^+) = M(x^-) = 0$. Denoting $(BOOL^{PS^\pm})^*$ the subset of all interpretations I from $BOOL^{PS^\pm}$ s.t. for every $x \in PS$, we do not have $I(x^+) = I(x^-) = 0$, every preferred model M of $\Sigma^\pm$ belongs to $(BOOL^{PS^\pm})^*$.

On this ground, every interpretation I from $(BOOL^{PS^\pm})^*$ can be associated to the 3-interpretation $3(I)$ over PS defined by $3(I)(x) = I(x^+)$ and $3(I)(\neg x) = I(x^-)$ for every variable $x \in PS$. It is obvious that 3() is a bijection from both sets of interpretations. It is also obvious to show (by induction on the structure of $\alpha \in PROP_{PS}$) that I is a model of the monotone (positive) formula $\alpha^\pm$ from $PROP_{PS^\pm}$ iff $3(I)$ is a 3-model of α.

Finally, we know that $\{\Sigma\}{\approx}^{\mathcal{IP},\emptyset}_{PS}\gamma$ iff every preferred 3-model of Σ is a 3-model of γ, where the preferred 3-models of Σ are the minimal 3-models of it w.r.t. the partial ordering $\leq^3$ over 3-interpretations defined by $I \leq^3 J$ iff $\{x \in PS \mid I(x) = I(\neg x)\} \subseteq \{x \in PS \mid J(x) = J(\neg x)\}$. Since every preferred model of $\Sigma^\pm$ can be associated to a 3-model of Σ through the 3() bijection, and since this bijection is a morphism (i.e., for every pair of interpretations I and J from $(BOOL^{PS^\pm})^*$, we have $I \leq J$ iff $3(I) \leq^3 3(J)$), we obtain that M is a preferred model of $\Sigma^\pm$ iff $3(M)$ is a preferred model of Σ, which concludes the proof. ■

Proof of Proposition 3.10:
Because the variables $new_{i,j}$ can be viewed as names given to the formulas $\phi_{i,j}$ of B, we can define a bijective mapping $\flat$ from the set of classical interpretations over $PS \setminus \{new_{i,j} \mid \phi_{i,j} \in B\}$ to the set of 3-interpretations over PS s.t. every variable from $Var(B)$ is interpreted in a classical way, and every variable $new_{i,j}$ is interpreted as 1. Formally, let I be a classical interpretation over $Var(B)$; we define $\flat(I)$ by:

- When x is not a new variable, $\flat(I)(x) = I(x)$ and $\flat(I)(\neg x) = 1 - I(x)$.
- $\flat(I)(new_{i,j}) = 1$.
- $\flat(I)(\neg new_{i,j}) = 0$ when $I \models \phi_{i,j}$ holds, $\flat(I)(\neg new_{i,j}) = 1$ otherwise.

It is obvious that $\flat$ is bijective and that I and $\flat(I)$ give the same truth value to every formula that does not contain any new variable.

