Three-valued Logics for Inconsistency Handling

Sébastien Konieczny[1] and Pierre Marquis[2]

[1] IRIT/CNRS - Université Paul Sabatier
118 route de Narbonne
31062 Toulouse - FRANCE
konieczny@irit.fr
[2] CRIL/CNRS - Université d'Artois
SP 16 - rue de l'Université
62300 Lens - FRANCE
marquis@cril.univ-artois.fr

Abstract. While three-valued paraconsistent logic is a valuable framework for reasoning under inconsistency, the corresponding basic inference relation is too cautious and fails in discriminating in a fine-grained way the set of expected consequences of belief bases. To address both issues, we point out more refined inference relations. We analyze them from the logical and computational points of view and we compare them with respect to their relative cautiousness.

1 Introduction

Inconsistency appears very often in actual, large sized belief bases used by intelligent systems such as autonomous robots or infobots. As a consequence, autonomous belief based agents have to handle inconsistency some way, to prevent inference from trivializing. Indeed, paraconsistency is acknowledged for a while as an important feature of common-sense reasoning, strongly connected to several central AI issues, like reasoning with exceptions and counterfactuals, and helpful for many applications (for instance, model-based diagnosis).

In the following, the focus is laid on *three-valued paraconsistent logic*. The additional third value (called middle element) intuitively means "both true and false" and allows to still reasoning meaningfully with variables that are not embedded directly in a contradiction.

Several inference relations can be defined in three-valued paraconsistent logic. For any of them, trivialization is avoided by weakening the classical entailment relation. Compared with other approaches to inconsistency handling, like belief merging or those based on the selection of preferred consistent subbases where trivialization is avoided by weakening the belief base (while keeping classical entailment as the inference mechanism), the three-valued logic approach is helpful in the situation where a single source of inconsistent information must be treated. The basic inference relation from three-valued paraconsistent logic has the advantage to benefit from a simple semantics, where each connective is truth functional. Furthermore, the corresponding decision problem is only coNP-complete in the general case [8], while tractable fragments (for which classical entailment is intractable, like the CNF one) exist [6].

However, the basic inference relation from three-valued paraconsistent logic suffers from important drawbacks. One of them (shared by many paraconsistent inference relations) is the fact that inference is too cautious. In particular, it does not coincide with classical entailment in the case where the belief base is classically consistent. Another important drawback is that consequences of the underlying belief base with different epistemic status are not discriminated. Indeed, when some piece of information is derived from a belief base, we cannot state whether:

- It is *necessary*, which means that its negation cannot be the case.
- It is *plausible*, which means that its negation is not a consequence (intuitively, there is no reason to question the piece of information, even if it cannot be completely discarded that it could be false).
- It is *possible*, which means that its negation also is a consequence (there are some arguments in favor of the piece of information, and other arguments against it).

Such a myopia can be really problematic, especially because the set of all consequences of a paraconsistent relation based on three-valued logic can be classically inconsistent. Thus, in some situations, missing the consequences of the belief base that are only possible can be a good point. A still more cautious agent would even prefer to focus only on the necessary consequences of its beliefs since it forms a classically consistent set.

This paper contributes to fill the gap. Inference relations are pointed out, for which the separation between different epistemic status of consequences is handled. Following Priest [21] and others (e.g. [2]), cautiousness is avoided by focusing on preferential refinements of the basic inference relation. Typically, a principle of *inconsistency minimization* is at work: roughly, the worlds that are as close as possible to classical interpretations are preferred. While minimization is understood with respect to set inclusion in Priest's LP_m logic, other minimization schemes can be taken into account, especially those for which some variables are more important than others [18].

Then three mechanisms are suggested to refine both the basic and the preferential inference relations of three-valued logic. The first one is based on a principle of *argumentation*: α is derivable from the belief base Σ if α is a consequence of Σ but its negation is not. This way, possible consequences are avoided. The second one relies on a principle of *uncertainty minimization*: only those α that are evaluated to true (and not the middle value) in every model of Σ are kept. Thus, only the necessary consequences of Σ are kept. The third one can be viewed as a generalization of the second one, α is a consequence of Σ if for every interpretation ω, $\omega(\alpha)$ is "at least as true" and "at most as false" as $\omega(\Sigma)$ (which means that uncertainty about α decreased).

