On the complexity of paraconsistent inference relations

Sylvie Coste-Marquis and Pierre Marquis *

CRIL-CNRS/Université d'Artois
rue de l'Université - S.P. 16
F-62307 Lens Cedex - France
{coste,marquis}@cril.univ-artois.fr

Abstract. Reasoning in a non-trivial way from inconsistent pieces of information is a major challenge in artificial intelligence, and its importance is reflected by the number of techniques designed so far for dealing with inconsistency (especially the few ones reported in this handbook). Many of these techniques have been investigated in depth from a logical point of view, but far less to what concerns the computational complexity aspects. The purpose of this chapter is to present in a structured way the main complexity results identified so far for paraconsistent inference based on multi-valued propositional logics.

1 Introduction

Classical logic cannot be used as such to deal with inconsistent information since every formula can be derived from a contradiction. This well-known trivialization problem leads to consider that an inconsistent information base is totally meaningless: Since all contradictory formulas are equivalent and since a replacement (meta)theorem holds in classical logic, every contradictory base can be substituted by the boolean constant *false*, i.e., the unique irreducible contradiction. In order to escape from such a problem and offer a way to exploit inconsistent information bases in a cleverer way, *paraconsistent* inference relations must be considered. Using such inference relations, every formula is not a consequence anymore of every (classically) inconsistent formula.

The trivialization of classical entailment in presence of inconsistency appears as an important issue in AI for several reasons. A major reason is that it is actually a concrete problem: Like the beliefs of human beings in everyday life, the beliefs of artificial agents can easily be jointly inconsistent, especially when their databases are large or when they integrate information stemming from several sources (sensors, other agents, etc.). The problem is all the most salient when societies of agents are considered instead of individuals: Due to their different abilities and evolutions, agents have often conflicting beliefs about some topics,

* Many thanks to the anonymous reviewers of this chapter for many interesting comments. Many thanks as well to the Région Nord / Pas-de-Calais, the IUT de Lens and the IRCICA Consortium for their support.

even if they share some consistent pieces of knowledge. Would it be rational to consider the common belief of such a society as meaningless? Clearly, intuition gives a negative answer: Since inconsistencies are typically local (i.e., they are due to proper subparts of the information base), it would be definitely unreasonable to accept that everything (or nothing) follows from the available information base whenever this information base is inconsistent.

Assume for instance that a given agent has the following beliefs, encoded as a propositional formula Σ (it will serve as a running example throughout the chapter):

Example 1. $\Sigma = a \wedge \neg a \wedge (a \vee b) \wedge c \wedge (\neg c \vee d)$
The agent has contradictory beliefs about the truth value of a (for instance, because she received conflicting evidence about it, she has been told by a reliable acquaintance that a holds, and by a second as reliable acquaintance that a does not hold). She also has fully consistent pieces of beliefs, $a \vee b$, c and $\neg c \vee d$. Because Σ is inconsistent, every formula is a logical consequence of Σ. Nevertheless, because the unique contradiction in Σ concerns only a and $\neg a$, it does not seem always relevant to conclude that the agent has trivial beliefs. In particular, since Σ does not tell anything about atom e, an expected conclusion could be that e is undetermined (i.e., contingent) given Σ in the sense that neither e nor $\neg e$ are consequences of Σ. In addition, since c and $\neg c \vee d$ do not participate to any contradiction, it is not absurd to desire viewing them as consequences of Σ as well.

A second factor explaining the prominence of the trivialization problem in AI is the close connection it has with the problem of *handling exceptions*, another main challenge in AI. Indeed, an exception to a rule is just a case which contradicts the rule. As a matter of fact, many interesting approaches to inconsistency tolerant reasoning can be used to address the exception handling issue as well (as we will see, most valuable paraconsistent inference relations are non-monotonic ones), and the connections between belief revision (in the AGM framework) and rational inference relations are known for a while [1, 2]. Nevertheless, it must be noted that paraconsistency and non-monotony are two logically independent notions: One can find inference relations that are paraconsistent and monotonic (e.g., inference in LP, see Section 3.2), one can find as well inference relations that are paraconsistent and non-monotonic (e.g., inference in LP_m, see Section 3.3) while a third (non-empty) family gathers inference relations that are not paraconsistent and monotonic (e.g., classical entailment) and a last family consists of inference relations that are neither paraconsistent nor monotonic (e.g., circumscription).

In the following, we assume that the information at hand are *beliefs*, and the main goal is the *inference* one: Making explicit some pieces of beliefs hidden in implicit ones. This issue contrasts with the problem of aggregating mutually inconsistent goals (or preferences or desires). The latter problem is not less concrete than the inference problem from inconsistent beliefs — on the contrary, it is at the core of both multicriteria decision making and group decision making

— and connections exist between the two issues (e.g., belief merging operators can also be used to goal fusing in some situations). However, the main objectives are different ones: On the one hand, it consists in characterizing consequences of a belief base; on the other hand, it consists in determining a global preference relation over goals.

The significance of handling inconsistent beliefs in AI is reflected by the many approaches developed so far to address the paraconsistency issue under its various forms. Let us just mention paraconsistent logics, belief revision, belief merging, reasoning from preferred consistent subsets, knowledge integration, argumentative logics, purification, etc. All these words actually name many different techniques described in an abundant literature (see [3–5] for a survey).

The large number of such techniques can be explained by the fact that paraconsistency can be achieved in various ways, depending on the exact nature of the problem at hand (hence, the available information at the start). Thus, when Σ represents the (conflicting) beliefs of several agents, a merged base giving the beliefs of the group of agents can be designed by logically weakening some local belief bases (associated to the agents) in order to restore global consistency [6–10]. Several weakening mechanisms can be used: Dilation [11] is at work when merging operators based on a selection of models are considered (a i-dilation of a belief base ϕ for a "distance" d is obtained by adding to the set of its models all these interpretations which are at a distance at most $i \in I\!N$ from a model of ϕ), formula inhibition is used within merging operators based on a selection of formulas. Thus, when Σ is an inconsistent set of formulas, focusing on the consistent subsets – eventually restricted to the most preferred ones, when some preferential information can be exploited – is sufficient to give rise to several families of paraconsistent inference relations (depending on the entailment principle under consideration, e.g., the skeptical one) [12–18]. Such techniques are closely related to so-called syntax-based approaches to belief revision [19, 20] and to the framework for supernormal default reasoning with priorities from [21]. Variable forgetting can be used as well as a weakening mechanism [22]. Another way to give rise to paraconsistent inference relations consists in associating to each inferred formula a justification under the form of a subset of Σ used to derive it, and by reasoning on such arguments whenever some mutually inconsistent formulas can be derived. This is the basic idea of argumentative logics, see e.g., [23–25]. It can be sophisticated through the use of preferential information in order to consider some arguments more acceptable than some other arguments [26].

Now, there are several definitions of what a paraconsistent logic is; in this chapter, the following (quite restricted) definition is considered: A propositional logic $\langle L, \vdash_L \rangle$ where L is a propositional language and $\vdash_L$ a binary relation over L is said to be paraconsistent whenever there exists a belief base $\Sigma \in L$ and a query $\gamma \in L$ s.t. Σ is classically inconsistent and $\Sigma \not\vdash_L \gamma$ nevertheless. Thus, within a paraconsistent logic, the trivialization issue of inference is addressed in some situations when the input consists of a *single inconsistent formula* (representing the belief base) and a query.

Compared with the other approaches listed above, paraconsistent logics (*stricto sensu*) give *more basic* ways to address the trivialization issue. Indeed, belief revision, belief merging, knowledge integration, reasoning from preferred inconsistent subsets and purification need some extra-logical information in order to be well-defined and avoid trivializing. Such extra-logical information can be rather poor (a splitting between the belief base and the revision formula in the belief revision setting, a set (or multi-set) organization of the elementary (or agents) beliefs in a belief merging scenario) or rather sophisticated (preference relations over the beliefs, knowledge gathering actions for purification) but they are required. Thus, none of the approaches listed above will be considered further in this chapter.[1]

Several (non mutually exclusive) techniques can be used to define an inference relation that avoid trivialization from an inconsistent formula (see [5] for a brilliant survey of them). One of them consists in *making the truth value of* α *(more or less) independent from the truth value of* $\neg\alpha$ (e.g., this is what is done in C_ω logic [28] and the whole family of C_i logics). A second one consists in *restricting the proof theory* of classical logic so as to retain only a subset of the classical proofs as admissible (this is at work in quasi-classical logic [29, 30]). A third one consists in *preventing inconsistent belief bases from having no model*, through the consideration of more general notions of interpretations. Several *multi-valued logics* are related to this line of research (among others, see [31–42]).

Just like belief revision or belief merging gather many operators from which many inference relations are induced, there is no universal consensus on what paraconsistent inference should be but several intuitions of what is expected. While paraconsistent inference relations designed so far share some common principles (for instance, any irrelevant piece of belief would not be considered as a consequence of some inconsistent belief bases – stepping back to the previous example, e would never be considered as a conclusion from Σ), they do not give the same closure (set of consequences) for all inconsistent belief bases Σ. On the previous example, some approaches lead to consider b as a conclusion of Σ (given that the conjunction Σ has both $\neg a$ and $a \vee b$ among its conjuncts) while others reject such a conclusion (given the logically strongest conjunct a of Σ, $a \vee b$ is considered irrelevant). Both approaches look reasonable in some specific contexts (thus, it can be the case that $a \vee b$ is an explicit belief whose origin differs from the one of a and that the agent wants to keep it but the opposite conclusion also makes sense in other situations). This tends to show that a purely descriptive approach to paraconsistency (i.e., modeling paraconsistent inference by considering how human beings reason in presence of inconsistency) is hard to be achieved when a so simple input (an inconsistent formula, nothing else) is considered.

[1] The definition of paraconsistency we used also sets aside several logics that are usually considered as paraconsistent logics, like Jaśkowski's discussive logic [27], one of the first logic defined so far to address the inconsistency handling issue.

This raises the issue of how to evaluate and compare paraconsistent inference relations. Several independent criteria can be used for such a purpose. A first criterion consists in investigating the logical properties of the relations. Many normative postulates which should be satisfied by common-sense inference relations have been proposed so far. For instance, it is interesting to determine whether the paraconsistent inference relation under consideration $\vdash_L$ satisfies reflexivity, transitivity, ... It is also valuable to determine whether $\vdash_L$ is subclassical (i.e., a subset of classical entailment) whenever the fragment of L from which both the belief bases and the queries are built up can be viewed as well as a fragment of classical propositional logic. It is interesting to determine whether usual metatheorems hold (substitution theorem, deduction theorem, compactness theorem, ...) and whether sound and complete axiomatisations (à la Hilbert or à la Gentzen) have been exhibited. All those logical features give a first dimension for a comparison. In the following, we will not focus on them but let the reader look at the bibliography for more about the logical aspects of the inference relations we consider (the logical dimension of paraconsistent relations is already presented with much details in several papers from the literature, in particular [5]).

Another criterion is the cautiousness one; it is important to consider this dimension together with the logical one since it is quite easy to design strongly rational inference relations which are useless just because they are too cautious (they do not lead to consider as consequences of a base many ones which are expected).

In this chapter, we mainly consider a third criterion, orthogonal to the previous ones but very relevant from an AI perspective: *The computational complexity one.* Indeed, in order to be able to simulate a paraconsistent inference relation on a computer, the computational resources (time and space) required must be taken into account. Complexity theory offers theoretical tools (especially, complexity classes and reductions) for a fine-grained classification of the computational difficulties. Particularly, it helps in identifying the complexity sources within a problem; polynomial reductions among problems are useful when they show the existence of a translation from the decision problem associated to a given paraconsistent relation to the decision problem associated to another one; this is particularly significant when sophisticated algorithms have been designed for the latter: It is then possible to take advantage of them for solving the former decision problems, avoiding thus an expensive software development (furthermore, if there is no guarantee that it is computationally valuable to do so from the practical side, it can actually be the case). Anyway, identifying complexity results for paraconsistent inference relations gives a third dimension allowing the comparison of these relations. From an AI point of view, it can make sense to prefer a more tractable inference relation over a less tractable one, even if the latter exhibits a more rational behavior than the former.

Quite surprisingly, the complexity of paraconsistent inference relations has been investigated only for the past few years. This contrasts with many other approaches to inconsistency tolerant reasoning for which many complexity results have been obtained; indeed, complexity results for belief revision can be found in [43, 44], complexity results for belief merging in [45], complexity of reasoning from preferred consistent subsets in [44, 18, 46], complexity of argumentative reasoning in [47–49].

It is also amazing to compare the (few) number of papers dedicated to the study of the complexity of paraconsistent inference relations to the (quite large) number of papers where such inference relations have been pointed out. Subsequently, we do not claim that this chapter covers exhaustively the complexity results for all the paraconsistent inference relations from the literature:[2] Much work remains to be done at that time to reach such a goal. The present chapter can be considered as a first step toward this objective. The focus is laid on paraconsistent inference relations based on propositional multi-valued logics. Especially, we do not consider first-order paraconsistent logics in the following (the extension to the first-order case does not raise any technical difficulty in general and, as expected, the corresponding inference relations are typically not recursive, see [5] for details). Furthermore, we refrain ourselves to presenting the various proof systems (typically, tableau-based) that have been designed so far for the inference relations considered in this chapter, unless the existence of such systems gives a reduction from which a membership result or a hardness one can be derived (again, we let the reader look at the bibliography to learn more about proof systems for paraconsistent inference relations).

The rest of this chapter is organized as follows. After some formal preliminaries (Section 2), we present in Section 3 many paraconsistent inference relations based on multi-valued logics; we focus on relations introduced in [31–36, 51, 29, 37, 38, 4, 39, 41, 42] for which complexity results have been identified in [34, 52–55, 41, 42]. Each relation is illustrated on the running example and its main logical properties are given. Then Section 4 presents complexity results for such relations in the general case and in some restricted cases. Finally, we conclude the chapter in Section 5.

2 Formal preliminaries

In this section, we briefly present the basic notions about propositional logic and complexity theory necessary to understand the rest of the chapter. More details about such notions can be found e.g., in [56] and [57, 58].

2.1 Propositional logic

Given a denumerable set PS of propositional symbols, $PROP^2_{PS}$ denotes the propositional language built up from PS, the boolean constants *true* and *false*, and the connectives $\neg$, $\vee$, $\wedge$ in the standard way. $\Rightarrow$ can also be introduced as a connective, where $\alpha \Rightarrow \beta$ is a short for $(\neg\alpha) \vee \beta$. The elements of $PROP^2_{PS}$ are

[2] For instance, we do not consider here inference relations based on multimodal logics for representing incoherent beliefs [50].

called formulas. $Var(\Sigma)$ denotes the set of propositional symbols occurring in the formula Σ. The size of a formula Σ noted $|\Sigma|$, is the number of occurrences of symbols and connectives used to write it.

For every subset V of PS, L_V is the set of literals built up from the propositional symbols of V. A negative literal is a literal of the form $\neg x$, where $x \in PS$. A symbol from PS is also called a positive literal. The complementary literal of a positive literal $l = x$ (resp. a negative literal $l = \neg x$) is $\bar{l} = \neg x$ (resp. $\bar{l} = x$). Every finite disjunction of literals is called a clause and every finite conjunction of literals is called a term. A positive clause contains only positive literals. A CNF formula is a (finite) conjunction of clauses, also viewed as a set of clauses when it is convenient. A k-CNF formula is a CNF formula in which every clause contains no more than k literals (where k is a non-negative integer). A Krom formula is a CNF formula in which every clause contains at most two literals. A Horn (resp. reverse Horn) CNF formula is a CNF formula in which each clause contains at most one positive (resp. negative) literal. A formula Σ is renamable Horn CNF if and only if there exists a substitution σ from L_{PS} to L_{PS} s.t. $\sigma(l) = l$ for every literal l of L_{PS} except those of a set L, and for every literal l of L, $\sigma(l) = \bar{l}$ and $\sigma(\bar{l}) = l$, and $\sigma(\Sigma)$ is a Horn CNF formula. A belief base is a finite, conjunctively-interpreted, set of formulas from $PROP_{PS}^2$. A DNF formula is a finite disjunction of terms. An NNF (Negation Normal Form) formula is a formula from $PROP_{PS}^2$ s.t. the scope of any occurrence of $\neg$ in the formula is a propositional symbol.

Formulas are interpreted in the classical way. An interpretation over $PROP_{PS}^2$ is a mapping I which associates every propositional symbol to one of the two truth values of $TWO = \{0, 1\}$. The semantics of a formula Σ is a boolean function: A truth value can be associated to Σ as long as an interpretation I is considered. $I(\Sigma)$ denotes the truth value taken by Σ within I; it is defined in the usual compositional way. When I is s.t. $I(\Sigma) = 1$, I is said to be a model of Σ, noted $I \models^2 \Sigma$; otherwise, I is a counter-model of Σ. When a formula has no model (resp. no counter-model), it is said to be unsatisfiable or inconsistent (resp. valid). The binary relation $\models^2$ over $PROP_{PS}^2$ is defined by $\Sigma \models^2 \gamma$ if and only if every model of Σ is a model of γ. $\models$ is referred to as logical entailment: Whenever $\Sigma \models^2 \gamma$ holds, γ is said to be a logical consequence of Σ. Whenever $\Sigma \models^2 \Phi$ and $\Phi \models^2 \Sigma$ both hold, Σ and Φ are said to be logically equivalent, noted $\Sigma \equiv^2 \Phi$.