Moreover, I is a classical model of a subset B' of B iff $\flat(I)$ is a $S' - 3$ model of B_{new} where $S' = Var(B) \cup \{new_{i,j} \mid \phi_{i,j} \in B'\}$. Indeed, by construction, $\flat(I)$ satisfies every $new_{i,j}$. If I satisfies $\phi_{i,j}$, then its extension $\flat(I)$ satisfies $\phi_{i,j}$, hence it satisfies also $\neg new_{i,j} \vee \phi_{i,j}$. If I does not satisfy $\phi_{i,j}$, then $\flat(I)$ satisfies $\neg new_{i,j}$, hence it satisfies also $\neg new_{i,j} \vee \phi_{i,j}$. Accordingly, $\flat(I)$ is a 3-model of B_{new}. Clearly enough, every variable from $Var(B)$ is interpreted classically in $\flat(I)$. This is also the case for the new variables $new_{i,j}$ s.t. I satisfies $\phi_{i,j}$, hence for the new variables $new_{i,j}$ s.t. $\phi_{i,j} \in B$ since I is a classical model of B. Thus, $\flat(I)$ is a $S' - 3$-model of B_{new} where $S' = Var(B) \cup \{new_{i,j} \mid \phi_{i,j} \in B'\}$. Conversely, if $\flat(I)$ is a $S' - 3$ model of B_{new} where $S' \subseteq Var(B) \cup \{new_{i,j} \mid \phi_{i,j} \in B\}$, then its restriction I is a classical model of $B' = \{\phi_{i,j} \in B \mid new_{i,j} \in S'\}$. Indeed, since every variable $\neg new_{i,j}$ of B_{new} must be interpreted as 0 in $\flat(I)$ whenever $new_{i,j} \in S'$, there is only one way for $\flat(I)$ to satisfy $\neg new_{i,j} \vee \phi_{i,j}$ for such $new_{i,j}$: $\flat(I)$ must satisfy $\phi_{i,j}$. Because no $\phi_{i,j}$ of B contains a new variable and $\flat(I)$ is an extension of I to such new variables, it must be the case that I satisfies every $\phi_{i,j}$ for every $new_{i,j} \in S'$. Because no variable except the new ones is interpreted in a non-classical way, I is a classical model of $B' = \{\phi_{i,j} \in B \mid new_{i,j} \in S'\}$.

As an immediate consequence, we obtain that B' is a consistent subbase of B iff $S' = Var(B) \cup \{new_{i,j} \mid \phi_{i,j} \in B'\}$ is a consistent subset of S w.r.t. B_{new} and $Var(B)$. Because the total pre-order over the new variables given by S is isomorphic to the total pre-order over the corresponding explicit beliefs given by B, we have $B' \in B_{\mathcal{P}}$ iff $S' \in \mathcal{S}_{\mathcal{P}}(B_{new}, S, Var(B))$ with $S' = Var(B) \cup \{new_{i,j} \mid \phi_{i,j} \in B'\}$ for every selection policy $\mathcal{P}$ among $\mathcal{PO}$, $\mathcal{LO}$, $\mathcal{IP}$ and $\mathcal{LE}$.

Finally, provided that γ is a formula from $PROP_{PS}$ that does not contain any new variable, we have $B {\mid\!\sim}^{\mathcal{P}} \gamma$ iff every classical model I of a preferred subbase B' of $B_{\mathcal{P}}$ satisfies γ iff every $S' - 3$ model $\flat(I)$ of B_{new} satifies γ where $S' \in \mathcal{S}_{\mathcal{P}}(B_{new}, S, Var(B))$ iff $B_{new} \approx_S^{\mathcal{P}, Var(B)} \gamma$. ■

Proof of Proposition 3.11:
By definition, we have: $B \approx_S^{\mathcal{P}, S_0} \gamma$, iff $\forall S' \in \mathcal{S}_{\mathcal{P}}(B, S, S_0)$, $B \models_{S'}^3 \gamma$.

Now, if we rename all the literals of B such that x becomes a new variable x^+ and $\neg x$ becomes a new variable x^-, we can consider the base $B^{\pm}$ obtained and add formulas to recover our notion of interpretation and preferred models. First, we have to add formulas to express the fact that we must at least have x^+ or x^- so as to guarantee the notion of three-interpretation. Then, since variables of S_0 have to be classically interpreted, we add formulas expressing that we have at least $\neg x^+$ or $\neg x^-$ for all the corresponding variables x of S_0. The last thing we have to do is interpreting classically the maximum subsets of variables depending on the stratification of S and the selection policy $\mathcal{P}$. For each stratum of S, we build a new stratum in which each variable x of S is replacing by $\neg x^+ \vee \neg x^-$ and then add this stratum to the new belief base.