Our contribution is a systematic investigation of these inference relations along two fundamental dimensions:

- The *logical* dimension. We check each inference relation against the three-valued counterpart of system P (for preferential), a normative set of postulates that interesting valuable inference relations should satisfy [13, 15].
- The *computational* dimension. We identify the computational complexity of the decision problem corresponding to our inference relations.

We also compare all the inference relations that are considered with respect to cautiousness. This additional dimension is orthogonal to the two other ones since an inference relation can satisfy high standards for both the logical and the computational dimensions without being interesting if it is very cautious. The results from our analysis constitute a base line from which an inference relation offering the best compromise with respect to what is expected (cautiousness, myopia, logical properties, complexity) can be elected.

The rest of this paper is organized as follows. In Section 2, both the syntax and the semantics of the three-valued logic we are concerned with are presented; the corresponding (basic) inference relation is defined. In Section 3, more refined inference relations are described, and analyzed both from their logical side and from their computational side; they are also compared with respect to cautiousness. In Section 4 we study some special cases of preferential inference relations of interest. As a conclusion, we briefly discuss in Section 5 the results of the paper and we give some hints about future work.

2 A Three-valued Paraconsistent Logic

In the following, we consider a three-valued paraconsistent logic. Let's take some space to explain the meaning of the third truth value. We do not embrace here the paradigm of many-valued logics where additional truth values are "in between" true and false. We stay in the classical paradigm where the truth value of a formula is true or false. In fact, even with two truth values, there are more than two epistemic attitudes about a formula (see [10]), there are four distinct ones depending on whether or not we can prove truth or falsity of the formula from the base:

- {} We cannot prove the truth nor the falsity of the formula. This is typically the case when there is not enough information in the belief base to conclude.
- {0} We can prove the falsity of the formula (but not its truth), so the formula is "false" in the usual meaning.
- {1} We can prove the truth of the formula (but not its falsity), so the formula is "true" in the usual meaning.
- {0,1} We can prove both the truth and the falsity of the formula, i.e. the formula is "contradictory" in the usual meaning.

The problem is that in classical logic if a formula is contradictory, then it pollutes all the belief base (*"ex falso quodlibet sequitur"*). We want to avoid this contagion, and so we will use a logic in which a truth value "both" denotes that a formula can be proved at the same time "true" and "false" in the belief base. This will allow to highlight contradictory formulas, but still reasoning "reasonably" about the other formulas. Thus the third truth value "$\top$" we will use has to be understood as some encoding of the epistemic attitude {0,1}, and not as a truth value like "0" and "1".

2.1 Syntactical Aspects

Several presentations of the logic are possible, depending on the chosen set of connectives. We focus on one that is functionally complete [2], contrariwise to several presentations based on more restricted fragments, like those reported in [9], [20], [11] or [16].

Definition 1 (language). *$\mathcal{L}^*$ is the propositional language over a finite set $\mathcal{L}$ of propositional symbols, generated from the constant symbols true, $false$ and both, the unary connective $\neg$ and the binary connectives $\vee$, $\wedge$, and $\supset$.*

We will write propositional symbols a, b, ... and formulas will be denoted by lower case Greek letters α, β, ... A belief base, that will be denoted by an upper case Greek letter such as Σ, is a finite set of formulas (conjunctively interpreted).

Clearly enough, the fragment of $\mathcal{L}^*$ built up from $\neg$, $\vee$, $\wedge$ coincides with a standard language for classical propositional logic. It is referred to as $\{\neg, \wedge, \vee\}$ *fragment*. A proper subset of this fragment is composed by the *CNF formulas*, i.e. the (finite) conjunctions of clauses, where a clause is a (finite) disjunction of literals (the symbols from $\mathcal{L}$, possibly negated).