2.2 Computational complexity

A decision problem (encoded as a language) belongs to P (resp. NP) if and only if there exists a deterministic (resp. non-deterministic) Turing machine which can classify every instance of it in a number of computational steps polynomially bounded in the input size. The decision problems of P are usually considered as efficiently solvable.

Because a deterministic Turing machine can be considered as a non-deterministic one, the inclusion $\mathsf{P} \subseteq \mathsf{NP}$ is established. However, the converse is the famous open problem: $\mathsf{P} \stackrel{?}{=} \mathsf{NP}$ (which is conjectured false). Among all the problems in NP, the hardest ones are those from which every problem in NP can be polynomially many-one reduced: Such problems are referred to as NP-complete. If any of them has a polynomial (deterministic) algorithm, then $\mathsf{P} = \mathsf{NP}$ holds. Accordingly, it is believed that it is impossible to solve NP-complete problems in deterministic polynomial time. SAT, the problem of determining whether a propositional formula in CNF is satisfiable, is the prototypical NP-complete problem. Its complementary problem UNSAT (which consists in determining whether a propositional formula in CNF is unsatisfiable) is not necessarily in NP (in contrast to P, NP is not known to be closed under complementation). UNSAT is assigned to the class coNP which contains the complementary problems to problems in NP. It is conjectured that $\mathsf{NP} \neq \mathsf{coNP}$.

Let X be a class of decision problems. P^X (resp. NP^X) is the class of all decision problems that can be solved in polynomial time using a deterministic (resp. non-deterministic) Turing machine which can use an oracle for deciding the membership to X for "free" (i.e., within a constant, unit time). The classes Δ^p_k, Σ^p_k and Π^p_k (with $k \in \mathbb{N}$) can be defined by:

- $\Delta^p_0 = \Sigma^p_0 = \Pi^p_0 = \mathsf{P}$,
- $\Delta^p_{k+1} = \mathsf{P}^{\Sigma^p_k}$,
- $\Sigma^p_{k+1} = \mathsf{NP}^{\Sigma^p_k}$,
- $\Pi^p_{k+1} = \mathsf{co}\Sigma^p_{k+1}$.

Thus, $\Sigma^p_1 = \mathsf{NP}$ and $\Pi^p_1 = \mathsf{coNP}$. For any integer k, the complexity class D^p_k contains every decision problem that belongs to the intersection of a language from Σ^p_k and a language from Π^p_k. The *polynomial hierarchy* PH is the union of all Σ^p_k (for k integer). A decision problem is said to be at the k^{th} level of the polynomial hierarchy if and only if it belongs to Δ^p_{k+1}, and is either Σ^p_k-hard or Π^p_k-hard. While it is easy to check that $\Delta^p_k \subseteq \Sigma^p_k$, $\Delta^p_k \subseteq \Pi^p_k$, $\Sigma^p_k \subseteq D^p_k$, $\Pi^p_k \subseteq D^p_k$, $D^p_k \subseteq \Delta^p_{k+1}$ hold for every k, it is unknown whether the inclusions are proper (but it is strongly conjectured that they are). Indeed, it is strongly believed that the polynomial hierarchy does not collapse (at any level), i.e., is a truly infinite hierarchy (for every integer k, $\mathsf{PH} \neq \Sigma^p_k$).

3 Multi-Valued logics and related frameworks

In this section, we successively describe the syntax and the semantics of many paraconsistent multi-valued logics introduced so far, and some related propositional systems [29–39, 41, 42]. On this ground, we present a dozen of paraconsistent inference relations (both monotonic ones and non-monotonic ones). Their complexity will be given and discussed in the following section.

3.1 Syntax and semantics

The logics considered in this chapter are typically based on the following propositional languages, or on proper fragments of them:

Definition 1 **($PROP^4_{PS}$ $PROP^3_{PS}$ $PROP^2_{PS}$).**
Let PS be a finite set of propositional symbols.

- *$PROP^4_{PS}$ is the propositional language over PS inductively generated from the constant symbols true, false, both, and unknown and the connectives $\neg$, $\vee$, $\wedge$, $\supset$, $\oplus$, and $\otimes$.*
- *$PROP^3_{PS}$ is the propositional language over PS inductively generated from the constant symbols true, false and both and the connectives $\neg$, $\vee$, $\wedge$, $\supset$, and $\oplus$.*
- *$PROP^2_{PS}$ is the propositional language over PS inductively generated from the constant symbols true and false and the connectives $\neg$, $\vee$, $\wedge$, and $\supset$.*

As we will see, the constant symbols *both*, *unknown*, the consensus connective $\otimes$ and the gullability connective $\oplus$ are not interpreted in a classical way, using only the two standard "truth values" 0 (falsum) and 1 (verum). This explains why they are not used to generate formulas from $PROP^2_{PS}$, which coincides with a standard language for classical propositional logic. We can observe from the definition that $PROP^2_{PS}$ is a proper subset of $PROP^3_{PS}$ and that $PROP^3_{PS}$ is a proper subset of $PROP^4_{PS}$; none of the constants *both*, *unknown* (resp. *unknown*) and the connectives $\oplus$, $\otimes$ (resp. $\otimes$) can be used to generate formulas from $PROP^2_{PS}$ (resp. $PROP^3_{PS}$) because the set of "truth values" under consideration must be closed under the connectives.

Proper fragments of $PROP^2_{PS}$ consist of the set of formulas generated from PS and the three connectives $\neg$, $\vee$, $\wedge$ (the so-called $\{\neg, \vee, \wedge\}$ fragment also referred to as the monotonic fragment), and the subset of it containing all CNF formulas.

We must now explain how the formulas of $PROP^4_{PS}$ and its subsets can be interpreted. Obviously, a key feature of multi-valued logics from the semantics point of view is the presence of non-standard "truth values", i.e., different from the classical ones (0, denoting falsity and 1 denoting truth). The two additional "truth values" considered in such a setting are denoted $\perp$ and $\top$. Intuitively, $\perp$ denotes lack of information, while $\top$ indicates inconsistency. The latter "truth value" $\top$ is central when the goal is to achieve paraconsistency.

Definition 2 **(Interpretations).**
A 4-interpretation (resp. a 3-interpretation, a 2-interpretation) over PS is a total function I from PS to $FOUR = \{0, 1, \top, \perp\}$ (resp. $THREE = \{0, 1, \top\}$, $TWO = \{0, 1\}$).

Some alternative definitions of the set of "truth values" are sometimes given; for instance, $FOUR$ can be defined as $TWO \times TWO$, or as the power set of TWO. All these definitions are equivalent to the one given here, the intuition is that $\perp$ means "neither true nor false" while $\top$ means "both true and false".

Whatever the definition, it is useful to associate $FOUR$ with two (partial) orderings, the usual *truth ordering* $\leq_t$ for which 0 is the least element, 1 is the greatest element and $\top$ and $\perp$ are incomparable, and the *knowledge ordering* $\leq_k$ for which $\perp$ is the least element, $\top$ is the greatest element and 0 and 1 are

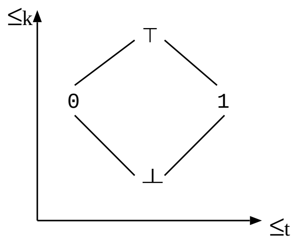

Fig. 1. FOUR.

incomparable. $FOUR$ with $\leq_t$ and $\leq_k$ is a bilattice, called Belnap's billatice. It is often represented by the double Hasse diagram given on Figure 1.

For every 4-interpretation I over PS, we define $I(true) = 1$, $I(false) = 0$, $I(both) = \top$, $I(unknown) = \bot$. Thus, the non-classical constants are mainly used to denote the corresponding non-classical "truth values" within the object language. Table 1 gives the truth table of the remaining connectives.

α	β	$\neg\alpha$	$\alpha \wedge \beta$	$\alpha \vee \beta$	$\alpha \supset \beta$	$\alpha \otimes \beta$	$\alpha \oplus \beta$
0	0	1	0	0	1	0	0
0	1	1	0	1	1	$\bot$	$\top$
0	$\top$	1	0	$\top$	1	0	$\top$
0	$\bot$	1	0	$\bot$	1	$\bot$	0
1	0	0	0	1	0	$\bot$	$\top$
1	1	0	1	1	1	1	1
1	$\top$	0	$\top$	1	$\top$	1	$\top$
1	$\bot$	0	$\bot$	1	$\bot$	$\bot$	1
$\top$	0	$\top$	0	$\top$	0	0	$\top$
$\top$	1	$\top$	$\top$	1	1	1	$\top$
$\top$	$\top$	$\top$	$\top$	$\top$	$\top$	$\top$	$\top$
$\top$	$\bot$	$\top$	0	1	$\bot$	$\bot$	$\top$
$\bot$	0	$\bot$	0	$\bot$	1	$\bot$	0
$\bot$	1	$\bot$	$\bot$	$\bot$	1	$\bot$	1
$\bot$	$\top$	$\bot$	0	1	1	$\bot$	$\top$
$\bot$	$\bot$	$\bot$	$\bot$	$\bot$	1	$\bot$	$\bot$

Table 1. Truth tables.

It can be observed that the semantics of the gullability connective $\oplus$ (resp. the consensus connective $\otimes$) is the supremum (resp. infimum) w.r.t. $\leq_k$; thus, $\oplus$ (resp. $\otimes$) plays a role similar to the one of the classical connective $\vee$ (resp. $\wedge$) but w.r.t. the knowledge ordering (and not the truth one).

Now, the semantics $I(\phi)$ of a formula ϕ from $PROP^4_{PS}$ in I is defined compositionally in the usual way (every connective is truth functional).

The semantics of a formula from $PROP^3_{PS}$ (resp. $PROP^2_{PS}$) in a 3-interpretation (resp. a 2-interpretation) is obtained by considering the reductions of the previous truth tables to $THREE$ (resp. TWO).

A further step to obtain inference relations consists in defining the set of *designated values.* Unlike many other multi-valued logics, the set considered in $FOUR$ and its restriction $THREE$ is $\{1, \top\}$. Accordingly, a formula ϕ is satisfied in a world I when its semantics in I is at least 1 w.r.t. the knowledge ordering.

Once this is stated, the following notions of models can be defined:

Definition 3 (Models).
*Let I be a 4-interpretation (resp. a 3-interpretation, a 2-interpretation) I over PS and ϕ be a formula from $PROP^4_{PS}$ (resp. $PROP^3_{PS}$, $PROP^2_{PS}$). I is a 4-*model *(resp. a 3-*model*, a 2-*model*) of ϕ if and only if $I(\phi) \in \{1, \top\}$. We note 4-mod(Σ) (resp. 3-mod(Σ)) the set of all 4-models (resp. 3-models) of a formula Σ from $PROP^4_{PS}$ (resp. $PROP^3_{PS}$).*

Now, equivalence can be defined: Two formulas ϕ and ψ from $PROP^4_{PS}$ (resp. $PROP^3_{PS}$) are said to be equivalent, noted $\phi \equiv^4 \psi$ (resp. $\phi \equiv^3 \psi$) if and only if they have the same set of 4-models (resp. 3-models). Note that a stronger notion of equivalence exists (two formulas are strongly equivalent when their semantics coincide in every interpretation) and that a replacement metatheorem holds for the strong notion but not for the weak one in the general case. For instance, while $a \vee \neg a \equiv^3 b \vee \neg b$ holds, it is not the case that $\neg(a \vee \neg a) \equiv^3 \neg(b \vee \neg b)$ ($a \vee \neg a$ and $b \vee \neg b$ are not strongly equivalent in $THREE$).

Because $\top$ is a designated value, it can be the case that a formula from $PROP^2_{PS}$ that is classically inconsistent has some 4-models (and this is what was expected in the purpose of avoiding trivialization). For instance:

Example 2. Let $\Sigma = a \wedge \neg a \wedge (a \vee b) \wedge c \vee (\neg c \wedge d)$. Σ has no 2-model. The 4-interpretation given by $I(a) = \top$, $I(b) = 1$, $I(c) = 1$ and $I(d) = 1$ is a 4-model of Σ. This interpretation can also be considered as a 3-model of Σ.

Contrariwise to what this example may suggest, it is not the case that inconsistency is always avoided in $FOUR$ or in $THREE$; some formulas in $PROP^4_{PS}$ (resp. $PROP^3_{PS}$) have no 4-models (the simplest one reduces to the constant *false*). Hence, only a weak form of paraconsistency is achieved in the general case (some inconsistent theories are not trivial ones, but not all of them). Nevertheless, it is easy to prove (by structural induction) that the interpretation mapping every symbol from PS to $\top$ is a 3-model (hence, a 4-model) of every formula from the $\{\neg, \vee, \wedge\}$ fragment.

Other connectives can be considered in the languages $PROP^4_{PS}$ and $PROP^3_{PS}$, as syntactic sugars. In particular:

- $\phi \Leftrightarrow \psi =_{def} (\phi \supset \psi) \wedge (\psi \supset \phi)$;
- $\phi \rightarrow \psi =_{def} (\phi \supset \psi) \wedge (\neg\psi \supset \neg\phi)$;
- $\phi \leftrightarrow \psi =_{def} (\phi \rightarrow \psi) \wedge (\psi \rightarrow \phi)$;
- $\phi! =_{def} \phi \wedge \neg\phi$;

– $\Box\phi =_{def} (\neg\phi \supset false) \wedge \neg(\phi \supset \neg\phi)$;
– $\phi \leq \psi =_{def} (\Box\phi \wedge \Box\psi) \vee (\Box\neg\phi \wedge \Box\neg\psi) \vee (\neg\Box\psi \wedge \neg\Box\neg\psi)$;
– $\diamond\phi =_{def} \neg\Box\neg\phi$;
– $\sim \phi =_{def} \Box\neg\phi$;
– $\odot\phi =_{def} \Box\phi \vee \Box\neg\phi$.

The language considered in [32] is the restriction of $PROP^2_{PS}$ where $\supset$ is not nestable. The language used in [39] to define $FOUR$ is the extension of $PROP^4_{PS}$ where $\rightarrow$ and $\leftrightarrow$ are used as additional (binary) connectives. The language of the logic of 3-inference of [34] (resp. the language of LP [35] or equivalently the language of RP [33]) is the propositional language over PS generated from the constant *false* and the connectives $\neg$, $\wedge$, $\vee$ (resp. from the connectives $\neg$, $\wedge$, $\vee$); $\supset$ and $\Leftrightarrow$ are also introduced but as syntactic sugars ($\phi \supset \psi =_{def} \neg\phi \vee \psi$ and $\phi \Leftrightarrow \psi =_{def} (\phi \supset \psi) \wedge (\psi \supset \phi)$). The language of LP_m [35,36] is the propositional language over PS generated from the connectives $\neg$, $\wedge$, $\vee$ and !.[3] The language of J_3 is the propositional language over PS generated from the connectives $\neg$, $\wedge$, $\vee$, $\supset$, $\Leftrightarrow$, $\Box$, $\diamond$, $\sim$, $\odot$. The language considered in [37] is the propositional language over PS generated from the constants *true*, *false* and the connectives $\neg$, $\wedge$, $\vee$, $\supset$, $\Leftrightarrow$, $\Box$, $\leq$. The language considered in [38] is the propositional language over PS generated from the connectives $\neg$, $\wedge$, $\vee$, $\supset$ and $\Leftrightarrow$.

The semantics of all such derived connectives can be easily defined by considering their truth tables. For instance, $\Box\phi$ means that ϕ is *necessarily* true, i.e., for every 4-interpretation I over PS, $I(\Box\phi) = 1$ if $I(\phi) = 1$ and $I(\Box\phi) = 0$ otherwise. All these additional connectives can be directly incorporated in $PROP^3_{PS}$ and $PROP^2_{PS}$ since both $THREE$ and TWO are closed w.r.t. any of them. When the truth values of TWO are considered, only, $\supset$ and $\rightarrow$ coincide and are equivalent to classical material implication (especially, we have $\phi \supset \psi \equiv^2 (\neg\phi) \vee \psi$ for every $\phi, \psi \in PROP^2_{PS}$); in the same vein, $\Leftrightarrow$ and $\leftrightarrow$ coincide and are equivalent to classical equivalence, and we have $\phi \equiv^2 \Box\phi \equiv^2 \diamond\phi$, $\neg\phi \equiv^2 \sim \phi$, $\phi! \equiv^2 false$, and $\odot\phi \equiv^2 true$ for every $\phi \in PROP^2_{PS}$. None of these equivalences holds as soon as $\top$ is considered as an additional "truth value".

The main observation here is that the incorporation of such additional connectives into $PROP^4_{PS}$ or $PROP^3_{PS}$ does not increase the expressiveness of the corresponding logic:[4] As $PROP^2_{PS}$ is functionally complete for TWO, $PROP^4_{PS}$ is functionally complete for $FOUR$ (Theorem 3.8 from [39]) and $PROP^3_{PS}$ is functionally complete for $THREE$ (Theorem 5.1 from [39]). This contrasts with the language of J_3 [31] (and its restrictions like LP [35], RP [33] and the logic of 3-inference of Levesque [34]) which includes neither $\oplus$ nor *both*, and despite the

[3] To be more precise, ! is introduced as a notation of the metalanguage in [35], and an element of the object language in [36].