Thus, we obtain the following stratified belief base:
$(B^{\pm} \cup \bigcup_{x \in PS} \{x^+ \vee x^-\} \cup \bigcup_{x \in S_0} \{\neg x^+ \vee \neg x^-\}, \{\neg x^+ \vee \neg x^- \mid x \in S_1\}, \{\neg x^+ \vee \neg x^- \mid x \in S_j\})$ on which we use the standard coherence based approach with the same selection policy $\mathcal{P}$ such that we have: $B \approx_S^{\mathcal{P}, S_0} \gamma$ iff $(B^{\pm} \cup \bigcup_{x \in PS} \{x^+ \vee x^-\} \cup \bigcup_{x \in S_0} \{\neg x^+ \vee \neg x^-\}, \{\neg x^+ \vee \neg x^- \mid x \in S_1\}, \{\neg x^+ \vee \neg x^- \mid x \in S_j\}) {\mid\!\sim}^{\mathcal{P}} \gamma^{\pm}$. ■

Proof of Proposition 3.12:
Let B be a belief base from $PROP_{PS}^3$, $S = (S_0, \ldots, S_j)$ be a stratification of

a subset of PS, where S_0 is s.t. $B \not\models^3_{S_0} false$ and γ be a CNF formula from $PROP_{PS}$. Let us first show that checking whether $B \models^3_S \gamma$ given B, S and γ is in coNP. To do it, it is sufficient to consider the complementary problem, and to show that it can be solved in polynomial time by the following nondeterministic algorithm:
(1) guess an $S-3$ interpretation I (over the variables of $Var(B) \cup S \cup Var(\gamma)$),
(2) check in polynomial time that $I(B) > 0$ and $I(\gamma) = 0$.
As to the hardness result, it is sufficient to reduce UNSAT to LITERAL $\models^3_S$. This can be done easily by associating to every CNF formula Σ from $PROP_{PS}$ in polynomial time the formula Σ' obtained by replacing every literal l in Σ by the square literal $\Box l$. Now, Σ is unsatisfiable iff $\Sigma' \models^3_\emptyset new$ where $new \in PS \backslash Var(\Sigma)$, and this concludes the proof. ■

Proof of Proposition 3.13:
Since $\models^3_S$ clearly satisfies the right-and property, we can assume w.l.o.g. that γ is a clause. Then we extend the proof of Lemma A.2 from [58] to the language of J_3. The basic idea is to translate in polynomial time each instance of $B \overset{?}{\models^3_S} \gamma$ into an instance of $B^* \overset{?}{\models} \gamma^*$ where B^* and γ^* are classical propositional formulas, and the positive instances are preserved. Here B^* is obtained by replacing in any NNF of B each positive (resp. negative) occurrence of variable x in a non-square literal by a positive occurrence of variable x^+ (resp. x^-) and each positive (resp. negative) occurrence of a square literal $\Box l$ by a positive (resp. negative) occurrence of variable Nl. The resulting formula is then conjoined with (1) $\bigwedge_{x \in Var(B) \cup Var(\gamma)} (x^+ \vee x^-) \wedge \bigwedge_{x \in S} (\neg x^+ \vee \neg x^-)$, and finally conjoined with (2) $(\neg Nx \vee x^+) \wedge (\neg Nx \vee \neg x^-) \wedge (\neg N\neg x \vee x^-) \wedge (\neg N\neg x \vee \neg x^+)$ for each occurrence of a square literal $\Box x$ or $\Box\neg x$ in the NNF of B under consideration. γ^* is the clause obtained by replacing in γ each positive (resp. negative) occurrence of variable x by a positive occurrence of variable x^+ (resp. x^-). This is the same translation at work in Lemma A.2 from [58], except that the conjunct (2) is new. This conjunct encodes the connection between each square literal $\Box l$ occurring in B and the corresponding literal l. Thus, if $\Box x$ (encoded as Nx) is true, then x (encoded as x^+) must be true and $\neg x$ (encoded as x^-) must be false. Similarly, if $\Box\neg x$ (encoded as $N\neg x$) is true, then x (encoded as x^+) must be false and $\neg x$ (encoded as x^-) must be true. If x is true, then $\Box\neg x$ must be false, while if x is false, $\Box x$ must be false. Thus, we have $B \models^3_S \gamma$ iff $B^* \models \gamma^*$. Because every variable occurring in B^* or γ^* is pure (i.e., it occurs only positively or negatively), except those stemming from S and those of the form Nx and $N\neg x$, the conclusion easily follows (it is sufficient to perform case analysis on these variables, only). ■

Proof of Proposition 3.14:

- Membership. For each selection policy $\mathcal{P}$, membership is an easy consequence of the polynomial reduction to the corresponding inference problem in the standard coherence-based approach that is established by Proposition 3.11 and membership results that can be found in the literature (see Theorem 8 from [50] (or Corollary 1 from [16]), Theorem 6.5 from [51], Theorem 14 from [48], Theorem 20 from [49]). Indeed, all the complexity classes considered here are closed under polynomial (many-one) reduction.