2.2 Semantical Aspects

Definition 2 (interpretation). *An interpretation ω over $\mathcal{L}$ is a total function from $\mathcal{L}$ to the set of truth values $\{0, 1, \top\}$. The set of all interpretations over $\mathcal{L}$ is noted $\mathcal{W}$.*

Whatever the interpretation ω from $\mathcal{W}$, we have $\omega(true) = 1$, $\omega(false) = 0$ and $\omega(both) = \top$. All the connectives are truth functional ones and the semantics of a formula α from $\mathcal{L}^*$ in an interpretation is defined in the obvious compositional way given the following truth tables (Table 1).

When working with more than two truth values, one has to set the set of designated values, i.e. the set of values that a formula can take to be considered as satisfied. Since we want to define a paraconsistent logic, we choose $\mathcal{D} = \{1, \top\}$: intuitively, a formula is satisfied if it is "at least true" (but it can also be false!).

We are now ready to define notions of models and of consequences:

Definition 3 (model, consequence).

- *ω is a model of Σ, denoted $\omega \models \Sigma$, iff for all $\alpha \in \Sigma$, $\omega(\alpha) \in \mathcal{D}$. $mod(\Sigma)$ denotes the set of models of Σ.*
- *α is a consequence of Σ, noted $\Sigma \models \alpha$, iff every model of Σ is a model of α.*
- *Σ is consistent iff it has at least one model ($mod(\Sigma) \neq \emptyset$).*

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

Table 1. Truth tables.

It is easy to see that many additional connectives can be defined as syntactic sugars in the three-valued logic we focus on. For instance:

- $\alpha \leftrightarrow \beta =_{def} (\alpha \supset \beta) \wedge (\beta \supset \alpha)$. $\leftrightarrow$ is an equivalence operator. $\alpha \leftrightarrow \beta$ is evaluated to a designated truth value iff both α and β are evaluated to a designated truth value, or none of them is (or, equivalently, α and β have the same set of models).
- $\Box\alpha =_{def} (\neg\alpha) \supset false$. $\Box\alpha$ is evaluated to 1 when α is evaluated to 1, otherwise $\Box\alpha$ is evaluated to 0. Thus, $\Box$ is a necessity operator.
- $\diamond\alpha =_{def} \neg\Box\neg\alpha$. $\diamond$ is the (dual) possibility operator.
- $\odot\alpha = (\Box\alpha) \vee (\Box\neg\alpha)$. $\odot\alpha$ is evaluated to 1 when α is evaluated classically (to 1 or 0), otherwise $\odot\alpha$ is evaluated to 0.

One can observe that the implication connective $\supset$ does not coincide with usual material implication $\Rightarrow$[1]. This does not change anything when truth values 0 and 1 are considered, only, so the $\supset$ connective can be considered as one of the possible generalizations of the implication connective of classical logic. It is the "right" generalization of classical implication, because $\supset$ is the internal implication connective [3] for the defined inference relation in the sense that a deduction (meta)theorem holds for it: $\Sigma \wedge \alpha \models \beta$ iff $\Sigma \models \alpha \supset \beta$.

An interesting feature of the inference relation $\models$ is that inconsistency cannot occur in the $\{\neg, \wedge, \vee\}$ fragment (this is not the case when the full language is considered since, for instance, $false$ and $\Box(a \wedge \neg a)$ are inconsistent formulas). Indeed, every formula from the $\{\neg, \wedge, \vee\}$ fragment has at least one model [6].

As evoked in the introduction, $\models$ has several other interesting properties. From the computational side, deciding it is "only" coNP-complete in the general case [8], and is even in P in the CNF fragment [16, 6]. From the logical side, it satisfies all the (three-valued counterparts of) postulates from system P and is even monotonic.

But the price to be paid is a very weak inference relation. Especially, it is well-known that disjunctive syllogism is not satisfied: $a \wedge (\neg a \vee b) \not\models b$. Subsequently, the set of consequences from a classically consistent belief base does

[1] $\alpha \Rightarrow \beta =_{def} (\neg\alpha) \vee \beta$.

not necessarily coincide with its classical deductive closure (we will see how to circumvent this).