[4] Nevertheless, like in classical logic, it may have an impact on the succinctness of the language; for instance, despite the fact that every classical propositional formula has an equivalent CNF form, it is well-known that a formula like $a_1 \Leftrightarrow a_2 \Leftrightarrow \ldots \Leftrightarrow a_n$ where every $a_i \in PS$ cannot be represented by an equivalent CNF formula of size polynomial in n.

presence of other connectives, is not expressive enough to enable the representation of a formula equivalent to *both* (see [59]). The functional completeness of $PROP^4_{PS}$ (resp. $PROP^3_{PS}$) w.r.t. *FOUR* (resp. *THREE*) explains why the focus has been laid on them; in some sense, all the logics considered in this paper (including classical logic) are restrictions of *FOUR*.

3.2 Monotonic inference relations

Let us now define some paraconsistent inference relations. The simplest one are based on model containment.

Definition 4 ($\models^4$-inference and $\models^3$-inference).
Let Σ and γ be two formulas from $PROP^4_{PS}$ (resp. $PROP^3_{PS}$). We note $\Sigma \models^4 \gamma$ (resp. $\Sigma \models^3 \gamma$) if and only if every 4-model (resp. 3-model) of Σ is a 4-model (resp. 3-model) of γ.

Example 3. Let $\Sigma = a \wedge \neg a \wedge (a \vee b) \wedge c \vee (\neg c \wedge d)$. We have $\Sigma \models^4 a$, $\Sigma \models^4 \neg a$ and $\Sigma \models^4 c$, but we have neither $\Sigma \models^4 b$, nor $\Sigma \models^4 d$. Similar conclusions can be derived using $\models^3$ instead of $\models^4$. However, $\models^3$ and $\models^4$ do not coincide: $true \models^3 a \vee \neg a$ holds while $true \models^4 a \vee \neg a$ does not hold since interpreting a as $\perp$ leads to interpret $a \vee \neg a$ as $\perp$ as well.

Since J_3, LP, RP[5] and Levesque's logic of 3-inference can be viewed as restricted cases of *THREE* (what differs is the underlying language which is a subset of $PROP^3_{PS}$), the corresponding inference relations are just restrictions of $\models^3$:

Definition 5 ($\models_{J_3}$-inference, $\models_L$-inference and $\models_{LP}$-inference).
Let Σ and γ be two formulas from the language of J_3 (resp. Levesque's logic of 3-inference, LP). We note $\Sigma \models_{J_3} \gamma$ (resp. $\Sigma \models_L \gamma$, $\Sigma \models_{LP} \gamma$) if and only if every 3-model of Σ is a 3-model of γ.

$\mid\!\sim\, = \models^4$, (resp. $\models^3$, $\models_{J_3}$, $\models_L$ and $\models_{LP}$) is obviously four-valued (resp. three-valued) *preferential*, i.e., it satisfies the following properties (provided that $\vdash$ in **(RW)** denotes $\models^4$ (resp. $\models^3$)):

(Ref)	$\alpha \mid\!\sim \alpha$	Reflexivity
(LLE)	If α and β are strongly equivalent and $\alpha \mid\!\sim \gamma$, then $\beta \mid\!\sim \gamma$	Left Logical Equivalence
(RW)	If $\alpha \mid\!\sim \beta$ and $\beta \vdash \gamma$, then $\alpha \mid\!\sim \gamma$	Right Weakening
(Or)	If $\alpha \mid\!\sim \gamma$ and $\beta \mid\!\sim \gamma$, then $\alpha \vee \beta \mid\!\sim \gamma$	Or
(Cut)	If $\alpha \wedge \beta \mid\!\sim \gamma$ and $\alpha \mid\!\sim \beta$, then $\alpha \mid\!\sim \gamma$	Cautious Cut
(CM)	If $\alpha \mid\!\sim \beta$ and $\alpha \mid\!\sim \gamma$, then $\alpha \wedge \beta \mid\!\sim \gamma$	Cautious Monotony

These properties have been stated in the framework of classical logic [60], but they can be extended to multi-valued settings in a straightforward way as above (such an extension has also been considered in [39]).

[5] Actually, we have $\models_{RP} = \models_{LP}$. In order to save some space, we will mainly focus on $\models_{LP}$ in the following.

$\models^4$, $\models^3$, $\models_{J_3}$, $\models_L$ and $\models_{LP}$ are even *monotonic* since the set of 4-models (resp. 3-models) of $\Sigma \wedge \phi$ is a subset of the set of 4-models (resp. 3-models) of Σ. As the previous example illustrates, the relations $\models^4$, $\models^3$, $\models_{J_3}$, $\models_L$ and $\models_{LP}$ are also *paraconsistent* [39]: There exist classically inconsistent formula $\Sigma \in PROP^2_{PS}$ whose closure w.r.t. any of the above mentioned relation is not the whole language $PROP^2_{PS}$. Furthermore, the restrictions of such relations over $PROP^2_{PS}$ are proper subsets of the classical entailment relation (noted $\models^2$ here instead of $\models$ for homogeneity). Stated otherwise, $\models^4$, $\models^3$, $\models_{J_3}$, $\models_L$ and $\models_{LP}$ are sub-classical over $PROP^2_{PS}$ (they can be viewed as approximations by default of classical entailment).

As usual when approximations are concerned, an important question is: What is the quality of the approximation? In this context, what are the classical consequences missed? And specifically, what happens when Σ is classically consistent?

Unfortunately, the news are not so good: $\models^4$, $\models^3$, $\models_{J_3}$, $\models_L$ and $\models_{LP}$ typically miss many expected consequences. A reason is that the disjunctive syllogism inference rule of classical logic fails for each of them. Thus, we do *not* have

$$c \wedge (\neg c \vee d) \models^4 d$$

because every 4-interpretation I s.t. $I(c) = \top$ and $I(d) = 0$ is a 4-model of $c \wedge (\neg c \vee d)$, but not a 4-model of d. Thus, on the running example, d is not found as a consequence of Σ despite the fact it comes logically from the sub-formulas c and $\neg c \vee d$ of Σ which are not involved in a contradiction. Since the disjunctive syllogism inference rule fails, none of the relations above coincides with $\models^2$ on $PROP^2_{PS}$ in the situation Σ is classically consistent - while we would expect it.

3.3 Non-Monotonic inference relations

Inference based on preferred models In order to circumvent such difficulties, more refined inference relations have been pointed out. The principle is to focus on some preferred models of Σ in order to keep as much information as possible. Preferential information typically characterize a preordering $\leq$ over the set of interpretations, and the $\leq$-preferred models of a base Σ are defined as the models of Σ that are minimal w.r.t. $\leq$. This approach leads to the paraconsistent inference relations $\models^4_{I1}$, $\models^4_{I2}$, $\models_{LP_m}$, $\models_{BS}$ and other variants.

Let us first consider $\models^4_{I1}$ and $\models^4_{I2}$ [39]. Each of these relations is characterized by a specific preference criterion. The first one consists in giving more credit to the 4-models of Σ which minimize the amount of inconsistent beliefs in Σ. The second one consists in preferring the 4-models of Σ which are as close as possible to its classical models. In a formal way, two (partial) preorderings $\leq_1$ and $\leq_2$ over the set of 4-interpretations over PS are used:

- $I \leq_1 J$ if and only if $\{x \in PS \mid I(x) \in \{\top\}\} \subseteq \{x \in PS \mid J(x) \in \{\top\}\}$;
- $I \leq_2 J$ if and only if $\{x \in PS \mid I(x) \in \{\top, \bot\}\} \subseteq \{x \in PS \mid J(x) \in \{\top, \bot\}\}$.

The preferred 4-models of Σ w.r.t. the first (resp. second) preference criterion is the set of 4-models of Σ which are minimal w.r.t. $\leq_1$ (resp. $\leq_2$).

Definition 6 **($\models^4_{I1}$-inference and $\models^4_{I2}$-inference).**
Let Σ and γ be two formulas from $PROP^4_{PS}$. We note $\Sigma \models^4_{I1} \gamma$ (resp. $\Sigma \models^4_{I2} \gamma$) if and only if every 4*-model of Σ minimal in its set w.r.t. $\leq_1$ (resp. $\leq_2$) is a* 4*-model of γ.*

Example 4. Let $\Sigma = a \wedge \neg a \wedge (a \vee b) \wedge c \wedge (\neg c \vee d)$. We have $\Sigma \models^4_{I1} d$ and $\Sigma \models^4_{I2} d$ (while we do *not* have $\Sigma \models^4 d$). On the contrary, we have neither $\Sigma \models^4_{I1} b$ nor $\Sigma \models^4_{I2} b$; the reason is that $a \wedge (a \vee b)$ is strongly equivalent to a in $FOUR$ (this shows that Σ is independent from b, so the impossibility to derive b is in some sense expected).

In [39], it is shown that both $\models^4_{I1}$ and $\models^4_{I2}$ are valuable inference relations since they are (four-valued) preferential, i.e., any of them satisfies reflexivity, left logical equivalence, right weakening, or, cautious monotony and cautious cut. They are also paraconsistent and sub-classical over $PROP^2_{PS}$. Contrariwise to $\models^4$, they are *non-monotonic*: While $\neg a \wedge (a \vee b) \models^4_{I1} b$ and $\neg a \wedge (a \vee b) \models^4_{I2} b$ hold, none of these two relations holds any longer when $\neg a \wedge (a \vee b)$ is logically strengthened by conjoining it with a. Hence, the scope of disjunctive syllogism remains limited whenever these inference relations are considered. Actually, this is mandatory under the paraconsistency requirement whenever the inference relation is s.t. $\alpha \vee \beta$ is a consequence of α since *ex contradictione quodlibet sequitur* comes from their unrestricted interaction (known as Lewis independent argument, cf. [5]):

$$\dfrac{\neg\alpha,\ \dfrac{\alpha}{\alpha \vee \beta}\ \texttt{disj.intro}}{\beta}\ \texttt{disj. syllogism}$$

Similarly, none of $\models^4_{I1}$ and $\models^4_{I2}$ is a transitive relation (and obviously, this contrasts with $\models^4$). Thus, while $\neg a \wedge a \models^4_{I1} \neg a \wedge (a \vee b)$ and $\neg a \wedge (a \vee b) \models^4_{I1} b$, we do *not* have $\neg a \wedge a \models^4_{I1} b$ (the same example can be used to show that $\models^4_{I2}$ is not transitive, *mutatis mutandis*).

Nevertheless, every classical model of a formula Σ of $PROP^4_{PS}$ is minimal in the set of 4-models of Σ w.r.t. any of the two preorderings. Therefore, both $\models^4_{I1}$ and $\models^4_{I2}$ coincide with $\models^2$ on the $PROP^2_{PS}$ fragment whenever Σ is classically consistent.

Finally, each of $\models^4_{I1}$ and $\models^4_{I2}$ is a proper superset of $\models^4$ but those relations do not coincide (see counterexamples in [39]).

Similar preference criteria have been considered in a three-valued framework [35, 36] in order to design the logic LP_m, less cautious than LP. Formally, the partial preordering $\leq_{LP_m}$ over the set of 3-interpretations over PS defined by $I \leq_{LP_m} J$ if and only if $\{x \in PS \mid I(x) \in \{\top\}\} \subseteq \{x \in PS \mid J(x) \in \{\top\}\}$ is considered:[6]

Definition 7 **($\models_{LP_m}$-inference).**
Let Σ and γ be two formulas from the language of LP_m. We note $\Sigma \models_{LP_m} \gamma$

[6] Obviously, the sole difference between $\leq_1$ - equal to $\leq_2$ when $\perp$ is not allowed - and $\leq_{LP_m}$ is the underlying set of non-classical interpretations.

if and only if every 3*-model of* Σ *that is minimal in its set w.r.t.* $\leq_{LP_m}$ *is a* 3*-model of* γ*.*

Example 5. Let $\Sigma = a \wedge \neg a \wedge (a \vee b) \wedge c \wedge (\neg c \vee d)$. We have $\Sigma \models_{LP_m} d$ (while we do *not* have $\Sigma \models^3 d$). But, we do *not* have $\Sigma \models_{LP_m} b$.

$\models_{LP_m}$ is (three-valued) preferential, non-monotonic, paraconsistent and subclassical over $PROP^2_{PS}$. Unlike $\models^3$, it is not transitive. Like $\models^4_{I1}$ and $\models^4_{I2}$, it coincides with $\models^2$ on the $PROP^2_{PS}$ fragment whenever Σ is classically consistent.

Other preference criteria result from giving more significance to the syntax of the database (especially, considering an additional comma connective). Indeed, when Σ is a finite set of formulas, it is possible to give more credit to the 3-models of Σ which maximize (w.r.t. $\subseteq$) the subset of formulas from Σ interpreted to 1. Formally, a partial preordering $\leq_{BS}$ over the set of 3-interpretations over PS can be defined as follows: $I \leq_{BS} J$ if and only if $\{\phi \in \Sigma \mid I(\phi) = 1\} \supseteq \{\phi \in \Sigma \mid J(\phi) = 1\}$. On this ground, another inference relation can be defined:

Definition 8 ($\models_{BS}$-inference).
Let Σ *be a finite set of formulas and* γ *be a formula from the language considered in [37]. We note* $\Sigma \models_{BS} \gamma$ *if and only if every* 3*-model of every formula from* Σ *minimal in its set w.r.t.* $\leq_{BS}$ *is a* 3*-model of* γ*.*

Example 6. Let $\Sigma = \{a, \neg a, (a \vee b), c, (\neg c \vee d)\}$. We have $\Sigma \models_{BS} d$. We also have $\Sigma \models_{BS} b$ (while we do *not* have $\Sigma \models_{LP_m} b$).

$\models_{BS}$ is three-valued preferential if left logical equivalence is stated as: If $E \cup \{\alpha\} \mid\!\sim \gamma$ and $\vdash (\alpha \leftrightarrow \beta)$ then $E \cup \{\beta\} \mid\!\sim \gamma$ (a set is required instead of a formula). It is also paraconsistent and coincides with $\models^2$ on the $PROP^2_{PS}$ fragment whenever Σ is a classically consistent set.

Other inference relations Other inference relations have been investigated in [53, 41]. A basic idea consists in characterizing inference in other ways than set containment, e.g., using *reductio ad absurdum.*[7]

Definition 9 ($\models^{4,inc}$-inference and $\models^{3,inc}$-inference).
Let Σ*,* γ *be two formulas from* $PROP^4_{PS}$*.* $\Sigma \models^{4,inc} \gamma$ *(resp.* $\Sigma \models^{3,inc} \gamma$*) if and only if* $\Sigma \wedge \neg\gamma$ *has no* 4*-model (resp. no* 3*-model).*

Such inference relations make sense when the language is not restricted to the monotonic fragment (or a subset of it); otherwise, they trivialize to the empty relation since every formula from the monotonic fragment has a 4-model and a 3-model.

Other relations can be defined by taking advantage of additional mechanisms. Among such mechanisms are the following ones:

[7] Another approach would consist in taking advantage of material implication and in defining $\models^{4,\supset}$ by $\Sigma \models^{4,\supset} \gamma$ if and only if $\models^4 \Sigma \supset \gamma$, and similarly for the three-valued setting. However, it is easy to prove that $\models^{4,\supset}$ coincides with $\models^4$, while $\models^{3,\supset}$ coincides with $\models^3$.

- considering only argumentative consequences of the belief bases.
- selecting the consequences of the belief base that are necessarily true.
- selecting as consequences of the belief base formulas that are so to speak "at least as true" as the belief base.

Definition 10 ($\models^{4,arg}_{\leq}$-inference, $\models^{4,1}_{\leq}$-inference and $\models^{4,t}_{\leq}$-inference). *Let $\leq$ be a binary relation over the set of all 4-interpretations over PS. Let Σ, γ be two formulas from $PROP^4_{PS}$. We define $\Sigma \models^4_{\leq} \gamma$ by $\forall I \in \min(4\text{-}mod(\Sigma), \leq)$, $I(\gamma) \in \{1, \top\}$. Then:*

- *$\Sigma \models^{4,arg}_{\leq} \gamma$ if and only if $\Sigma \models^4_{\leq} \gamma$ and $\Sigma \not\models^4_{\leq} \neg\gamma$.*
- *$\Sigma \models^{4,1}_{\leq} \gamma$ if and only if $\forall I \in \min(4\text{-}mod(\Sigma), \leq)$, $I(\gamma) = 1$.*
- *$\Sigma \models^{4,t}_{\leq} \gamma$ if and only if $\forall I \in \min(4\text{-}mod(\Sigma), \leq)$, $I(\Sigma) \leq_t I(\gamma)$.*

Example 7. Let $\Sigma = a \wedge \neg a \wedge (a \vee b) \wedge c \wedge (\neg c \vee d)$. We have:

- $\Sigma \models^{4,arg}_{\leq_1} c$ and $\Sigma \models^{4,arg}_{\leq_1} d$
- $\Sigma \models^{4,1}_{\leq_1} c$ and $\Sigma \models^{4,1}_{\leq_1} d$
- $\Sigma \models^{4,t}_{\leq_1} c$ and $\Sigma \models^{4,t}_{\leq_1} d$ and $\Sigma \models^{4,t}_{\leq_1} a$
- $\Sigma \models^{4,arg}_{\leq_2} c$ and $\Sigma \models^{4,arg}_{\leq_2} d$
- $\Sigma \models^{4,1}_{\leq_2} c$ and $\Sigma \models^{4,1}_{\leq_2} d$
- $\Sigma \models^{4,t}_{\leq_2} c$ and $\Sigma \models^{4,t}_{\leq_2} d$ and $\Sigma \models^{4,t}_{\leq_2} a$

Let us now present some logical properties for these three inference relations. $\models^{4,arg}_{\leq}$, $\models^{4,1}_{\leq}$ and $\models^{4,t}_{\leq}$ satisfy or, cautious cut and cautious monotony but they satisfy neither monotony nor right weakening in the general case. Neither $\models^{4,arg}_{\leq}$ nor $\models^{4,1}_{\leq}$ satisfy reflexivity in general, but $\models^{4,t}_{\leq}$ does. $\models^{4,1}_{\leq}$ satisfies left logical equivalence but none of $\models^{4,arg}_{\leq}$ or $\models^{4,t}_{\leq}$ does.