- Hardness. For each selection policy $\mathcal{P}$ under consideration, hardness is a direct consequence of the polynomial reduction from the corresponding inference problem in the standard coherence-based approach that is established by Proposition 3.10, and hardness results that can be found in the literature [9] (see Theorem 8 from [50] (or Corollary 1 from [16]), Theorem 6.5 from [51], Theorem 14 from [48], Theorem 20 from [49]).

■

Proof of Proposition 3.15:
When $|S| \leq b$, there is only 2^b subsets of S, which is a constant when b is a constant. Accordingly, the naive algorithm that consists in generating $\mathcal{S}_{\mathcal{P}}(B, S, S_0)$ by filtering out the preferred consistent elements S' of 2^S, and testing for every S' whether or not $B \models^3_{S'} \gamma$ holds runs in time $\mathcal{O}(2^{2 \cdot b + 2 \cdot |Nec(B)|} \cdot |B| \cdot |\gamma|)$, hence in $\mathcal{O}(|B| \cdot |\gamma|)$ for a constant b and a bounded $|Nec(B)|$, for every B and a CNF query γ. ■

Acknowledgements

The authors would like to thank the reviewers for their helpful comments. The authors have been partly supported by the IUT de Lens, the Université d'Artois, the Région Nord / Pas-de-Calais under the TACT-TIC project, and by the European Community FEDER Program. Nadège Porquet is granted by the Région Nord / Pas-de-Calais.

References

[1] L. Amgoud and C. Cayrol. On the acceptability of arguments in preference-based argumentation frameworks. *Proc. UAI'98*, 1-7, 1998.
[2] O. Arieli and A. Avron. The value of four values. *Artificial Intelligence*, 102:97–141, 1998.
[3] C. Baral and S. Kraus and J. Minker. Combining multiple knowledge bases. *IEEE Transactions on Knowledge and Data Engineering* 3(2):208–220, 1991.
[4] N. Belnap. A useful four-valued logic. In *Modern Uses of Multiple-Valued Logic*, Reidel, 8-37, 1977.

[9] Actually, previous complexity results typically concern the FORMULA $\mid\!\sim^{\mathcal{P}}$ problem, i.e., queries γ are unrestricted propositional formulas. Nevertheless, it is easy to modify the corresponding hardness proofs to show that the complexity lower bounds are also the case for both the corresponding CLAUSE $\mid\!\sim^{\mathcal{P}}$ and LITERAL $\mid\!\sim^{\mathcal{P}}$ problems.

[5] S. Benferhat, C. Cayrol, D. Dubois, J. Lang and H. Prade. Inconsistency management and prioritized syntax-based entailment. *Proc. IJCAI'93*, 640-645, 1993.

[6] S. Benferhat, D. Dubois, and H. Prade. How to infer from inconsistent beliefs without revising? *Proc. IJCAI'95*, 1449-1455, 1995.

[7] S. Benferhat. Computing specificity in default reasoning. In *Handbook of DRUMS*, vol. 5, Kluwer Academic, 147-177, 2000.

[8] Ph. Besnard and A. Hunter. Quasi-classical logic: non-trivializable classical reasoning from inconsistent information. *Proc. ECSQARU'95*, LNAI 946, Springer-Verlag, 44-51, 1995.

[9] Ph. Besnard and A. Hunter. Introduction to actual and potential contradictions. In *Handbook of DRUMS*, vol. 2, Kluwer Academic, 1-11, 1998.