Furthermore, $\models$ does not make any distinction between consequences with different epistemic status. For example, consider the following belief base:

$$\Sigma = \{(\Box a) \wedge b \wedge c \wedge \neg c\}.$$

From the belief base Σ, both a, b and c can be derived, whereas they have quite different status. a for example is *necessary*, since in each model of Σ, a is necessarily true. b is *plausible*, since we have some evidence about its truth but no evidence at all about its falsity. c is only *possible*, since we have contradictory pieces of evidence about it.

3 A Study of Refined Consequence Relations

3.1 Refining Basic Inference

In order to avoid both weak and myopic inference relations, four mechanisms can be exploited:

- taking advantage of some preferential information to focus on a subset of the set of models of the belief base.
- considering only argumentative consequences of the belief bases.
- selecting those consequences of the belief base that are necessarily true.
- selecting as consequences of the belief base formulas that are so to speak "more true" than the belief base.

Formally, the four principles above give rise to the following inference relations. Let $\leq$ be any binary relation on $\mathcal{W}$; $\min(mod(\Sigma), \leq)$ denotes the set $\{\omega \in mod(\Sigma) \mid \nexists \omega' \in mod(\Sigma)\ \omega' \leq \omega \text{ and } \omega \not\leq \omega'\}$.

Definition 4 (refined inference relations).

- *Let* $\leq$ *be a standard binary relation on* $\mathcal{W}$*:* [2]
  $\Sigma \models^{\leq} \alpha$ *iff* $\forall \omega \in \min(mod(\Sigma), \leq)$, $\omega \models \alpha$.
- $\Sigma \models_{arg} \alpha$ *iff* $\Sigma \models \alpha$ *and* $\Sigma \not\models \neg\alpha$.
- $\Sigma \models_1 \alpha$ *iff* $\forall \omega \in mod(\Sigma)$, $\omega(\alpha) = 1$.
- $\Sigma \models_t \alpha$ *iff* $\forall \omega \in \mathcal{W}$, $\omega(\Sigma) \leq_t \omega(\alpha)$, *where the truth ordering* $\leq_t$ *is given by the reflexive-transitive closure of* $0 <_t \top <_t 1$.

Combining the first mechanism with any of the three other ones results in some additional inference relations:

Definition 5 (refined inference relations). *Let* $\leq$ *be a standard binary relation on* $\mathcal{W}$*:*

- $\Sigma \models^{\leq}_{arg} \alpha$ *iff* $\Sigma \models^{\leq} \alpha$ *and* $\Sigma \not\models^{\leq} \neg\alpha$.
- $\Sigma \models^{\leq}_{1} \alpha$ *iff* $\forall \omega \in \min(mod(\Sigma), \leq)$, $\omega(\alpha) = 1$.
- $\Sigma \models^{\leq}_{t} \alpha$ *iff* $\forall \omega \in \min(mod(\Sigma), \leq)$, $\omega(\Sigma) \leq_t \omega(\alpha)$.

[2] We will call *standard* any binary relation on $\mathcal{W}$ whose definition is independent from the belief base Σ under consideration.

3.2 Cautiousness

Assuming that the belief base is consistent[3], we have derived the following results:

Theorem 1. *The inclusions between inference relations reported in Figure 1 hold.*

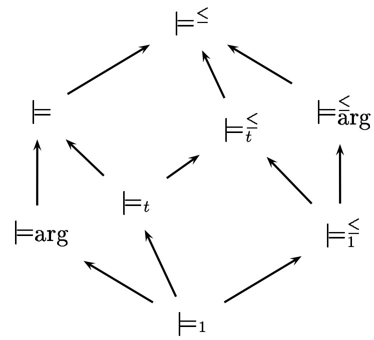

Fig. 1. Cautiousness (assuming Σ consistent)

Figure 1 gives a Hasse diagram of inclusion in the set of our inference relations. An arrow $X \to Y$ means that relation X is strictly more cautious than Y, i.e. $X \subsetneq Y$. Arrows that would stem from transitivity of inclusion are omitted.