Again, similar inference relations can be stated in a three-valued setting. [41] investigated the logical properties of such relations, as well as their relative cautiousness, in the general case (i.e., for any $\leq$). A similar analysis has also been conducted in the specific case preorderings $\leq$ are considered. In particular, we considered both the universal ordering (i.e., the minimal 3-models of Σ are all its 3-models), $\leq_{LP_m}$, $\leq_{BS}$ (assuming that Σ is a finite set of formulas), and the variant of the latter based on cardinality instead on set containment. All the relations under consideration are paraconsistent and non-monotonic (when the universal ordering is not chosen).

One of the purposes for designing these relations is to derive paraconsistent inference relations more discriminating toward their consequences. Indeed, the three-valued inference relations depicted before suffer from a relative myopia: They do not make sufficient distinctions between the conclusions they provide given a belief base Σ. For instance, let $\Sigma = \Box a \wedge b \wedge c \wedge \neg c$. Both a, b and

c are consequences of Σ but they have different epistemic status w.r.t. Σ: a is necessary in Σ (it must be true, and only true, in every model of Σ), b is plausible since we have some evidence about its truth but no evidence about its falsity, and c is only possible since we have contradictory evidence about it.

Beyond S-3 logic As we will see in the next section, taking a preference relation into account in order to retain only the 3-models that are as close as possible to the classical ones (in a certain sense) has a computational cost. The complexity of the corresponding inference relations is at least one level higher in the polynomial hierarchy compared with $\models^3$. This just reflects that preference handling is a source of complexity (this is not a major surprise for any reader aware of the complexity of many preferential (and not paraconsistent) inference relations): Looking for preferred models requires searching an exponential space, and this source of complexity is typically orthogonal to the other source (exponentially many preferred models are possible).

In order to avoid such a complexity shift while refining $\models^3$ nevertheless, an approach is S-3 logic [51]. In a nutshell, instead of preferring only those models of the base Σ that are "as classical as possible" (or "as consistent as possible"), all the models of Σ that are classical over a prespecified set of symbols are kept. Thus, S-3-models are 3-models that are classical over the set S of propositional symbols. More formally, given a subset S of PS, an S-3-model I of $\Sigma \in PROP^3_{PS}$ is any 3-model of Σ s.t. $\forall x \in S$, $I(x) \neq \top$. Whenever S is fixed (and a part of the input), determining whether a given 3-model I of Σ is an S-3-model of it can be achieved in polynomial time (while the problem of determining whether a given 3-model I of Σ is a 3-model of it that is minimal w.r.t. $\leq_{LP_m}$ cannot be achieved in polynomial time – under the standard assumptions of complexity theory).

Given such a notion of S-3-model, another inference relation can be defined:

Definition 11 ($\models^3_S$-inference).
Let S be a subset of PS. Let Σ and γ be two formulas from $PROP^3_{PS}$. We note $\Sigma \models^3_S \gamma$ if and only if every S-3-model of Σ is an S-3-model of γ.

Now, the problem with S-3-logic is that a classically inconsistent belief base Σ may also be S-3-inconsistent, even if it is 3-consistent (especially when it is from the monotonic fragment). This happens when S contains all the propositional symbols that are brought into play in a conflict of Σ:

Example 8. Let $\Sigma = a \wedge \neg a \wedge (a \vee b) \wedge c \wedge (\neg c \vee d)$. Σ has no $\{a\}$-3-model.

In order to deal with the inconsistencies that may be revealed when using S-3 inference in $PROP^3_{PS}$, it was suggested in [42] to focus on some subsets S' of S, the ones for which the corresponding inference relations are not trivial. The approach is similar to the standard coherence-based approach to inconsistency handling, except that the inference relation is weakened by removing symbols from S instead of removing explicit beliefs from Σ:

Definition 12 (Consistent subsets).
Let Σ be a belief base from $PROP^3_{PS}$. Let $S \subseteq PS$ and $S_0 \subseteq S$ s.t. $\Sigma \not\models^3_{S_0} false$. A consistent subset *$S'$ of S w.r.t. Σ and S_0 is a subset of S containing S_0 and s.t. $\Sigma \not\models^3_{S'} false$. $\mathcal{S}(\Sigma, S, S_0)$ denotes the set of all consistent subsets of S w.r.t. Σ and S_0.*

In this definition, S_0 is a given set of symbols which must be interpreted classically. The condition $\Sigma \not\models^3_{S_0} false$ means that the set of consequences of Σ is not the whole language $PROP^3_{PS}$. Note that the existence of such a set S_0 (possibly empty) is ensured provided that Σ is 3-consistent.

Example 9. Let $\Sigma = a \wedge \neg a \wedge (a \vee b) \wedge c \wedge (\neg c \vee d)$. Every subset S of PS s.t. $a \notin S$ is a consistent subset of Σ.

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 $\mathcal{S}(\Sigma, S, S_0)$. In [42], the authors adhered to a skeptical approach: Any piece of belief is considered as a consequence of Σ given S and S_0 if and only if it is S'-3 entailed by Σ for every preferred consistent subset S' of S w.r.t. Σ and S_0.

Definition 13 ($\approx^{\mathcal{P},S_0}_S$-inference).
Let Σ be a belief base from $PROP^3_{PS}$ and S_0 be a subset of PS s.t. $\Sigma \not\models^3_{S_0} false$. Let $\mathcal{P}$ be a selection policy s.t. $\mathcal{S}_{\mathcal{P}}(\Sigma, S, S_0)$ is a subset of $\mathcal{S}(\Sigma, S, S_0)$ and let γ be a formula from $PROP_{PS}$. γ is a consequence *of Σ w.r.t. $\mathcal{P}$ and S_0, noted $\Sigma \approx^{\mathcal{P},S_0}_S \gamma$, if and only if $\forall S' \in \mathcal{S}_{\mathcal{P}}(\Sigma, S, S_0)$, $\Sigma \models^3_{S'} \gamma$.*

Example 10. Let $\Sigma = a \wedge \neg a \wedge (a \vee b) \wedge c \wedge (\neg c \vee d)$. If $S = (\emptyset, \{b, c, d\})$ and $\mathcal{S}_{\mathcal{P}}(\Sigma, S, S_0) = \{\{b, c, d\}\}$, we have $\Sigma \approx^{\mathcal{P},S_0}_S d$. We also have $\Sigma \approx^{\mathcal{P},S_0}_S a$, $\Sigma \approx^{\mathcal{P},S_0}_S \neg a$ but $\Sigma \not\approx^{\mathcal{P},S_0}_S b$.

Many selection policies for consistent subsets that are similar to the ones defined for consistent subbases in the standard coherence-based approach to inconsistency handling [16, 17, 61] are given in [42]. Thus, the possibilistic policy $\mathcal{PO}$ (or best-out), the linear-order policy $\mathcal{LO}$, the inclusion-preference policy $\mathcal{IP}$ (or discrimin) and the lexicographic policy $\mathcal{LE}$ (or leximin) have been defined for stratified S (i.e., S is given by a totally ordered partition of its propositional symbols); such policies consist in selecting the minimal elements in $\mathcal{S}_{\mathcal{P}}(\Sigma, S, S_0)$ w.r.t. the corresponding preorderings.

Every $\approx^{\mathcal{P},S_0}_S$ inference relation is a proper subset of $\models^3_S$ and is equal to it when $\Sigma \not\models^3_S false$ for many reasonable $\mathcal{P}$ (including $\mathcal{PO}$, $\mathcal{LO}$, $\mathcal{IP}$, $\mathcal{LE}$). The family of $\approx^{\mathcal{P},S_0}_S$ inference relations also includes $\models^3$ (just set S to $\emptyset$). All these relations are paraconsistent when Σ is assumed to be 3-consistent. However, in contrast to $\models^3_S$ that is monotonic, the $\approx^{\mathcal{P},S_0}_S$ inference relations are typically non-monotonic:

Example 11. Let $\Sigma = \neg a \wedge (a \vee b)$. If $S = (\emptyset, \{a, b\})$ and $\mathcal{S}_{\mathcal{P}}(\Sigma, S, S_0) = \{\{a, b\}\}$, we have $\Sigma \approx^{\mathcal{P},S_0}_S b$. Now, if $\Sigma = a \wedge \neg a \wedge (a \vee b)$, then $\{a, b\}$ may not belong any longer to $\mathcal{S}_{\mathcal{P}}(\Sigma, S, S_0)$. If this set becomes $\mathcal{S}_{\mathcal{P}}(\Sigma, S, S_0) = \{\{a\}\}$, then we have $\Sigma \not\approx_S^{\mathcal{P}, S_0} b$.

3.4 Other related systems for paraconsistent inference

We finally describe two logical approaches to paraconsistency that are related to the propositional multi-valued logics framework: Quasi-classical logic [29, 4] and the signed systems for paraconsistent inference [38].

Quasi-classical logic As explained before, *ex contradictione quodlibet sequitur* is obtained whenever full disjunctive syllogism (or more generally resolution) and disjunction introduction are allowed inference rules. Removing any of those rules is a drastic way to escape from trivialization in presence of inconsistency. However, it clearly leads to too cautious inference relations.

Quasi-classical logic, as defined in [29] and in [30], relies on the idea that it is sufficient to limit the way resolution and disjunction introduction interact in order to avoid trivialization in presence of inconsistency. In this logic, an application of the resolution rule can never follow an application of the disjunction introduction rule; additionally, applications of the resolution rule that would produce an inconsistency (i.e., the empty clause) are forbidden.

The language of quasi-classical propositional logic is the restriction of $PROP_{PS}^2$ where no constant symbols occur.

Quasi-classical logic has a simple proof theory; when considering CNF formulas only, it is sufficient to consider the resolution rule and the disjunction introduction rule (disjunction and conjunction being implicitly taken as associative and commutative). Starting from a CNF formula Σ, a CNF γ is derivable if and only if each clause of it is derivable, and a clause is derivable if and only if it has a resolution proof from Σ in which the empty clause does not occur and no applications of the resolution rule follow an application of the disjunction introduction rule.

Example 12. Let $\Sigma = a \wedge \neg a \wedge (a \vee b) \wedge c \wedge (\neg c \vee d)$. Here is a proof of d from Σ:

$$\frac{c, \neg c \vee d}{d}\ \texttt{resolution}$$

Here is a proof of b from Σ:

$$\frac{\neg a, a \vee b}{b}\ \texttt{resolution}$$

Here is a proof of $d \vee e$ from Σ:

$$\frac{\dfrac{c, \neg c \vee d}{d}\ \texttt{resolution}}{d \vee e}\ \texttt{disj.intro.}$$

Contrastingly, neither the empty clause nor e are derivable from Σ in quasi-classical logic.

Quasi-classical logic also has a semantics w.r.t. which the proof theory is sound and complete. In order to define the quasi-classical logic entailment relation $\models_{QC}$, we first need to make precise the underlying notion of interpretation, as well as the notions of *strong satisfaction* and *weak satisfaction*. The first part is easy: An interpretation in quasi-classical logic (QC interpretation for short) is a 4-interpretation over PS. For the other notions, our presentation slightly departs from the one reported in [29, 4] so as to keep as much as possible the notations used in this chapter:

Definition 14 (Strong satisfaction).
The notion of strong satisfaction *of a formula Σ (from the language of quasi-classical logic) by a QC interpretation I, noted $I \models_s \Sigma$, is defined as follows:*

- *If $\alpha \in PS$, then $I \models_s \alpha$ if and only if $I(\alpha) \in \{1, \top\}$;*
- *If $\alpha \in PS$, then $I \models_s \neg\alpha$ if and only if $I(\alpha) \in \{0, \top\}$;*
- *If $\alpha_1 \vee \ldots \vee \alpha_n$ is a clause, then $I \models_s \alpha_1 \vee \ldots \vee \alpha_n$ if and only if there exists $i \in 1 \ldots n$ s.t. $I \models_s \alpha_i$ and for every $i \in 1 \ldots n$, if $I \models_s \bar{\alpha}_i$ then $I \models_s \alpha_1 \vee \ldots \vee \alpha_{i-1} \vee \alpha_{i+1} \vee \ldots \vee \alpha_n$;*
- *$I \models_s \alpha \wedge \beta$ if and only if $I \models_s \alpha$ and $I \models_s \beta$;*
- *$I \models_s (\neg\neg\alpha) \vee \beta$ if and only if $I \models_s \alpha \vee \beta$;*
- *$I \models_s \neg(\alpha \wedge \beta) \vee \gamma$ if and only if $I \models_s \neg\alpha \vee \neg\beta \vee \gamma$;*
- *$I \models_s \neg(\alpha \vee \beta) \vee \gamma$ if and only if $I \models_s (\neg\alpha \wedge \neg\beta) \vee \gamma$;*
- *$I \models_s \alpha \vee (\beta \wedge \gamma)$ if and only if $I \models_s (\alpha \vee \beta) \wedge (\alpha \vee \gamma)$;*
- *$I \models_s \alpha \wedge (\beta \vee \gamma)$ if and only if $I \models_s (\alpha \wedge \beta) \vee (\alpha \wedge \gamma)$;*
- *$I \models_s (\alpha \supset \beta) \vee \gamma$ if and only if $I \models_s \neg\alpha \vee \beta \vee \gamma$;*
- *$I \models_s \neg(\alpha \supset \beta) \vee \gamma$ if and only if $I \models_s (\alpha \wedge \neg\beta) \vee \gamma$.*

Clearly enough, the definition for disjunction is more restricted than the classical definition. Actually, the notion of strong satisfaction implicitly translates the formula into CNF whenever disjunction applies to more sophisticated subformulas than literals. As to clauses $\alpha_1 \vee \ldots \vee \alpha_n$, it asks at least one disjunct to be satisfied, and in addition, when the interpretation strongly satisfies a complementary literal $\bar{\alpha}_i$ ($i \in 1 \ldots n$), it must satisfy the subclause obtained by removing α_i in the clause. This is necessary to restore the connection between a literal and its complementary literal which is decoupled when 4-interpretations are considered (remember that $\models^4$ does not satisfy the disjunctive syllogism).

The notion of strong satisfaction can be relaxed to the less demanding notion of weak satisfaction defined by:

Definition 15 (Weak satisfaction).
The notion of weak satisfaction *of a formula Σ (from the language of quasi-classical logic) by a QC interpretation I, noted $I \models_w \Sigma$, is defined as follows:*

- *If $\alpha \in PS$, then $I \models_w \alpha$ if and only if $I(\alpha) \in \{1, \top\}$;*
- *If $\alpha \in PS$, then $I \models_w \neg\alpha$ if and only if $I(\alpha) \in \{0, \top\}$;*
- *$I \models_w \alpha \vee \beta$ if and only if $I \models_w \alpha$ or $I \models_w \beta$;*
- *$I \models_w \alpha \wedge \beta$ if and only if $I \models_w \alpha$ and $I \models_w \beta$;*

– $I \models_w (\neg\neg\alpha) \vee \beta$ *if and only if* $I \models_w \alpha \vee \beta$;
– $I \models_w \neg(\alpha \wedge \beta) \vee \gamma$ *if and only if* $I \models_w \neg\alpha \vee \neg\beta \vee \gamma$;
– $I \models_w \neg(\alpha \vee \beta) \vee \gamma$ *if and only if* $I \models_w (\neg\alpha \wedge \neg\beta) \vee \gamma$;
– $I \models_w (\alpha \supset \beta) \vee \gamma$ *if and only if* $I \models_w \neg\alpha \vee \beta \vee \gamma$;
– $I \models_w \neg(\alpha \supset \beta) \vee \gamma$ *if and only if* $I \models_w (\alpha \wedge \neg\beta) \vee \gamma$.

When $I \models_s \Sigma$ holds, I is said to be a *strong model* of Σ. When $I \models_w \Sigma$ holds, I is said to be a *weak model* of Σ. Obviously, every strong model of Σ is a weak model of Σ but the converse does not hold in the general case.

It is not very difficult to show that every CNF formula Σ has a strong model (hence a weak one as well), even if it is classically inconsistent (the interpretation I s.t. $I(x) = \top$ for every $x \in PS$ does the job).

Definition 16 ($\models_{QC}$-inference).
Let Σ and Φ be two formulas from the language of quasi-classical logic. We have $\Sigma \models_{QC} \Phi$ if and only if for every QC interpretation I, if $I \models_s \Sigma$, then $I \models_w \Phi$.

A detailed study of the logical properties of quasi-classical inference is given in [30]; it is shown that quasi-classical entailment satisfies reflexivity, monotony, and-introduction, or-elimination and consistency preservation, but none of supraclassicality, right modus ponens, conditionalization, deduction, cut, transitivity, unit cumulativity, right weakening and left logical equivalence is satisfied.