[10] Ph. Besnard and J. Lang. Graded paraconsistency - reasoning with inconsistent and uncertain knowledge. In *Frontiers of Paraconsistent Logics*, Studies in Logic and Computation, vol. 8, Research Studies Press, 75-94, 2000.

[11] Ph. Besnard and T. Schaub. Circumscribing inconsistency. *Proc. IJCAI'97*, 150-155, 1997.

[12] Ph. Besnard and T. Schaub. Signed systems for paraconsistent reasoning. *Journal of Automated Reasoning*, 20:191-213, 1998.

[13] A. Bondarenko, P. Dung, R.A. Kowalski and F. Toni. An abstract, argumentation-theoretic framework for default reasoning. *Artificial Intelligence*, 93(1-2):63-101, 1997.

[14] G. Brewka. Preferred subtheories: an extended logical framework for default reasoning. *Proc. IJCAI'89*, 1043–1048, 1989.

[15] M. Cadoli and M. Schaerf. On the complexity of entailment in propositional multivalued logics. *Annals of Mathematics and Artificial Intelligence*, 18:29–50, 1996.

[16] C. Cayrol and M.-C. Lagasquie-Schiex. On the complexity of non-monotonic entailment in syntax-based approaches. *Proc. ECAI'94 Workshop on Algorithms, Complexity and Commonsense Reasoning*, 1994.

[17] C. Cayrol and M.-C. Lagasquie-Schiex. Non-monotonic syntax-based entailment: a classification of consequence relations. *Proc. ECSQARU'95*, LNAI 946, Springer-Verlag, 107–114, 1995.

[18] C. Cayrol, M.-C. Lagasquie-Schiex and Th. Schiex. Nonmonotonic reasoning: from complexity to algorithms. *Annals of Mathematics and Artificial Intelligence*, 22(3-4):207–236, 1998.

[19] S. Chopra and R. Parikh. An inconsistency tolerant model for belief representation and belief revision. *Proc. IJCAI'99*, 192-197, 1999.

[20] S. Coste-Marquis and P. Marquis. Complexity results for paraconsistent inference relations. *Proc. KR'02*, 61-72, 2002.

[21] J.M. Crawford and D.W. Etherington. A non-deterministic semantics for tractable inference. *Proc. AAAI'98*, 1998.

[22] J.M. Crawford and B. Kuipers. Towards a theory of access-limited logic for knowledge representation. *Proc. KR'89*, 67–78, 1989.

[23] N. da Costa. On the theory of inconsistent formal systems. *Notre Dame Journal of Formal Logic*, 15:497-510, 1974.

[24] M. Dalal. Anytime families of tractable propositional reasoners. *Proc. AI&Math'96*, 42–45, 1996.

[25] M. Davis, G. Logemann and D. Loveland. A machine program for theorem proving. *Journal of the Association for Computing Machinery* 5:394-397, 1962.

[26] I.M.L. D'Ottaviano and N.C.A. da Costa. Sur un problème de Jaśkowski. *Comptes Rendus de l'Académie des Sciences de Paris* 270(A):1349–1353, 1970.

[27] O. Dubois, P. André, Y. Boufkhad and J. Carlier. SAT versus UNSAT. *Second DIMACS Challenge*, DIMACS Series in Discrete Mathematics and Theoretical Computer Science, American Mathematical Society, 415-436, 1996.