As comments to these cautiousness results one can note that, unsurprisingly, the preferential inference relations usually contain their original counterparts, except that $\models_{\text{arg}}$ is not contained in $\models_{\text{arg}}^{\leq}$. The fact that all the inference relations are contained in $\models$ or $\models^{\leq}$ is not surprising since they aim at separating consequences with different epistemic status: $\models_{\text{arg}}$, $\models_1$, $\models_t$ are less cautious than $\models$. Similarly, the relations $\models_{\text{arg}}^{\leq}$, $\models_1^{\leq}$ and $\models_t^{\leq}$ are less cautious than $\models^{\leq}$. The strong principle of uncertainty minimization (focusing on truth value 1) is more demanding than both the weak one (based on the $\leq_t$ pre-order) and argumentation ($\models_1$ is included in $\models_t$ and in $\models_{\text{arg}}$) and this is still the case when some preference information is taken into account ($\models_1^{\leq}$ is included in $\models_t^{\leq}$ and $\models_{\text{arg}}^{\leq}$).

3.3 Logical Properties

Following seminal works in non-monotonic logic [12, 17, 13, 15], a set of normative properties that a non-monotonic inference relation should satisfy has been given in [13]. This set of properties is called system P (for Preferential).

[3] Without this assumption, many of the inference relations trivialize, either because they coincide with the total relation $\mathcal{L}^* \times \mathcal{L}^*$ or with the empty one (argumentative relations). The only change with respect to cautiousness is that $\models_1$ (resp. $\models_1^{\leq}$) is no longer included in $\models_{\text{arg}}$ (resp. $\models_{\text{arg}}^{\leq}$).

Definition 6 (system P). *An inference relation $\mid\!\sim$ is preferential if it satisfies the following properties (system* P*):*

(Ref)	$\alpha \mid\!\sim \alpha$	*Reflexivity*
(LLE)	*If* $\models \alpha \leftrightarrow \beta$ *and* $\alpha \mid\!\sim \gamma$*, then* $\beta \mid\!\sim \gamma$	*Left Logical Equivalence*
(RW)	*If* $\models \beta \supset \gamma$ *and* $\alpha \mid\!\sim \beta$*, 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* $\alpha \mid\!\sim \gamma$	*Cut*
(CM)	*If* $\alpha \mid\!\sim \beta$ *and* $\alpha \mid\!\sim \gamma$*, then* $\alpha \wedge \beta \mid\!\sim \gamma$	*Cautious Monotony*

Those properties have been stated in the framework of classical logic, but as we work here in a three-valued setting, we have to consider that $\models$ denotes the three-valued inference relation, as given in Definition 3. In the same vein, we refer to the "classical" three-valued implication connective $\supset$, and the equivalence connective $\leftrightarrow$ in the properties above. Following Arieli and Avron [2], we call the relations satisfying those properties three-valued preferential relations.

We will say that a relation is monotonic if it satisfies the following:

(Mon)	If $\alpha \mid\!\sim \gamma$, then $\alpha \wedge \beta \mid\!\sim \gamma$	Monotony

Theorem 2. *The logical properties of* P *satisfied by the inference relations considered in the paper are as given in Table 2.*

	Ref	LLE	RW	Or	Cut	CM	Mon
$\models$	✓	✓	✓	✓	✓	✓	✓
$\models^{\leq}$	✓	✓	✓	✓	✓	✓	
$\models_{arg}$		✓		✓	✓	✓	
$\models_{1}$		✓		✓	✓	✓	✓
$\models^{\leq}_{arg}$		✓		✓	✓	✓	
$\models^{\leq}_{1}$		✓		✓	✓	✓	
$\models_{t}$	✓			✓	✓	✓	✓
$\models^{\leq}_{t}$	✓			✓	✓	✓	

Table 2. Logical Properties of the Inference Relations

One can note that $\models^{\leq}$ inference relations, that are the less cautious ones, satisfy all properties of system P, and that all the preferential inference relations have the same logical properties as their original counterpart (except monotony of course). (Or),(Cut), (CM) are always satisfied. We can also observe that, unsurprisingly, (RW) is lost for all relations aiming at discriminating consequences obtained by $\models$ and $\models^{\leq}$.

3.4 Computational Complexity

We assume the reader familiar with some basic notions of complexity, especially the complexity classes coNP and Π_2^p of the polynomial hierarchy PH and the class BH_2 of the Boolean hierarchy (see [19] for a survey).