An interesting feature of quasi-classical logic is that it coincides with classical entailment in the clausal case when the CNF belief base Σ is consistent and the CNF query γ does not contain any valid clauses.

Finally, let us note that quasi-classical logic has been extended recently to quasi-classical possibilistic logic [62]. This logic extends both possibilistic logic and quasi-classical logic, whilst preserving their merits. Thus, conflicts taking place at the same level of certainty are handled as in quasi-classical logic, while the remaining conflicts are handled as in possibilistic logic. $\models_{Q\Pi L}$ denotes the corresponding paraconsistent inference relation.

Signed systems for paraconsistent inference To conclude this section, let us consider the paraconsistent inference relations introduced in [38] in a default logic approach.

In Besnard and Schaub's work [38], the language used is the set of NNF formulas from $PROP^2_{PS}$, i.e., formulas built up from the connectives $\neg$, $\wedge$, $\vee$, only, and for which the scope of every occurrence of $\neg$ is a propositional symbol from PS.[8] Every formula Σ is associated to a default theory $\langle \Sigma^{\pm}, D_{\Sigma} \rangle$ where:

– $\Sigma^{\pm}$ is a formula in the language $PROP^2_{PS^{\pm}}$ 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^-.

[8] The NNF assumption can be relaxed through the notion of polarity (see [38]) but every occurrence of a subformula $\phi \Leftrightarrow \psi$ must be replaced first by $(\phi \supset \psi) \wedge (\psi \supset \phi)$.

- $D_\Sigma = \{\delta_x \mid x \in PS\}$ is a set of default rules

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

Negation is given a special treatment; first, every literal is rendered independent from its negation through renaming; then the corresponding dependence relations are re-introduced in a parsimonious way, so that no inconsistency occurs. Based on the extensions of the default theory $\langle \Sigma^\pm, D_\Sigma \rangle$, several paraconsistent consequence relations can be defined, especially the relations $\vdash_s$ and $\vdash_s^\pm$:

Definition 17 ($\vdash_s$-inference and $\vdash_s^\pm$-inference).
Let Σ and γ be two formulas from $PROP_{PS}^2$.

- *γ is a* skeptical unsigned consequence *of Σ, noted $\Sigma \vdash_s \gamma$, if and only if γ belongs to every extension of $\langle \Sigma^\pm, D_\Sigma \rangle$.*
- *γ is a* skeptical signed consequence *of Σ, noted $\Sigma \vdash_s^\pm \gamma$, if and only if $\gamma^\pm$ belongs to every extension of $\langle \Sigma^\pm, D_\Sigma \rangle$.*

Example 13. Let $\Sigma = a \wedge \neg a \wedge (a \vee b) \wedge c \wedge (\neg c \vee d)$. Σ is associated with:

- $\Sigma^\pm = a^+ \wedge a^- \wedge (a^+ \vee b^+) \wedge c^+ \wedge (c^- \vee d^+)$.
- $D_\Sigma = \{\frac{: x^+ \Leftrightarrow \neg x^-}{(x \Leftrightarrow x^+) \wedge (\neg x \Leftrightarrow x^-)} \mid x \in PS\}$.

This default theory has two extensions. We have:

- $\Sigma \vdash_s c \wedge d \wedge (a \vee b)$ but $\Sigma \not\vdash_s a$, $\Sigma \not\vdash_s \neg a$ and $\Sigma \not\vdash_s b$.
- $\Sigma \vdash_s^\pm c \wedge d \wedge a \wedge \neg a \wedge (a \vee b)$ but $\Sigma \not\vdash_s^\pm b$.

Reflexivity is satisfied by $\vdash_s^\pm$ while it is not ensured for $\vdash_s$. On the contrary, right weakening is ensured for $\vdash_s$ but not for $\vdash_s^\pm$. Both $\vdash_s^\pm$ and $\vdash_s$ coincide with classical entailment as soon as Σ is (classically) consistent.

Other inference relations can be obtained by considering other inference mechanisms (credulous inference, prudent inference), other ways to render each propositional symbol independent from its negation, other kinds of extensions when some preferential information (under the form of a total preordering over PS) are available. It must be noted that, unsurprisingly, taking advantage of such preferential information does not lead to a complexity shift (see [55] for details).

4 Complexity results

4.1 Main results

In this section, we report the complexity of inference problems for the different (paraconsistent) inference relations presented in the previous section. For each inference relation considered $\vdash_L$, the decision problem is defined as follows:

Definition 18 (Decision problem for $\vdash_L$).
Input: *A pair $\langle \Sigma, \gamma \rangle$ (plus a possibly stratified subset S of PS for the relations for which this is relevant) of formulas from the language L of the logic $\langle L, \vdash_L \rangle$ considered.*
Question : *Does $\Sigma \vdash_L \gamma$ hold?*

Complexity results for most of the relations come from [34, 52] for the monotonic ones and from [53, 41] for the non-monotonic ones. [54] gives complexity results for quasi-classical logic and [62] for quasi-possibilistic logic. The complexity of the inference relations based on signed systems have been investigated in [53, 55]. Finally, [42] presents complexity results for the $\approx_S^{\mathcal{P},S_0}$ relations.

Proposition 1. *The complexity of the decision problem for $\vdash_L$ has been identified as follows:*

- coNP*-complete when $\vdash_L$ is $\models^4$, $\models^3$, $\models_{J_3}$, $\models_{LP}$, $\models_L$, $\models_S^3$, $\models^{4,inc}$, $\models^{3,inc}$.*
- *$\Delta_2^p[\mathcal{O}(log\ n)]$-complete when $\vdash_L$ is $\approx_S^{\mathcal{PO},S_0}$.*
- *Δ_2^p-complete when $\vdash_L$ is $\approx_S^{\mathcal{LO},S_0}$ or $\approx_S^{\mathcal{LE},S_0}$.*
- *Π_2^p-complete when $\vdash_L$ is $\models_{I1}^4$, $\models_{I2}^4$, $\models_{LP_m}$, $\models_{BS}$, $\vdash_s^{\pm}$, $\vdash_s$, $\approx_S^{\mathcal{IP},S_0}$, $\models_{\leq_1}^{4,1}$, $\models_{\leq_2}^{4,1}$, $\models_{\leq_1}^{4,t}$, $\models_{\leq_2}^{4,t}$, $\models_{\leq_{LP_m}}^{3,1}$ or $\models_{\leq_{LP_m}}^{3,t}$.*
- *D_2^p-complete when $\vdash_L$ is $\models_{\leq_1}^{4,arg}$, $\models_{\leq_2}^{4,arg}$ or $\models_{\leq_{LP_m}}^{3,arg}$.*

To conclude this section, let us sketch some of the proofs of the results given in the proposition above (the remaining ones can be found in [52, 54, 42]).

All the membership proofs are easy. For any relation $\vdash_L$ among $\models^4$, $\models^3$, $\models_{J_3}$, $\models_{LP}$, $\models_L$, $\models_S^3$, membership can be easily proven by showing that the corresponding complementary problems are in NP, thanks to the following non-deterministic algorithm:

```
1. Guess an interpretation I over Var(Σ) ∪ Var(γ);
2. Check that I is a model of Σ but not a model of γ
```

By interpretation (resp. model) here, it is meant a 4-interpretation (resp. a 4-model) when $\models^4$ is concerned, an S-3-interpretation (resp. an S-3-model) when $\models_S^3$ is concerned, and a 3-interpretation (resp. a 3-model) otherwise.

The membership results for $\models^{4,inc}$, $\models^{3,inc}$ comes from a similar non-deterministic algorithm (only step. 2 changes and consists in checking that I is a model of $\Sigma \wedge \neg\gamma$).

For any relation $\vdash_L$ among $\models_{I1}^4$, $\models_{I2}^4$, $\models_{LP_m}$, $\models_{BS}$, membership is proven by showing that the corresponding complementary problems are in Σ_2^p. This is achieved through the following non-deterministic algorithm:

```
1. Guess an interpretation I over Var(Σ) ∪ Var(γ);
2. Check that I is a model of Σ minimal w.r.t. ≤
   through a call to an NP oracle;
3. Check that I is not a model of γ
```

This time, by interpretation (resp. model), it is meant a 4-interpretation (resp. a 4-model) when $\models^4_{I1}$ or $\models^4_{I2}$ is concerned, a 3-interpretation (resp. a 3-model) otherwise.

Indeed, checking that I is not minimal w.r.t. any relation $\leq$ that can be decided in deterministic polynomial time can be easily done in non-deterministic polynomial time, just by guessing an interpretation J over $Var(\Sigma) \cup Var(\gamma)$, and by checking in deterministic polynomial time that J is a model of Σ, and that $J \leq I$ holds while $I \leq J$ does not hold. Obviously, $\leq = \leq_1$, $\leq_2$, $\leq_{LP_m}$ or $\leq_{BS}$ can be decided in deterministic polynomial time.

A similar non-deterministic algorithm can be used for $\models^{4,1}_{\leq_1}$, $\models^{4,1}_{\leq_2}$, $\models^{3,1}_{\leq_{LP_m}}$ (resp. $\models^{4,t}_{\leq_1}$, $\models^{4,t}_{\leq_2}$, $\models^{3,t}_{\leq_{LP_m}}$) (only step. 3 changes and consists in checking that $I(\gamma) \neq 1$ (resp. $I(\Sigma) \not\leq_t I(\gamma)$)).

As to $\vdash^{\pm}_{s}$, $\vdash_s$, the membership results are are direct consequences of Theorem 5.2 from [63] (see also [64]) showing that skeptical default reasoning is in Π^p_2 in the general case.

As to the relations $\approx^{\mathcal{PO},S_0}_S$, $\approx^{\mathcal{LO},S_0}_S$, $\approx^{\mathcal{IP},S_0}_S$ and $\approx^{\mathcal{LE},S_0}_S$, the membership results come from modular and faithful polytime translations from the inference problem for such relations to the corresponding inference problems from stratified belief bases (Proposition 3.11 from [42]), and the membership results for the latter decision problems as reported in [44, 18].

As to the argumentative relations $\models^{4,arg}_{\leq_1}$, $\models^{4,arg}_{\leq_2}$ or $\models^{3,arg}_{\leq_{LP_m}}$, the membership results come directly from the definitions and the membership results for $\models^4_{\leq_1}$, $\models^4_{\leq_2}$ and $\models^3_{\leq_{LP_m}}$.

All the hardness proofs come from the corresponding restrictions to the CNF case, except to what concerns $\models^4$, $\models^3$, $\models_{J_3}$, $\models_{LP}$, $\models_L$, $\models^{4,inc}$, $\models^{3,inc}$.

As to $\models^4$, hardness comes from the following polytime reduction from the canonical NP-complete problem SAT; let Σ be a CNF formula s.t. $Var(\Sigma) = \{x_1, \ldots, x_n\}$; we have that Σ is classically consistent if and only if $\Sigma \wedge \bigwedge_{i=1}^n (x_i \vee \neg x_i) \not\models^4 \bigvee_{i=1}^n (x_i \wedge \neg x_i)$ (see [52]).

As to $\models^3$, $\models_{J_3}$, $\models_{LP}$, $\models_L$, hardness comes from the following polytime reduction from SAT; let Σ be a CNF formula s.t. $Var(\Sigma) = \{x_1, \ldots, x_n\}$; we have that Σ is classically consistent if and only if $\Sigma \not\models^3 \bigvee_{i=1}^n (x_i \wedge \neg x_i)$ (see [52]).

As to $\models^{4,inc}$ and $\models^{3,inc}$, we associate in polynomial time to any CNF Σ the formula $\Sigma_\Box$ obtained by replacing in Σ every occurrence of any literal l by $\Box l$. Now, Σ has a 2-model if and only if $\Sigma_\Box$ has a 4-model if and only if $\Sigma_\Box$ has a 3-model, due to the truth table of $\Box$. Accordingly, Σ is classically consistent if and only if $\Sigma_\Box \not\models^4 new$ if and only if $\Sigma_\Box \not\models^3 new$, where $new \in PS \setminus Var(\Sigma)$.

4.2 Restrictions

Let us now give more specific results, obtained by imposing further restrictions to the decision problem.

The CNF case Significant decreases in complexity can be obtained for some inference relations when Σ and γ are restricted to CNF formulas (from $PROP^2_{PS}$).

Proposition 2. *The complexity of the decision problem for $\vdash_L$ has been identified as follows in the case Σ and γ are CNF formulas (from $PROP^2_{PS}$):*

- *trivial (i.e., in time $\mathcal{O}(1)$) when $\vdash_L$ is $\models^{4,inc}$ or $\models^{3,inc}$ because such relations are empty.*
- *in* P *when $\vdash_L$ is $\models^4$ or $\models^3$, $\models_{J_3}$, $\models_{LP}$ or $\models_L$.*
- coNP*-complete when $\vdash_L$ is $\models^3_S$.*
- *$\Delta^p_2[\mathcal{O}(log\ n)]$-complete when $\vdash_L$ is $\approx^{\mathcal{PO},S_0}_S$.*
- *Δ^p_2-complete when $\vdash_L$ is $\approx^{\mathcal{LO},S_0}_S$ or $\approx^{\mathcal{LE},S_0}_S$.*
- *Π^p_2-complete when $\vdash_L$ is $\models^4_{I1}$, $\models^4_{I2}$, $\models_{LP_m}$, $\models_{BS}$, $\vdash^{\pm}_s$, $\vdash_s$, $\approx^{\mathcal{IP},S_0}_S$, $\models^{4,1}_{\leq_1}$, $\models^{4,1}_{\leq_2}$, $\models^{4,t}_{\leq_1}$, $\models^{4,t}_{\leq_2}$, $\models^{3,1}_{\leq_{LP_m}}$ or $\models^{3,t}_{\leq_{LP_m}}$.*
- *D^p_2-complete when $\vdash_L$ is $\models^{4,arg}_{\leq_1}$, $\models^{4,arg}_{\leq_2}$ or $\models^{3,arg}_{\leq_{LP_m}}$.*

Again, let us sketch some of the proofs of the results given in the proposition above. The Π^p_2-hardness proof for $\models_{LP_m}$ is fully detailed; it is a bit long, but on the one hand, it is central (in the sense that the other Π^p_2-hardness proofs are based on it) and on the other hand, it exhibits an interesting connection with the notion of conflict (a minimally inconsistent subset of formulas), which is at the core of many approaches to inconsistency tolerant reasoning based on formula inhibition.

Membership to P for $\models^4$ or $\models^3$, (and their restrictions $\models_{J_3}$, $\models_{LP}$ or $\models_L$) intuitively comes from the fact that the disjunctive syllogism rule fails for each of them. Thus, when Σ and γ are CNF formulas, $\Sigma \models^4 \gamma$ holds (resp. $\Sigma \models^3 \gamma$ holds) if and only if every clause of γ is subsumed by a clause from Σ (resp. every clause of γ is subsumed by a clause from Σ or it contains a pair of complementary literals) [34].

Hardness for $\models^3_S$ simply comes from the fact that classical entailment is coNP-hard and coincides with $\models^3_S$ when $S = PS$.

Hardness for $\approx^{\mathcal{PO},S_0}_S$ comes from the $\Delta^p_2[\mathcal{O}(log\ n)]$-hardness of inference from stratified belief bases interpreted under the possibilistic policy (Theorem 6.5 from [44]) and a (modular and faithful) polytime translation of this inference problem into the decision problem for $\approx^{\mathcal{PO},S_0}_S$-inference (Proposition 3.10 in [42]).

Similarly, hardness for $\approx^{\mathcal{LO},S_0}_S$ (resp. $\approx^{\mathcal{LE},S_0}_S$) comes from the Δ^p_2-hardness of inference from stratified belief bases interpreted under the linear-order policy (Theorem 5.9 from [44]) (resp. the lexicographic policy (see [18])) and a (modular and faithful) polytime translation of such inference problems from stratified belieb bases into the corresponding decision problem for $\approx^{\mathcal{LO},S_0}_S$-inference (resp. $\approx^{\mathcal{LE},S_0}_S$) (again, see Proposition 3.10 in [42]).

Hardness for $\models^4_{I1}$ and $\models^4_{I2}$ comes from hardness for $\models_{LP_m}$, together with the fact that in the fragment $\{\neg, \wedge, \vee, true, false\}$, we have $\Sigma \models^4_{I2} \gamma$ if and only if $\Sigma \models_{LP_m} \gamma$ and when in addition γ is a CNF formula that does not contain any valid clause, we have $\Sigma \models^4_{I1} \gamma$ if and only if $\Sigma \models_{LP_m} \gamma$ (Proposition 5.4 from [39]).

Hardness for $\vdash^{\pm}_s$ and $\approx^{\mathcal{IP},S_0}_S$, comes from the fact that $\vdash^{\pm}_s$ and $\approx^{\mathcal{IP},\emptyset}_{PS}$ coincide with $\models_{LP_m}$ on the monotonic fragment (see Propositions 3.8 and 3.9 in [42] and Lemma 5).