[28] P. Dung. On the acceptability of arguments and its fundamental role in nonmonotonic reasoning, logic programming and n-person games. *Artificial Intelligence* 77:321-357, 1995.
[29] M. Elvang-Goransson and A. Hunter. Argumentative logics: reasoning from classically inconsistent information. *Data and Knowledge Engineering* 16:125–145, 1995.
[30] R.L. Epstein. The semantic foundation of logic. Vol. I: propositional logics. *Kluwer Academic Publisher*, 1990.
[31] R. Fagin and J.D. Ullman and M.Y. Vardi. On the semantics of updates in databases. *Proc. PODS'83*, 352–355, 1993.
[32] M. Finger and R. Wassermann. Tableaux for approximate reasoning. *Proc. IJCAI'01 Workshop on Inconsistency in Data and Knowledge*, 2001.
[33] A.M. Frisch. Inference without chaining. *Proc. IJCAI'87*, 515–519, 1987.
[34] M. Ginsberg. Counterfactuals. *Artificial Intelligence* 30:35–79, 1986.
[35] M. Ginsberg. The computational value of nonmonotonic reasoning. *Proc. KR'91*, 262–268, 1991.
[36] C. Hewitt. Metacritique of McDermott and the logicist approach. *Computational Intelligence* 3, 1987.
[37] A. Hunter. Reasoning with contradictory information using quasi-classical logic. *Journal of Logic and Computation* 10(5), 2000.
[38] D. Kayser. Une étude du "raisonnement par approximation" de Schaerf et Cadoli (In French). *Revue d'intelligence artificielle*, vol. 11, 4:449–461, 1997.
[39] S. Konieczny and R. Pino Pérez. On the logic of merging. *Proc. KR'98*, 488-498, 1998.
[40] S. Konieczny. On the difference between merging knowledge bases and combining them. *Proc. KR'00*, 135-144, 2000.
[41] S. Kraus, D. Lehmann and M. Magidor. Nonmonotonic reasoning, preferential models and cumulative logics. *Artificial Intelligence* 44:167–207, 1990.
[42] H.J. Levesque. A logic of implicit and explicit belief. *Proc. AAAI'84*, 198–202, 1984.
[43] H.J. Levesque. Knowledge representation and reasoning. *Annual Reviews in Computer Science* 1:255–287, 1986.
[44] H.J. Levesque. A knowledge-level account of abduction. *Proc. IJCAI'89*, 1061-1067, 1989.
[45] F. Lin. Reasoning in the presence of inconsistency. *Proc. AAAI'87*, 139-143, 1987.
[46] J. Lin. Integration of weighted knowledge bases. *Artificial Intelligence*, 83(2):363-378, 1996.
[47] P. Marquis and N. Porquet. Computational aspects of quasi-classical entailment, *Journal of Applied Non-Classical Logics*, 2001 (to appear).
[48] B. Nebel. Belief revision and default reasoning: syntax-based approaches. *Proc. KR'89*, 301–311, 1989.
[49] B. Nebel. Syntax-based approaches to belief revision. In *Belief Revision*, P. Gärdenfors (ed.), Cambridge Tracts in Theoretical Computer Science, vol. 29, 52–88, 1992.
[50] B. Nebel. Base revision operations and schemes: semantics, representation and complexity. *Proc. ECAI'94*, 341–345, 1994.
[51] B. Nebel. How hard is it to revise a belief base? In *Handbook of DRUMS*, vol. 3, Kluwer Academic, 77-145, 1998.
[52] Ch. H. Papadimitriou. *Computational Complexity*, Addison-Wesley, 1994.
[53] G. Pinkas and R.P. Loui. Reasoning from inconsistency: a taxonomy of principles for resolving conflict. *Proc. KR'92*, 709-719, 1992.
[54] G. Priest. Reasoning about truth. *Artificial Intelligence* 39:231-244, 1989.
[55] G. Priest. Minimally inconsistent LP. *Studia Logica* 50:321-331, 1991.
[56] N. Rescher and R. Manor. On inference from inconsistent premises. *Theory and Decision* 1:179–219, 1970.
[57] P.Z. Revesz. On the semantics of arbitration. *International Journal of Algebra and Computation* 7(2):133-160, 1997.

[58] M. Schaerf and M. Cadoli. Tractable reasoning via approximation. *Artificial Intelligence* 74(2):249-310, 1995.

[59] B. Selman and H.A. Kautz. Knowledge compilation and theory approximation. *Journal of the Association for Computing Machinery* 43(2):193-224, 1996.

[60] A. ten Teije and F. van Harmelen. Computing approximate diagnoses by using approximate entailment. *Proc. KR'96*, 256–265, 1996.

[61] R. Wassermann. On structured belief bases. In H.Rott and M.Williams (eds.) *Frontiers in Belief Revision*, Applied Logic Series, vol. 22, Kluwer Academic Publisher, 2001.