We have derived the following results:

Inference	Complexity of its decision problem
$\models$	coNP-complete
$\models_1$	coNP-complete
$\models_t$	coNP-complete
$\models_{\mathrm{arg}}$	BH_2-complete
$\models^{\leq}$	in Π_2^p
$\models_{\mathrm{arg}}^{\leq}$	in Δ_3^p
$\models_{1}^{\leq}$	in Π_2^p
$\models_{t}^{\leq}$	in Π_2^p

Table 3. Complexity Results

Theorem 3. *The complexity results reported in Table 3 hold (where it is assumed that $\leq$ can be tested in (deterministic) polynomial time).*

In the light of these results, the following observations can be done. First, all the inference relations considered in this paper are intractable, the complexity varying from the first level to the second level of the polynomial hierarchy. Second, focusing on the necessary consequences does not imply a complexity shift ($\models_1$ is just as complex as $\models$, and $\models_1^{\leq}$ is just as complex as $\models^{\leq}$). As it is the case in other frameworks [7], the corresponding argumentative versions of inference relations are mildly harder. Finally, as expected, preferring some models may definitely lead to more complex inference (under the standard assumptions of complexity theory). Especially, when $\leq$ is the preference pre-ordering proposed by Priest (preferring interpretations that are as classical as possible with respect to set inclusion), the decision problems associated to $\models^{\leq}$ and $\models_1^{\leq}$ are complete for Π_2^p, and this is still the case when the belief base Σ is from the CNF fragment and the query is an atom. Intuitively, when Priest's preference relation is considered, two independent sources of complexity must be dealt with: the first one lies in the number of preferred models and the second one in the difficulty to check whether a model is preferred. The decision problem associated to $\models_{\mathrm{arg}}^{\leq}$ is both Σ_2^p-hard and Π_2^p-hard, showing that it is not in $\Sigma_2^p \cup \Pi_2^p$, unless the polynomial hierarchy collapses.

4 Preference Based Inference Relations

In the previous section, for the inference relations $\models^{\leq}$, $\models_{\mathrm{arg}}^{\leq}$ and $\models_1^{\leq}$ we consider a binary relation $\leq$ with no special property. But one can expect this relation to express some kinds of preferences and thus to have some specific properties. A first intuitive need is to work with a pre-order for example, since preference relations are often transitive ones.

In this section we will investigate more deeply and compare four particular relations. Two of them prefer the most defined interpretations, i.e. the interpretations with a maximum of classical truth values. The first one takes the maximum for set inclusion (it was defined by Priest in [21]), the second one for cardinality. The last two relations give preference to interpretations that satisfy the more formulas in the belief base (w.r.t. set inclusion and cardinality). They are inspired by an inference relation defined by Besnard and Schaub [5]. Before defining formally those relations, we first need the following definitions:

Definition 7 (inconsistent set, satisfaction set).

- *The inconsistent set of an interpretation* ω *is* $\omega! = \{a \in \mathcal{L} \mid \ \omega(a) = \top\}$.
- *The satisfaction set of an interpretation* ω *given a belief base* Σ *is* $\mathcal{S}_\Sigma(\omega) = \{\alpha \in \Sigma \mid \ \omega(\alpha) = 1\}$.

Definition 8 (preferential inference relations). *Let* ω *and* ω' *be two interpretations from* $\mathcal{W}$. $\leq_P$, $\leq_{CP}$, $\leq^\Sigma_{BS}$ *and* $\leq^\Sigma_{CBS}$ *are defined by:*

- $\omega \leq_P \omega'$ *iff* $\omega! \subseteq \omega'!$
- $\omega \leq_{CP} \omega'$ *iff* $card(\omega!) \leq card(\omega'!)$
- $\omega \leq^\Sigma_{BS} \omega'$ *iff* $\mathcal{S}_\Sigma(\omega) \supseteq \mathcal{S}_\Sigma(\omega')$
- $\omega \leq^\Sigma_{CBS} \omega'$ *iff* $card(\mathcal{S}_\Sigma(\omega)) \geq card(\mathcal{S}_\Sigma(\omega'))$

The corresponding preferential inference relations $\models_{\mathrm{P}}$, $\models_{\mathrm{CP}}$, $\models^\Sigma_{\mathrm{BS}}$, $\models^\Sigma_{\mathrm{CBS}}$ *are those obtained using respective pre-orders* $\leq_P$, $\leq_{CP}$, $\leq^\Sigma_{BS}$, $\leq^\Sigma_{CBS}$ *for the definition of the preferential relation* $\models^{\leq}$.