The Π^p_2-hardness results for $\models_{LP_m}$, $\vdash_s$, $\models_{BS}$ remain to be proven. Let us do it in a gentle way, through a number of intermediate lemmata. We first need the following notions:

Definition 19. *Let $\Sigma = \{\phi_1, \ldots, \phi_n\}$ be a finite set of formulas from $PROP^2_{PS}$. Let α be a formula from $PROP^2_{PS}$.*

- *A subset S of Σ is an* α-conflict *of Σ if and only if $S \cup \{\alpha\}$ is classically inconsistent but every proper subset of it is consistent. We note $Conf(\Sigma, \alpha)$ the set of all formulas ϕ_i from Σ where ϕ_i belongs to at least one α-conflict of it. When $\alpha \equiv^2$ true holds, an α-conflict of Σ is simply referred to as a* conflict *of Σ.*
- *A* var-conflict *of Σ is a subset V of $Var(\Sigma)$ s.t. there exists a conflict S of Σ satisfying $V = Var(S)$.*
- *A* minimal var-conflict *of Σ is a var-conflict of Σ that is minimal w.r.t. $\subseteq$.*

Given a CNF formula $\Sigma = \{\phi_1, \ldots, \phi_n\}$ (we represent it as the set of its clauses instead of the conjunction of them for simplicity), we are going to prove successively that:

(1) Determining whether $\phi_i \in \Sigma$ does not belong to any conflict of Σ is Π^p_2-hard (this holds even in the case when ϕ_i is a propositional symbol).
(2) Determining whether $p \in Var(\Sigma)$ does not belong to any minimal var-conflict of Σ is Π^p_2-hard.
(3) Determining whether $p \in Var(\Sigma)$ satisfies $\Sigma \models_{LP_m} p$ (resp. $\Sigma \vdash_s p$) is Π^p_2-hard.

We prove (1) by considering a polynomial reduction from the inference problem related to WIDTIO belief revision operator [65]. Let $\Sigma = \{\phi_1, \ldots, \phi_n\}$ be a finite set of formulas from $PROP^2_{PS}$. Let α be a formula from $PROP^2_{PS}$. Let us recall that the revised database $\Sigma \circ_W \alpha$ can be defined up to classical equivalence by the formula $\alpha \wedge \bigwedge_{\phi_i \in \Sigma \setminus Conf(\Sigma,\alpha)} \phi_i$. Eiter and Gottlob proved in [43] that the corresponding inference problem is Π^p_2-hard. Looking carefully at the proofs of Lemma 6.2 and Theorem 8.2 from [43], it is easy to show that $\Sigma \circ_W p \overset{?}{\models} q$ remains Π^p_2-hard in the restricted case when Σ is a classically consistent CNF formula, p and q are propositional variables from PS, $q \in \Sigma$, and $p \notin \Sigma$ (especially, we can assume without loss of generality that the matrix of the $2-QBF_\forall$ formula E used in the proof of Lemma 6.2 is a CNF formula). On this ground, we prove the following lemma:

Lemma 1. *Let $\Sigma = \{\phi_1, \ldots, \phi_n\}$ be a classically consistent CNF formula. Let p and q be two propositional symbols from PS, $q \in \Sigma$, and $p \notin \Sigma$. Determining whether q does not belong to any p-conflict of Σ is Π^p_2-hard.*

Proof of the lemma: It is sufficient to show that $\Sigma \circ_W p \models^2 q$ holds if and only if q does not belong to any p-conflict of Σ. In order to simplify a bit the writing, let us note $no - conflict(\Sigma, p) =_{def} \Sigma \setminus Conf(\Sigma, p)$ (this is the set of formulas from Σ not appearing in any p-conflict of it). The "only-if" way is obvious since $\Sigma \circ_W p \equiv^2 p \wedge \bigwedge_{\phi_i \in no-conflict(\Sigma,p)} \phi_i$ and $q \in no - conflict(\Sigma, p)$ by assumption. Conversely, let us assume that $\Sigma \circ_W p \models^2 q$ holds and there exists a p-conflict S of Σ s.t. $q \in S$. First of all, $\Sigma \circ_W p \models^2 q$ holds if and only if $no - conflict(\Sigma, p) \models^2 \neg p \vee q$. Now, let $S' = S \setminus \{q\}$. Since $S \models^2 \neg p$, we also have $S' \cup \{q\} \models^2 \neg p$, or equivalently $S' \models^2 \neg q \vee \neg p$. Moreover, by minimality of a conflict, we have $S' \not\models^2 \neg p$. Consider the subset $S' \cup no - conflict(\Sigma, p)$ of Σ. By monotony of $\models$, we have $S' \cup no - conflict(\Sigma, p) \models^2 (\neg q \vee \neg p) \wedge (\neg p \vee q)$. This is equivalent to state that $S' \cup no - conflict(\Sigma, p) \models^2 \neg p$. Accordingly, this implies that there exists a p-conflict $S'' \subseteq S' \cup no - conflict(\Sigma, p)$ of Σ. Since $S' \not\models^2 \neg p$, this implies that $S'' \cap no - conflict(\Sigma, p) \neq \emptyset$, contradiction. ■

Let us now show that considering conflicts (instead of p-conflicts) is sufficient:

Lemma 2. *Let $\Sigma = \{\phi_1, \ldots, \phi_n\}$ be a classically consistent CNF formula. Let p be a propositional symbol from PS and ϕ_i a clause from Σ. ϕ_i belongs to a p-conflict of Σ if and only if ϕ_i belongs to a conflict of $\Sigma \cup \{p\}$ and $\phi_i \not\equiv^2 p$.*

Proof of the lemma: Let us first show that S is a p-conflict of Σ if and only if $S \cup \{p\}$ is a conflict of $\Sigma \cup \{p\}$. Obviously, we have $S \models^2 \neg p$ if and only if $S \cup \{p\}$ is inconsistent. It remains to show that $\forall \phi_i \in S, S \setminus \{\phi_i\} \not\models^2 \neg p$ if and only if $\forall \phi_i \in S \cup \{p\}, (S \cup \{p\}) \setminus \{\phi_i\} \not\models^2 false$.

- "if" way: Assume that there exists $\phi_i \in S \cup \{p\}$ s.t. $(S \cup \{p\}) \setminus \{\phi_i\} \models^2 false$. If $\phi_i \neq p$, then we have $(S \cup \{p\}) \setminus \{\phi_i\} = (S \setminus \{\phi_i\}) \cup \{p\}$. Accordingly, we have $(S \setminus \{\phi_i\}) \cup \{p\} \models^2 false$, which contradicts our assumption (minimality of a p-conflict).
  Otherwise, $\phi_i = p$ and $(S \cup \{p\}) \setminus \{\phi_i\}$ is a subset of S, hence a subset of Σ. The fact that $(S \cup \{p\}) \setminus \{\phi_i\} \models^2 false$ contradicts the assumption that Σ is consistent.
- "only-if" way: Assume that there exists $\phi_i \in S$ s.t. $S \setminus \{\phi_i\} \models^2 \neg p$. This is equivalent to state that $(S \setminus \{\phi_i\}) \cup \{p\} \models^2 false$.
  If $\phi_i \neq p$, then we have $(S \setminus \{\phi_i\}) \cup \{p\} = (S \cup \{p\}) \setminus \{\phi_i\}$. Accordingly, we have $(S \cup \{p\}) \setminus \{\phi_i\} \models^2 false$, with $\phi_i \in S$, which contradicts our assumption (minimality of a conflict).
  Otherwise, $\phi_i = p$ and $(S \setminus \{\phi_i\}) \cup \{p\} = S$, hence a subset of Σ. The fact that $(S \setminus \{\phi_i\}) \cup \{p\} \models^2 false$ contradicts the assumption that Σ is consistent.

Once this is established, it is sufficient to show that p does not belong to any p-conflict of Σ in the case $p \in \Sigma$. Assume that this is not the case: Let S be a p-conflict of Σ s.t. $p \in S$. Let $S' = S \setminus \{p\}$. Since $S = S' \cup \{p\}$ and S is a p-conflict, we have $S' \cup \{p\} \models^2 \neg p$. This is equivalent to state that $S' \models^2 \neg p$ through the (meta)deduction theorem for propositional logic. Hence, S' is a p-conflict of Σ and a proper subset of S, which contradicts the existence of S. ■

We are now ready to prove (1): To every triple $\langle \Sigma = \{\phi_1, \ldots, \phi_n\}, p, \phi_i\rangle$ where Σ is a classically consistent CNF formula, $p \notin \Sigma$ and $\phi_i \in \Sigma$, we associate in polynomial time the ordered pair $\langle \Sigma \cup \{p\}, \phi_i\rangle$. From Lemma 2, ϕ_i does not occur in any p-conflict of Σ if and only if ϕ_i does not occur in any conflict of $\Sigma \cup \{p\}$ or $\phi_i = p$. Now, the case $\phi_i = p$ can be excluded since $p \notin \Sigma$. The Π_2^p-hardness result given in Lemma 1 completes the proof.

In order to prove (2), we associate in polynomial time to every ordered pair $\langle \Sigma = \{\phi_1, \ldots, \phi_n\}, \phi_i\rangle$ (where Σ is a CNF formula and ϕ_i a clause from Σ) the following ordered pair: $\langle \bigcup_{\phi_i \in \Sigma}\{\phi_i \vee \neg new_i, new_i\}, new_i\rangle$, where $\{new_1, \ldots, new_n\}$ are new symbols from $PS \setminus Var(\Sigma)$. We have that ϕ_i belongs to a conflict of Σ if and only if new_i belongs to a minimal var-conflict of $\bigcup_{\phi_i \in \Sigma}\{\phi_i \vee \neg new_i, new_i\}$. Indeed, $S \subseteq \Sigma$ is a conflict of Σ if and only if $S_{new} = \{\phi_j \vee \neg new_j, new_j \mid \phi_j \in S\}$ is a conflict of $\bigcup_{\phi_i \in \Sigma}\{\phi_i \vee \neg new_i, new_i\}$ (intuitively, this is just a naming operation, hence every conflict is mainly preserved and no new conflict is added). The point is that every conflict S_{new} of $\bigcup_{\phi_i \in \Sigma}\{\phi_i \vee \neg new_i, new_i\}$ gives rise to a *minimal* var-conflict $Var(S_{new})$ of it by construction (this is not the case in general). The Π_2^p-hardness result stated in (1) completes the proof.

In order to prove (3), we first prove the two following lemmata:

Lemma 3. *Let $\Sigma = \{\phi_1, \ldots, \phi_n\}$ be a CNF formula and $p \in Var(\Sigma)$. p belongs to a minimal var-conflict of Σ if and only if there exists a 3-model I of Σ that is minimal w.r.t. $\leq_{LP_m}$ in its set and s.t. $I(p) = \top$.*

Proof of the lemma: We assume that no clause ϕ_i in Σ contains both a literal and its negation. This assumption can be done without loss of generality since removing such classical tautologies from Σ does not change the set of its conflicts, hence the set of its minimal var-conflicts and does not change the set of its 3-models (every tautological clause is interpreted as 1 or $\top$ in *THREE*).

- "if" way: Let I be a 3-model of Σ that is minimal w.r.t. $\leq_{LP_m}$ in its set and s.t. $I(p) = \top$. Let us show that p belongs to a minimal var-conflict of Σ. Let $\Sigma_p =_{def} \{\phi_i \in \Sigma \mid \forall x \in Var(\phi_i)$ if $x \neq p$, then $I(x) \neq \top\}$ be the set of clauses from Σ whose symbols are classically interpreted in I, expect possibly for p. Σ_p can be partitioned into three sets:
  - The subset Σ_p^+ of Σ_p containing all the clauses in which p appears as a literal.
  - The subset Σ_p^- of Σ_p containing all the clauses in which $\neg p$ appears as a literal.
  - The subset Σ_p^* of Σ_p containing all the clauses in which p does not appear (as a symbol).

  Since I is minimal w.r.t. $\leq_{LP_m}$ in the set of 3-models of Σ and $I(p) = \top$, it must be the case that Σ_p is classically inconsistent (otherwise there would exist a 2-model J of Σ_p^+ and the 3-interpretation J' defined by $J'(x) = J(x)$ for every $x \in Var(\Sigma_p^+)$ and $J'(x) = I(x)$ otherwise would be a 3-model of Σ s.t. $J' \leq_{LP_m} I$ and $I \not\leq_{LP_m} J'$: The minimality of I would be questioned).

Accordingly, there exists at least one conflict of Σ_p, hence at least one conflict of Σ.
Now, Σ_p^* is classically satisfiable since the restriction of I to $Var(\Sigma_p^*)$ is a (partial but classical) model of it. Since p (resp. $\neg p$) is a *pure* literal in Σ_p^+ (resp. Σ_p^-) (i.e., $\neg p$ (resp. p) does not appear as a literal in any clause of the set), every conflict of Σ_p contains at least one clause from Σ_p^- and one clause from Σ_p^+. In particular, every conflict of Σ_p is s.t. $p \in Var(S)$.
It remains to show that p belongs to every minimal var-conflict of Σ. If it were not the case, for every conflict S of Σ_p, there would exist a conflict S' of Σ s.t. $p \notin Var(S')$ and $Var(S') \subset Var(S)$. Let ϕ_i be any clause from $\Sigma \setminus \Sigma_p$. By construction, there exists at least one propositional symbol $x_i \neq p$ occurring in ϕ_i s.t. $I(x_i) = \top$ and $x_i \notin Var(\Sigma_p)$. Every conflict S' containing ϕ_i is s.t. $x_i \in Var(S')$, and this prevents $Var(S') \subset Var(S)$ from being true since $Var(S) \subseteq Var(\Sigma_p)$ whenever S is a conflict of Σ_p. Accordingly, every conflict S' of Σ s.t. $Var(S') \subset Var(S)$ does not contain any clause from $\Sigma \setminus \Sigma_p$; in other words, S' is a conflict of Σ_p, hence it must satisfy $p \in Var(S')$, contradiction.

– "only-if" way: Let $VConf(\Sigma)$ be the set of all minimal var-conflicts of Σ. Let us first recall the following notion of a *minimal hitting set* (cf. [66]):

Definition 20 (minimal hitting set). *Let E be a finite set and C a set of subsets of E (i.e., $C \subseteq 2^E$). Let I be a subset of E.*

- *I is a* hitting set *of C if and only if $\forall S \in C, I \cap S \neq \emptyset$.*
- *I is a* minimal hitting set *of C if and only if I is a hitting set of C and no proper subset of I is a hitting set of C.*

If p is a propositional symbol belonging to a minimal var-conflict of Σ, then there exists a minimal hitting set H_p of $VConf(\Sigma)$ s.t. $p \in H_p$ (the minimality requirement is important here, the conclusion would not hold if any var-conflict were considered).
Once this is observed, it is sufficient to show that for every minimal hitting set H of $VConf(\Sigma)$, there exists a 3-model I_H of Σ that is minimal w.r.t. $\leq_{LP_m}$ and s.t. $\forall x \in Var(\Sigma)$, $I(x) = \top$ if and only if $x \in H$.
Indeed, let $\Sigma_H^* = \{\phi_i \in \Sigma \mid Var(\phi_i) \cap H = \emptyset\}$. Let us show that Σ_H^* is classically satisfiable. If it were not the case, Σ_H^* would include a conflict S so there would exist a minimal var-conflict V of Σ_H^* s.t. $V \subseteq Var(S)$. This means that there would exist at least one conflict S_V of Σ_H^* s.t. $V = Var(S_V)$. Since $\Sigma_H^* \subset \Sigma$, S_V would be a conflict of Σ, there would also exist a minimal var-conflict V' of Σ s.t. $V' \subseteq V$. Since H is a minimal hitting set of $VConf(\Sigma)$ and $V' \in VConf(\Sigma)$, there exists a variable x s.t. $x \in V'$ and $x \in H$. Since $V' \subseteq V \subseteq Var(S)$, we must have $x \in Var(S)$ but since $S \subseteq \Sigma_H^*$ and $x \in H$, no clause of S can contain x as a symbol, contradiction. We conclude that a 2-model I of Σ_H exists.
Now, since every clause ϕ_i from Σ except those from Σ_H^* contains at least one symbol from H (by definition of Σ_H^*), every 3-interpretation J s.t. $J(x) = \top$

for every $x \in H$ also satisfies $J(\phi_i) = \top$. Since $H \cap Var(\Sigma^*_H) = \emptyset$, the 3-interpretation I_H defined by $I_H(x) = \top$ whenever $x \in H$ and $I_H(x) = I(x)$ otherwise is well-defined. By construction, I_H is a 3-model of Σ.
It remains to show that I_H is minimal w.r.t. $\leq_{LP_m}$. This comes directly from the fact that H is a *minimal* hitting set of $VConf(\Sigma)$. Indeed, let us assume that I_H is not a minimal 3-model of Σ w.r.t. $\leq_{LP_m}$. Then there exists $x \in H$ and a 3-model $I_{H'}$ of Σ minimal w.r.t. $\leq_{LP_m}$ s.t. for every $y \in Var(\Sigma)$, if $I_{H'}(y) = \top$, then $I_H(y) = \top$, and $I_{H'}(x) \neq \top$. This implies that the set H' of propositional symbols interpreted as $\top$ in $I_{H'}$ is a proper subset of H. Since H is a minimal hitting set of $VConf(\Sigma)$, this implies that there exists at least one minimal var-conflict $V \in VConf(\Sigma)$ s.t. $H' \cap V = \emptyset$. This means that no conflict S_V of Σ s.t. $Var(S_V) = V$ contains an element of H' as a symbol. Since $I_{H'}$ restricted to the symbols outside H' behaves classically and $I_{H'}$ is a 3-model of Σ, $I_{H'}$ is a classical model of S_V. This contradicts the fact that S_V is a conflict.

■

Lemma 4. *Let $\Sigma = \{\phi_1, \ldots, \phi_n\}$ be a CNF formula and $p \in Var(\Sigma)$. Let $new \in PS \setminus Var(\Sigma)$. There exists a 3-model I of Σ that is minimal w.r.t. $\leq_{LP_m}$ in its set and s.t. $I(p) = \top$ if and only if $\Sigma \cup \{p \vee new, \neg p \vee new\} \not\models_{LP_m} new$ if and only if $\Sigma \cup \{p \vee new, \neg p \vee new\} \not\vdash_s new$.*

Proof of the lemma:

- We first prove that there exists a 3-model I of Σ that is minimal w.r.t. $\leq_{LP_m}$ in its set and s.t. $I(p) = \top$ if and only if $\Sigma \cup \{p \vee new, \neg p \vee new\} \not\models_{LP_m} new$. It is obvious that the expansion of Σ with the two clauses $p \vee new$ and $\neg p \vee new$ does not create any new conflict, hence any new minimal var-conflict. Accordingly, the 3-models of Σ that are minimal w.r.t. $\leq_{LP_m}$ are exactly the 3-models of $\Sigma \cup \{p \vee new, \neg p \vee new\}$ that are minimal w.r.t. $\leq_{LP_m}$ (see the "only-if" side of the proof of Lemma 3). Let I be a 3-model of Σ that is minimal w.r.t. $\leq_{LP_m}$ in its set and s.t. $I(p) = \top$. Since $new \notin Var(\Sigma)$, the minimality requirement ensures that there exists a 3-model I_0 of Σ that is minimal w.r.t. $\leq_{LP_m}$ and s.t. $I_0(new) = 0$. Since I_0 also is a 3-model of $\Sigma \cup \{p \vee new, \neg p \vee new\}$ that is minimal w.r.t. $\leq_{LP_m}$, it cannot be the case that $\Sigma \cup \{p \vee new, \neg p \vee new\} \models^2 new$. Conversely, let I_0 be a 3-model of $\Sigma \cup \{p \vee new, \neg p \vee new\}$ that is minimal w.r.t. $\leq_{LP_m}$ and s.t. $I_0(new) = 0$. Since I_0 is a 3-model of both $p \vee new$ and $\neg p \vee new$, it must be the case that I_0 is a 3-model of both p and $\neg p$, which requires that $I_0(p) = \top$. Since I_0 also is a 3-model of Σ that is minimal w.r.t. $\leq_{LP_m}$, the conclusion follows.
- Let us now prove that $\Sigma' \not\models_{LP_m} new$ if and only if $\Sigma' \not\vdash_s new$, where $\Sigma' = \Sigma \cup \{p \vee new, \neg p \vee new\}$. Let us first prove the following lemma:

  Lemma 5. *Let Σ and γ be two formulas from $PROP^2_{PS}$.*[9] *We have $\Sigma \vdash^{\pm}_s \gamma$ if and only if $\Sigma \models_{LP_m} \gamma$.*

[9] Note that $\supset$ and $\Leftrightarrow$ are considered as syntactic sugars here ($\phi \supset \psi =_{def} (\neg\phi) \vee \psi$ and $\phi \Leftrightarrow \psi =_{def} (\phi \supset \psi) \wedge (\psi \supset \phi)$).

Proof of the lemma: By definition, γ is a skeptical signed consequence of Σ if and only if $\gamma^{\pm}$ belongs to every extension of $\langle \Sigma^{\pm}, D_{\Sigma} \rangle$. From Theorems A.1 and A.2 of [38], this is equivalent to state that $\gamma^{\pm}$ belongs to every extension of supernormal default theory $\langle \Sigma^{\pm}, \{x^{+} \Leftrightarrow \neg x^{-} \mid x \in PS\} \rangle$. Now, from the well-known preferred models characterization of skeptical inference, this is still equivalent to state that every preferred (2-)model of $\Sigma^{\pm}$ is a 2-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 $TWO^{PS^{\pm}}$ defined by $I \leq J$ if and only if $\{x \in PS \mid I(x^{+}) = I(x^{-})\} \subseteq \{x \in PS \mid J(x^{+}) = J(x^{-})\}$ (indeed, I satisfies $x^{+} \Leftrightarrow \neg x^{-}$ if and only if $I(x^{+}) \neq I(x^{-})$).
By construction, $\Sigma^{\pm}$ is a monotone (positive) formula from $PROP^{2}_{PS^{\pm}}$. As a consequence, if I is a 2-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 propositional symbol from $PS^{\pm}$ (except possibly on x^{+} and x^{-}) is still a 2-model of $\Sigma^{\pm}$. This is specifically the case for any 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 symbol $x \in PS$, we do *not* have $M(x^{+}) = M(x^{-}) = 0$. Denoting $(TWO^{PS^{\pm}})^{*}$ the subset of all interpretations I from $TWO^{PS^{\pm}}$ 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 $(TWO^{PS^{\pm}})^{*}$.
Now, every interpretation I from $(TWO^{PS^{\pm}})^{*}$ 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 structural induction on $\phi \in PROP^{2}_{PS}$) that I is a 2-model of the monotone (positive) formula $\phi^{\pm}$ from $PROP^{2}_{PS^{\pm}}$ if and only if $3(I)$ is a 3-model of ϕ.
Finally, we know that $\Sigma \models_{LP_m} \gamma$ holds if and only if 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. $\leq_{LP_m}$. Since every preferred model of $\Sigma^{\pm}$ can be associated to a 3-model of Σ through the $3(.)$ bijection, and since for every pair of interpretations I and J from $(TWO^{PS^{\pm}})^{*}$, we have $I \leq J$ if and only if $3(I) \leq_{LP_m} 3(J)$, we obtain that M is a preferred model of $\Sigma^{\pm}$ w.r.t. $\leq$ if and only if $3(M)$ is a preferred 3-model of Σ w.r.t. $\leq_{LP_m}$, which concludes the proof. ∎
Thanks to this lemma, we know that we have that $\Sigma' \not\models_{LP_m} new$ if and only if $\Sigma' \not\vdash_{s} new$, where $\Sigma' = \Sigma \cup \{p \vee new, \neg p \vee new\}$ if and only if we have that $\Sigma' \not\vdash^{\pm}_{s} new$ if and only if $\Sigma' \not\vdash_{s} new$. This equivalence is quite obvious since new and new^{+} always appear in the same extensions of $\langle \Sigma'^{\pm}, D_{\Sigma'} \rangle$; indeed, every extension of $\langle \Sigma'^{\pm}, D_{\Sigma'} \rangle$ contains the consequent formula $(new \Leftrightarrow new^{+}) \wedge (\neg new \Leftrightarrow new^{-})$ because the corresponding justification $(new^{+} \Leftrightarrow \neg new^{-})$ is (classically) consistent with Σ' (no occurrence of $\neg new$ appears in Σ').

∎

(3) is a straightforward consequence of (2), Lemma 3 and Lemma 4.

The hardness of $\models_{BS}$ finally remains to be proven. Let us reduce the inference problem of LP_m to the inference problem for the logic of [37]. Let Σ be a CNF formula and γ a symbol. We associate to $\langle \Sigma, \gamma \rangle$ in polynomial time the following ordered pair $\langle \{\Sigma \wedge new \wedge \neg new\} \cup \bigcup_{p \in Var(\Sigma) \cup Var(\gamma)} \{p \vee \neg p\}, \gamma \rangle$, where $new \in PS \setminus (Var(\Sigma) \cup Var(\gamma))$. We have $\Sigma \models_{LP_m} \gamma$ if and only if $\{\Sigma \wedge new \wedge \neg new\} \cup \bigcup_{p \in Var(\Sigma) \cup Var(\gamma)} \{p \vee \neg p\} \models_{BS} \gamma$. Indeed, since new does not occur elsewhere in Σ or γ, we have $\Sigma \models_{LP_m} \gamma$ if and only if $\Sigma \wedge new \wedge \neg new \models_{LP_m} \gamma$. Accordingly, it is sufficient to consider the 3-models I of Σ s.t. $I(new) = \top$. Once this is stated, let us observe that the set S of 3-models I of Σ over $Var(\Sigma) \cup Var(\gamma) \cup \{new\}$ s.t. $I(new) = \top$ coincide with the set of 3-models of $\{\Sigma \wedge new \wedge \neg new\} \cup \bigcup_{p \in Var(\Sigma) \cup Var(\gamma)} \{p \vee \neg p\}$ over $Var(\Sigma) \cup Var(\gamma) \cup \{new\}$. Indeed, for every propositional symbol p and every 3-interpretation I, we always have $I(p \vee \neg p) \in \{1, \top\}$. Now, it remains to show that $\leq_{LP_m}$ and $\leq_{BS}$ coincides over S. To do it, it is sufficient to remark that for every $I \in S$, (1) we have $I(\Sigma \wedge new \wedge \neg new) = \top$ (this explains why new has been introduced), and (2) for every $p \in Var(\Sigma) \cup Var(\gamma)$, we have $I(p) = \top$ if and only if $I(p \vee \neg p) = \top$. This completes the proof of Π_2^p-hardness of $\models_{BS}$-inference in the case when the belief base is a finite set of CNF formulas and the query a symbol.

Hardness for $\models_{\leq_1}^{4,arg}$ and $\models_{\leq_2}^{4,arg}$ comes from the hardness result for $\models_{\leq_{LP_m}}^{3,arg}$ and the fact that $\models_{\leq_1}^{4}$, $\models_{\leq_2}^{4}$ and $\models_{\leq_{LP_m}}^{3}$ coincide when Σ is CNF and γ a literal (see Proposition 5.4 from [39]). Since $\models_{\leq_{LP_m}}^{3}$ is Π_2^p-hard in this situation, it is sufficient to point out a polytime reduction that maps every tuple $\langle \Sigma_1, p_1, \Sigma_2, p_2 \rangle$ (where Σ_1 and Σ_2 are CNF formulas and p_1 and p_2 are propositional symbols) to a pair $\langle \Sigma, \gamma \rangle$ (where Σ and γ are CNF formulas) s.t. $\Sigma \models_{\leq_{LP_m}}^{3,arg} \gamma$ if and only if $\Sigma_1 \models_{\leq_{LP_m}}^{3} p_1$ and $\Sigma_2 \not\models_{\leq_{LP_m}}^{3} p_2$. Without loss of generality, we can assume that Σ_1 and Σ_2 share no propositional symbols. Let us take $\Sigma = \Sigma_1 \wedge \Sigma_2 \wedge new_2 \wedge \neg new_2$ and $\gamma = (p_1 \vee new_1) \wedge (\neg p_2 \vee new_2)$, where new_1 and new_2 are new propositional symbols, not occurring in $Var(\Sigma_1) \cup Var(\Sigma_2) \cup \{p_1, p_2\}$.

Let us first assume that $\Sigma_1 \models_{\leq_{LP_m}}^{3} p_1$ and $\Sigma_2 \not\models_{\leq_{LP_m}}^{3} p_2$. If $\Sigma_1 \models_{\leq_{LP_m}}^{3} p_1$, then $\Sigma \models_{\leq_{LP_m}}^{3} p_1 \vee new_1$ because $Var(\Sigma_1) \cap Var(\Sigma_2) = \emptyset$; furthermore, we necessarily have $\Sigma \models_{\leq_{LP_m}}^{3} \neg p_2 \vee new_2$ because, by construction, every 3-model I of Σ is s.t. $I(new_2) = \top$; hence, $\Sigma \models_{\leq_{LP_m}}^{3} \gamma$. It remains to show that $\Sigma \not\models_{\leq_{LP_m}}^{3} \neg\gamma$. We have $\neg\gamma \equiv^3 (\neg p_1 \vee \neg new_2) \wedge (\neg new_1 \vee \neg new_2) \wedge (\neg p_1 \vee p_2) \wedge (\neg new_1 \vee p_2)$. Since $\Sigma_2 \not\models_{\leq_{LP_m}}^{3} p_2$, there exists a $\leq_{LP_m}$-preferred 3-model I_2 of Σ_2 s.t. $I_2(p_2) = 0$; since $new_1 \notin Var(\Sigma_2)$, the 3-interpretation I_2' that coincides with I_2 on every symbol, except possibly on new_1 and satisfies $I_2'(new_1) = 1$, also is a $\leq_{LP_m}$-preferred 3-model of Σ_2. Accordingly, $I_2'(\neg new_1 \vee p_2) = 0$ and $\Sigma_2 \not\models_{\leq_{LP_m}}^{3} \neg new_1 \vee p_2$; now, since $Var(\Sigma_1) \cap Var(\Sigma_2) = \emptyset$ and new_1 and new_2 are new propositional symbols, we obtain that $\Sigma \not\models_{\leq_{LP_m}}^{3} \neg new_1 \vee p_2$, hence $\Sigma \not\models_{\leq_{LP_m}}^{3} \neg\gamma$.

Conversely, let us assume that $\Sigma \models_{\leq_{LP_m}}^{3,arg} \gamma$ and prove that $\Sigma_1 \models_{\leq_{LP_m}}^{3} p_1$ and $\Sigma_2 \not\models_{\leq_{LP_m}}^{3} p_2$. If $\Sigma \models_{\leq_{LP_m}}^{3} p_1 \vee new_1$, then it must be the case that $\Sigma_1 \models_{\leq_{LP_m}}^{3} p_1$; otherwise, there would exist a $\leq_{LP_m}$-preferred 3-model I_1 of Σ_1

s.t. $I_1(p_1) = 0$; since $new_1 \notin Var(\Sigma_1)$, the 3-interpretation I_1' that coincides with I_1 on every symbol, except possibly on new_1 and satisfies $I_1'(new_1) = 0$, also is a $\leq_{LP_m}$-preferred 3-model of Σ_1. Accordingly, $I_1'(p_1 \vee new_1) = 0$, and since $Var(\Sigma_1) \cap Var(\Sigma_2) = \emptyset$ and new_1 and new_2 are new propositional symbols, we would have $\Sigma \not\models^3_{\leq_{LP_m}} p_1 \vee new_1$, contradiction. Finally, it is the case that $\Sigma_2 \not\models^3_{\leq_{LP_m}} p_2$. Otherwise, we would have $\Sigma_2 \models^3_{\leq_{LP_m}} (\neg p_1 \vee p_2) \wedge (\neg new_1 \vee p_2)$, hence $\Sigma \models^3_{\leq_{LP_m}} (\neg p_1 \vee p_2) \wedge (\neg new_1 \vee p_2)$. Since by construction every 3-model I of Σ is s.t. $I(new_2) = \top$, we also have $\Sigma \models^3_{\leq_{LP_m}} (\neg p_1 \vee \neg new_2) \wedge (\neg new_1 \vee \neg new_2)$. As a consequence, we obtain that $\Sigma \models^3_{\leq_{LP_m}} \neg\gamma$, contradiction.

The complexity of $\models_{QC}$ and $\models_{Q\Pi L}$ has been identified in [54] and [62], respectively, when both Σ and γ are CNF formulas:

Proposition 3. *The complexity of the decision problem for $\models_{QC}$ and $\models_{Q\Pi L}$ is* coNP*-complete when both Σ and γ are CNF formulas.*

The membership proof is by a reduction to classical entailment (more specifically, by a modular and faithful polytime translation to classical logic, see Propositions 1 and 2 in [54]).

Tractability results As to $\models^4$ or $\models^3$, tractability is achieved as soon as the query is a CNF formula when the belief base belongs to the monotonic fragment [52]. Tractability is still the case for $\models^4$ when the base is a DNF formula, while $\models^3$ remains coNP-complete in this situation [52].

Other tractability results for $\models^3_S$ and the $\approx^{\mathcal{P},S_0}_S$ relations can be obtained by imposing in addition that $|S|$ is bounded (see Theorem A.3 in [51] and Proposition 3.15 in [42]).

Proposition 4. *The complexity of the decision problem for $\vdash_L$ is in* P *in the case when Σ and γ are CNF formulas (from $PROP^2_{PS}$) and the size of $S \subseteq PS$ is bounded, when $\vdash_L$ is $\models^3_S$, $\approx^{\mathcal{PO},S_0}_S$, $\approx^{\mathcal{LO},S_0}_S$, $\approx^{\mathcal{IP},S_0}_S$ or $\approx^{\mathcal{LE},S_0}_S$.*

Clearly enough, under the CNF assumption – a strict subset of the monotonic fragment –, the key for tractability is the fact that the size of S is bounded (hence only polynomially many S-3-interpretations are to be considered).

Finally, the translation from the decision problem for $\models_{QC}$ (in the CNF case) to the decision problem for classical entailment reported in [54] gives two tractable subcases as by-products: The decision problem for $\models_{QC}$ for CNF queries γ is in P given that the base Σ is a Krom formula or a renamable Horn k-CNF formula (see Propositions 4 and 5 in [54]).

5 Conclusion

The main purpose of this chapter was to put together in a coherent way complexity results for paraconsistent inference relations defined in multi-valued propositional settings (or related frameworks), and already identified in the literature.

In a nutshell, the corresponding decision problems are located at the first or the second level of the polynomial hierachy; to be more precise, the relations under consideration are either as hard as classical entailment (coNP), or mildly harder ($\Delta_2^p[\mathcal{O}(log\ n)]$ or Δ_2^p), or finally strongly harder (Π_2^p-complete or D_2^p-complete) (under the usual assumptions of complexity theory); only few tractable cases have been exhibited.