Note that $\leq^\Sigma_{BS}$ and $\leq^\Sigma_{CBS}$ are non-standard binary relations since they depend on a belief base Σ. The logical properties have been studied for standard binary relations. Non-standard binary relations give operators that do not satisfy all properties ((LLE) is often missing).

Theorem 4. *The inclusions between inference relations reported in figure 2 hold.*

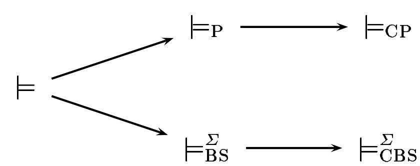

Fig. 2. Cautiousness

As the pre-orders $\leq_P$ and $\leq_{CP}$ are standard pre-orders, from Theorem 2 we get that the two inference relations $\models_{\mathrm{P}}$ and $\models_{\mathrm{CP}}$ are three-valued preferential relations (they satisfy all the logical properties).

Theorem 5. $\models_{\mathrm{P}}$ *and* $\models_{\mathrm{CP}}$ *satisfy (Ref), (LLE), (RW), (Or), (Cut) and (CM).* $\models^\Sigma_{\mathrm{BS}}$ *and* $\models^\Sigma_{\mathrm{CBS}}$ *satisfy (Ref), (RW), (Or), (Cut) and (CM).* [4]

One of the main interests of those four relations is that they coincide with classical inference when the belief base is classically consistent, and that they give meaningful inferences otherwise.

Finally, while the decision problems associated with $\models_{\mathrm{P}}$ and $\models^\Sigma_{\mathrm{BS}}$ are Π^p_2-complete, those associated to $\models_{\mathrm{CP}}$ and $\models^\Sigma_{\mathrm{CBS}}$ are in Δ^p_2, "only".

[4] Note that those two preference relations are based on sets of formulas, so one has to consider a generalization of the logical properties to sets of formulas (see e.g. [2]).

5 Conclusion

In this paper, we have investigated several three-valued inference relations for paraconsistent reasoning, both from the logical and the computational point of view. We have highlighted the fact that the basic three-valued inference relation can be refined with respect to the status of derived facts.

In the light of the results obtained, it appears that no relation is better than all the other ones with respect to both criteria. On the one hand, discriminating consequences does not imply a major complexity shift but leads to lose some valuable expected properties. On the other hand, avoiding too cautious relations can be achieved while keeping many interesting logical properties but typically leads to an increase in complexity. The choice of a good compromise depends mainly on what is expected in priority (avoiding cautiousness or myopia).

An interesting family of three-valued inference relations are those based on the selection of preferred models based on cardinality. Indeed, those relations have good logical properties, they also coincide with classical inference when the belief base is classically consistent (provided that a suitable preference relation has been chosen), and to give meaningful conclusions when the belief base is classically inconsistent. Furthermore, their decision problem remains at the first level of the polynomial hierarchy.

When working with more than two truth values a question that may arise is how many truth values are needed. In [2], Arieli and Avron showed that four-valued logics play a central role in billatice-based multi-valued logics, showing that more than four values are not necessary (the other logics can be characterized by a four-valued one). In such four-valued logics (e.g. [22, 1]), based on Belnap's seminal work [4], the fourth value $\perp$ means "not known", i.e. it denotes the epistemic attitude where we can prove neither the truth nor the falsity of the formula (see Section 2). Even if this truth value adds expressivity to the language and is useful to express ignorance (when one wants to be able to express the fact that an agent is agnostic about a formula), it does not help much as far as paraconsistency is concerned.