The complexity results above can be related to similar complexity results for belief revision [43, 44] and reasoning from preferred consistent subsets in [44, 18, 46]. Actually,the Π_2^p-hardness of the problem of determining whether a given formula is a consequence of a belief base revised according to the WIDTIO policy – or equivalently, whether it is a so-called free consequence of a stratified belief base – has been exploited in some of the hardness proofs for non-monotonic relations. Conversely, some reductions to inference problems considered in belief revision or when reasoning from preferred consistent subsets can be used to prove membership results (even if we did not systematically adhere to such an approach when presenting membership results).

Roughly, the derived results show the complexity of paraconsistent inference to vary from P in some restricted cases to the first level of the polynomial hierarchy for many monotonic relations and even to the second level of the polynomial hierarchy for the non-monotonic ones. This reflects that there is typically a computational price to be paid for preserving more beliefs through the exploitation of preferential information (contrariwise to the model checking problem when no preferences are considered, determining whether a given interpretation is a preferred model of a belief base is intractable). It is interesting to observe that tractability is obtained in restricted situations when classical entailment is not tractable (under the standard assumptions of complexity theory). For instance, $\models^4$, $\models^3$ and its restrictions are tractable when the belief base is a CNF formula and clausal queries are considered, only. Indeed, restricting the language to its clausal fragment is sufficient to lower the complexity of many monotonic relations (this can be easily explained by the fact that each clause from the belief base can be considered independent to the other ones in the sense that the disjunctive syllogism fails). Contrastingly, focusing on the clausal fragment is not sufficient to decrease the complexity of the non-monotonic relations.

Among the inference relations considered before, $\approx_S^{\mathcal{LE},S_0}$ and $\models_{QC}$ (when CNF formulas are considered) appear as achieving an interesting compromise in the sense that many expected consequences are preserved while complexity remains at the first level of the polynomial hierarchy. Furthermore, $\approx_S^{\mathcal{P},S_0}$ relations allow for resource-bounded reasoning by limiting the size of S.

Much work remains to be done in this area. First, the complexity of many propositional paraconsistent inference relations is still unknown; this is particularly the case for many proof-theoretically defined relations.

Second, the tractability islands are not so numerous, especially for the most sophisticated relations that have been considered, so it is important to identify ways to deal with the complexity of paraconsistent inference relations. In such an objective, one can improve existing inference engines so as to increase the set of instances that are feasible from the practical side. Importantly, the complexity results we have reported show the existence of polynomial reductions from decision problems for paraconsistent inference to complete problems at the first or at the second level of the polynomial hierarchy; among the latter problems are, on the one hand, validity problems for quantified boolean formulas (QBFs) – including UNSAT as a specific case – and on the other hand, for the problems located at the second level of the polynomial hierarchy, inference problems for many non-monotonic logics considered in AI (e.g., skeptical reasoning from default theories, forms of closed-world reasoning) for which automated deduction tools have been developed). If such reductions are practically feasible, it is possible to take advantage of existing QBFs solvers and other programs for non-monotonic inference to implement forms of paraconsistent inference. Observing that QBFs solvers are more and more efficient in practice over years, this possibility is particularly interesting. In fact, reductions (and among them, modular and faithful translations) have already been exploited to encode several non-monotonic paraconsistent inference problems into validity problems for QBFs or circumscriptive inference problems (see [55, 67, 68]). Of course, it would be nice to compare from the empirical side the performances of such reduction-based approaches with the ones achieved by implementations of some proof systems (usually, tableau-based) for paraconsistent inference.

Another approach consists in circumventing complexity, and three main techniques prevail: Restriction, approximation and compilation. We have mainly focused on the first technique, showing for instance that quasi-classical inference is tractable under clausal queries when the belief base is a Krom formula or a renamable Horn k-CNF formula. Many other tractable fragments for classical satisfiability appear as interesting candidates to be investigated. The second technique is at work in [42]; approximation is used for reaching two goals: Considering $\models^3_S$ with a bounded S instead of its superset $\models^2$ allows for ensuring tractability, while considering $\approx^{\mathcal{LE},S_0}_S$ instead of its superset $\models^3_S$ is sufficient for avoiding trivialization when it would occur. Similar ideas could be applied to other inference relations. Finally, it seems that the last technique mentioned above (compilation) has almost never been used in the purpose of enhancing paraconsistent inference; the principle would be to pre-process the belief base during an off-line compilation stage so as to obtain better performances when queries are considered at the on-line stage. In the case when a translation to inference problems from stratified belief bases exist (as for the $\approx^{\mathcal{P},S_0}_S$ relations), it is possible to take advantage of some recent work about the compilation of stratified belief bases (see in particular [69–71]). Developing other compilation techniques for paraconsistent inference or fitting the compilation techniques for classical inference to paraconsistent inference is another perspective for further research.

References

1. Gärdenfors, P.: Belief revision and nonmonotonic logic: Two sides of the same coin? In: Proceedings of the 9^{th} European Conference on Artificial Intelligence,

Stockholm (Sweden) (1990) 768–773
2. Gärdenfors, P., Makinson, D.: Relations between the logic of theory change and nonmonotonic logic. In: The Logic of Theory Change. (1990) 185–205
3. Besnard, P., Hunter, A.: Introduction to actual and potential contradictions. In: Handbook of Defeasible Reasoning and Uncertainty Management Systems. Volume 2. Kluwer Academic (1998) 1–11
4. Hunter, A.: Paraconsistent logics. In: Handbook of Defeasible Reasoning and Uncertainty Management Systems. Volume 2. Kluwer Academic (1998) 11–36
5. Priest, G.: Paraconsistent Logic. In: Handbook of Philosophical Logic. Volume 6. Kluwer Academic (2002) 287–393
6. Grant, J., Subrahmanian, V.: Reasoning in inconsistent knowledge bases. IEEE Transactions on Knowledge and Data Engineering **7** (1995) 177–189
7. Lin, J.: Integration of weighted knowledge bases. Artificial Intelligence **83** (1996) 363–378
8. Revesz, P.: On the semantics of arbitration. International Journal of Algebra and Computation **7** (1997) 133–160
9. Konieczny, S., Pino Pérez, R.: On the logic of merging. In: Proceedings of the 6^{th} International Conference on Knowledge Representation and Reasoning (KR'98), Trento (Italy) (1998) 488–498
10. Konieczny, S.: On the difference between merging knowledge bases and combining them. In: Proceedings of the 7^{th} International Conference on Knowledge Representation and Reasoning (KR'00), Breckenridge (CO) (2000) 135–144
11. Bloch, I., Lang, J.: Towards mathematical morpho-logics. In: Proceedings of the 8^{th} International Conference on Information Processing and Management of Uncertainty in Knowledge based Systems (IPMU'00), Madrid (Spain) (2000) 1405–1412
12. Rescher, N., Manor, R.: On inference from inconsistent premises. Theory and Decision **1** (1970) 179–219
13. Fagin, R., Ullman, J., Vardi, M.: On the semantics of updates in databases. In: Proceedings of the 2^{nd} ACM Symposium on Principles of Database Systems (PODS'83). (1983) 352–355
14. Ginsberg, M.: Counterfactuals. Artificial Intelligence **30** (1986) 35–79
15. Baral, C., Kraus, S., Minker, J.: Combining multiple knowledge bases. IEEE Transactions on Knowledge and Data Engineering **3** (1991) 208–220
16. Pinkas, G., Loui, R.: Reasoning from inconsistency: A taxonomy of principles for resolving conflict. In: Proceedings of the 3^{rd} International Conference on Knowledge Representation and Reasoning (KR'92), Cambridge (MA) (1992) 709–719
17. Benferhat, S., Cayrol, C., Dubois, D., Lang, J., Prade, H.: Inconsistency management and prioritized syntax-based entailment. In: Proceedings of the 13^{th} International Joint Conference on Artificial Intelligence (IJCAI'93), Chambéry (France) (1993) 640–645
18. Cayrol, C., Lagasquie-Schiex, M.C., Schiex, T.: Nonmonotonic reasoning: From complexity to algorithms. Annals of Mathematics and Artificial Intelligence **22** (1998) 207–236
19. Nebel, B.: Belief revision. In Gärdenfors, P., ed.: Cambridge Tracts in Theoretical Computer Science. Volume 29. Cambridge University Press, Cambridge (1992) 52–88
20. Nebel, B.: Base revision operations and schemes: Semantics, representation and complexity. In: Proceedings of the 11^{th} European Conference on Artificial Intelligence, Amsterdam (Netherlands) (1994) 341–345

21. Brewka, G.: Preferred subtheories: An extended logical framework for default reasoning. In: Proceedings of the 11^{th} International Joint Conference on Artificial Intelligence (IJCAI'89), Detroit (MI) (1989) 1043–1048
22. Lang, J., Marquis, P.: Resolving inconsistencies by variable forgetting. In: Proceedings of the 8^{th} International Conference on Knowledge Representation and Reasoning (KR'02), Toulouse (France) (2002) 239–250
23. Dung, P.: On the acceptability of arguments and its fundamental role in nonmonotonic reasoning, logic programming and n-person games. Artificial Intelligence **77** (1995) 321–357
24. Elvang-Goransson, M., Hunter, A.: Argumentative logics: Reasoning from classically inconsistent information. Data and Knowledge Engineering **16** (1995) 125–145
25. Bondarenko, A., Dung, P., Kowalski, R., Toni, F.: An abstract, argumentation-theoretic framework for default reasoning. Artificial Intelligence **93** (1997) 63–101
26. Amgoud, L., Cayrol, C.: On the acceptability of arguments in preference-based argumentation framework. In: Proceedings of the 14^{th} Conference on Uncertainty in Artificial Intelligence, Madison (WI) (1998) 1–7
27. Jaśkowski, S.: Propositional calculus for contradictory deductive systems. Studia Logica **24** (1969) 143–167
28. da Costa, N.C.A.: On the theory of inconsistent formal systems. Notre Dame Journal of Formal Logic **15** (1974) 497–510
29. Besnard, P., Hunter, A.: Quasi-classical logic: Non-trivializable classical reasoning from inconsistent information. In: Proceedings of the European Conference on Symbolic and Quantitative Approaches to Reasoning and Uncertainty. Volume 946 of LNAI., Fribourg (Switzerland), Springer-Verlag (1995) 44–51
30. Hunter, A.: Reasoning with contradictory information using quasi-classical logic. Journal of Logic and Computation **10** (2000) 677–703
31. D'Ottaviano, I., da Costa, N.: Sur un problème de Jaśkowski. Technical report, Comptes Rendus de l'Académie des Sciences de Paris (1970)
32. Belnap, N.: A useful four-valued logic. In: Modern Uses of Multiple-Valued Logic. Reidel (1977) 8–37
33. Frisch, A.: Inference without chaining. In: Proceedings of the 10^{th} International Joint Conference on Artificial Intelligence (IJCAI'87), Milan (Italy) (1987) 515–519
34. Levesque, H.: A knowledge-level account of abduction (preliminary version). In: Proceedings of the 11^{th} International Joint Conference on Artificial Intelligence (IJCAI'89), Detroit (MI) (1989) 1061–1067
35. Priest, G.: Reasoning about truth. Artificial Intelligence **39** (1989) 231–244
36. Priest, G.: Minimally inconsistent LP. Studia Logica **50** (1991) 321–331
37. Besnard, P., Schaub, T.: Circumscribing inconsistency. In: Proceedings of the 15^{th} International Joint Conference on Artificial Intelligence (IJCAI'97), Nagoya (Japan) (1997) 150–155
38. Besnard, P., Schaub, T.: Signed systems for paraconsistent reasoning. Journal of Automated Reasoning **20** (1998) 191–213
39. Arieli, O., Avron, A.: The value of four values. Artificial Intelligence **102** (1998) 97–141
40. Arieli, O., Avron, A.: A model-theoretic approach for recovering consistent data from inconsistent knowledge bases. Journal of Automated Reasoning **22** (1999) 263–309
41. Konieczny, S., Marquis, P.: Three-valued logics for inconsistency handling. In: Proceedings of the 8^{th} European Conference on Logics in Artificial Intelligence

(JELIA'02). Volume 2424 of Lecture Notes on Artificial Intelligence., Cosenza (Italy), Springer-Verlag (2002) 332–344
42. Marquis, P., Porquet, N.: Resource-bounded paraconsistent inference. Annals of Mathematics and Artificial Intelligence **39** (2003) 349–384
43. Eiter, T., Gottlob, G.: On the complexity of propositional knowledge base revision, updates, and counterfactuals. Artificial Intelligence **57** (1992) 227–270
44. Nebel, B.: Belief revision. In Dubois, D., Prade, H., eds.: Handbook of Defeasible Reasoning and Uncertainty Management Systems. Volume 3. Kluwer Academic (1998) 77–145
45. Konieczny, S., Lang, J., Marquis, P.: DA^2 merging operators. Artificial Intelligence **157** (2004) 49–79
46. Eiter, T., Lukasiewicz, T.: Default reasoning from conditional knowledge bases: Complexity and tractable cases. Artificial Intelligence **124** (2000) 169–241
47. Dimopoulos, Y., Nebel, B., Toni, F.: Preferred arguments are harder to compute than stable extensions. In: Proceedings of the 16^{th} International Joint Conference on Artificial Intelligence (IJCAI'99), Stockholm (Sweden) (1999) 36–41
48. Dimopolos, Y., Nebel, B., Toni., F.: Finding admissible and preferred arguments can be very hard. In: Proceedings of the 7^{th} International Conference on Knowledge Representation and Reasoning (KR'00), Breckenridge (CO) (2000) 53–61
49. Dunne, P., Bench-Capon, T.: Coherence in finite argument systems. Artificial Intelligence **141** (2002) 187–203
50. Meyer, J.J.C., der Hoek, W.V.: Modal logics for representing incoherent knowledge. In: Handbook of Defeasible Reasoning and Uncertainty Management Systems. Volume 2. Kluwer Academic (1998) 37–75
51. Schaerf, M., Cadoli, M.: Tractable reasoning via approximation. Artificial Intelligence **74** (1995) 249–310
52. Cadoli, M., Schaerf, M.: On the complexity of entailment in propositional multivalued logics. Annals of Mathematics and Artificial Intelligence **18** (1996) 29–50
53. Coste-Marquis, S., Marquis, P.: Complexity results for paraconsistent inference relations. In: Proceedings of the 8^{th} International Conference on Knowledge Representation and Reasoning (KR'02), Toulouse (France) (2002) 61–72
54. Marquis, P., Porquet, N.: Computational aspects of quasi-classical entailment. Journal of Applied Non-Classical Logics **11** (2001) 295–312
55. Besnard, P., Schaub, T., Tompits, H., Woltran, S.: Paraconsistent reasoning via Quantified Boolean Formulas I: Axiomatizing signed systems. In: Proceedings of the 8^{th} European Conference on Logics in Artificial Intelligence (JELIA'02), Cosenza (Italy) (2002) 320–331
56. Enderton, H.: A mathematical introduction to logic. Academic Press, New York (1972)
57. Garey, M., Johnson, D.: Computers and intractability: A guide to the theory of NP-completeness. Freeman (1979)
58. Papadimitriou, C.H.: Computational complexity. Addison–Wesley (1994)
59. Epstein, R.: The Semantic Foundations of Logic. In: Propositional Logics. Volume 1. Kluwer Academic (1990)
60. Kraus, S., Lehmann, D., Magidor, M.: Nonmonotonic reasoning, preferential models and cumulative logics. Artificial Intelligence **44** (1990) 167–207
61. Benferhat, S., Dubois, D., Prade, H.: How to infer from inconsistent beliefs without revising. In: Proceedings of the 14^{th} International Joint Conference on Artificial Intelligence (IJCAI'95), Montreal (Canada) (1995) 1449–1455
62. Dubois, D., Konieczny, S., Prade, H.: Quasi-possibilistic logic and its measures of information and conflict. Fundamenta Informaticae **57** (2003) 101–125

63. Gottlob, G.: Complexity results for nonmonotonic logics. Journal of Logic and Computation **2** (1992) 397–425
64. Stillman, J.: The complexity of propositional default logics. In: Proceedings of the 10^{th} National Conference on Artificial Intelligence, San Jose (CA) (1992) 794–799
65. Winslett, M.: Updating logical databases. In: Cambridge Tracts in Theoretical Computer Science. Cambridge University Press (1990)
66. Reiter, R.: A theory of diagnosis from first principles. Artificial Intelligence **32** (1987) 57–95
67. Arieli, O., Denecker, M.: Reducing preferential paraconsistent reasoning to classical entailment. Journal of Logic and Computation **13** (2003) 557–580
68. Besnard, P., Schaub, T., Tompits, H., Woltran, S.: Paraconsistent reasoning via Quantified Boolean Formulas II: Circumscribing inconsistent theories. In: Proceedings of the 7^{th} European Conference on Symbolic and Quantitative Approaches to Reasoning with Uncertainty, Aalborg (Denmark) (2003) 528–539
69. Benferhat, S., Kaci, S., Berre, D.L., Williams, M.A.: Weakening conflicting information for iterated revision and knowledge integration. Artificial Intelligence **153** (2004) 339–371
70. Coste-Marquis, S., Marquis, P.: On stratified belief base compilation. Annals of Mathematics and Artificial Intelligence **42** (2004) 399–442
71. Darwiche, A., Marquis, P.: Compiling propositional weighted bases. Artificial Intelligence **157** (2004) 81–113