A future work concerns the design of three-valued inference relations as "blackbox" tools for paraconsistent reasoning in the classical framework. The idea is, from a classical two-valued belief base (possibly inconsistent), to derive facts with a three-valued logic (allowing to rule out the ex falso quodlibet), and finally to come back to the classical framework by translating three-valued models into two-valued ones. Several means can be envisioned for the last step, in particular one can use a generalization of the forgetting operator [14] to three-valued logics.

Acknowledgements

The authors want to thank two anonymous referees for their careful reading. The second author has been partly supported by the IUT de Lens, the Université d'Artois, the Région Nord/Pas-de-Calais under the TACT-TIC project, and by the European Community FEDER Program.

References

1. O. Arieli and A. Avron. Four-valued diagnoses for stratified knowledge-bases. In *Proceedings of CSL'96*, LNCS 1258, pages 1–17, 1997.
2. O. Arieli and A. Avron. The value of four values. *Artificial Intelligence*, 102:97–141, 1998.
3. A. Avron. Simple consequence relations. *Information and Computation*, 92:105–139, 1991.
4. N.D. Belnap. *A useful four-valued logic*, pages 7–37. Modern Uses of Multiple-Valued Logic. Reidel, 1977.
5. Ph. Besnard and T. Schaub. Circumscribing inconsistency. In *Proceedings of IJCAI'97*, pages 150–155, 1997.
6. M. Cadoli and M. Schaerf. On the complexity of entailment in propositional multivalued logics. *Annals of Mathematics and Artificial Intelligence*, 18:29–50, 1996.
7. C. Cayrol and M.-C. Lagasquie-Schiex. Non-monotonic syntax-based entailment: a classification of consequence relations. In *Proceedings of ECSQARU'95*, LNAI 946, pages 107–114, 1995.
8. S. Coste-Marquis and P. Marquis. Complexity results for paraconsistent inference relations. In *Proceedings of KR'02*, pages 61–72, 2002.
9. I.M.L. D'Ottaviano and N.C.A. da Costa. Sur un problème de Jaśkowski. In *Comptes Rendus de l'Académie des Sciences de Paris*, volume 270, pages 1349–1353. 1970.
10. D. Dubois and H. Prade. Possibility theory, probability theory and multiple-valued logics : A clarification. *Annals of Mathematics and Artificial Intelligence*, 32:35–66, 2001.
11. A.M. Frisch. Inference without chaining. In *Proceedings of IJCAI'87*, pages 515–519, 1987.
12. D. M. Gabbay. Theoretical foundations for nonmonotonic reasoning in experts systems. In K. Apt, editor, *Logic and Models of Concurrent Systems*. Springer Verlag, 1985.
13. S. Kraus, D. Lehmann, and M. Magidor. Nonmonotonic reasoning, preferential models and cumulative logics. *Artificial Intelligence*, 44:167–207, 1990.
14. J. Lang and P. Marquis. Resolving inconsistencies by variable forgetting. In *Proceedings of KR'02*, pages 239–250, 2002.
15. D. Lehmann and M. Magidor. What does a conditional knowledge base entail? *Artificial Intelligence*, 55:1–60, 1992.
16. H.J. Levesque. A knowledge-level account of abduction (preliminary version). In *Proceedings of IJCAI'89*, pages 1061–1067, 1989.
17. D. Makinson. *Handbook of Logic in Artificial Intelligence and Logic Programming*, volume III, chapter General Pattern in nonmonotonic reasoning, pages 35–110. Clarendon Press, Oxford, 1994.
18. P. Marquis and N. Porquet. Resource-bounded inference from inconsistent belief bases. In *Proceedings of IJCAI'01*, pages 103–108, 2001.
19. Ch. H. Papadimitriou. *Computational Complexity*. Addison-Wesley, 1994.
20. G. Priest. Reasoning about truth. *Artificial Intelligence*, 39:231–244, 1989.
21. G. Priest. Minimally inconsistent LP. *Studia Logica*, 50:321–331, 1991.
22. E. Sandewall. A functional approach to non-monotonic logic. In *Proceedings of IJCAI'85*, pages 100–106, 1985.