

How to infer from inconsistent beliefs without revising?

Salem Benferhat, Didier Dubois and Henri Prade

Institut de Recherche en Informatique de Toulouse (I.R.I.T.) – C.N.R.S.

Université Paul Sabatier – Bât. 1R3, 118 route de Narbonne

31062 Toulouse Cedex, France

Abstract

This paper investigates several methods for coping with inconsistency caused by multiple source information, by introducing suitable consequence relations capable of inferring non-trivial conclusions from an inconsistent stratified knowledge base. Some of these methods presuppose a revision step, namely a selection of one or several consistent subsets of formulas, and then classical inference is used for inferring from these subsets. Two alternative methods that do not require any revision step are studied: inference based on arguments, and a new approach called *safely supported inference*, where inconsistency is kept local. These two last methods look suitable when the inconsistency is due to the presence of several sources of information. The paper offers a comparative study of the various inference modes under inconsistency.

1 Introduction

Inconsistency can be encountered in different reasoning tasks, in particular:

- when reasoning with exception-tolerant generic knowledge, where the knowledge base includes default rules and instanciated facts and later a new information is received that contradicts a plausible conclusion derived from the previous knowledge base;
- in abductive reasoning, for instance in model-based diagnosis, when observations conflict with the normal functioning mode of the system and the hypothesis that the components of the system are working well; this leads to diagnose what component(s) fail(s);
- when several consistent knowledge bases pertaining to the same domain, but coming from n different experts, are available. For instance, each expert is a reliable specialist in some aspect of the concerned domain but less reliable on other aspects. A straightforward way of building a global base Σ is to concatenate the knowledge bases K_i provided by each expert. Even if K_i is consistent, it is rather unlikely that $K_1 \cup K_2 \cup \dots \cup K_n$ will be consistent also.

This paper is primarily oriented towards the treatment of inconsistency caused by the use of multiple sources of information. Knowledge bases considered in this paper are all stratified, namely each formula in the knowledge base is associated with its level of certainty corresponding to the layer to which it belongs. The use of priorities among formulas has been shown to be very important to appropriately revise inconsistent knowledge bases (Fagin et al., 1983). In particular, Gärdenfors (1988) has proved that any revision process that satisfies natural requirements is implicitly based on priority ordering. In the context of

merging several knowledge bases, the introduction of priorities between pieces of information in Σ can be explained by the two following scenarios:

- Each consistent knowledge base K_i , issued from a source of information, is "flat" (i.e., without any priority between their elements). But we have a total pre-ordering between the sources of information according to their reliability. In this case merging different sources of information lead to a prioritized knowledge base Σ , where the certainty level of each formula reflects the reliability of the source. A particular case is when each piece of information in Σ is supported by a different source.
- All sources of information are equally reliable (and thus have the same level of reliability), but inside each consistent knowledge base K_i there exists a preference relation between pieces of information given by an expert, who rank-orders them according to their level of certainty. Here again, the combination of the different sources of information gives an uncertain knowledge base, provided that the scales of uncertainty used in each knowledge base K_i are commensurate.

This paper investigates two classes of approaches to deal with inconsistency in knowledge bases: coherence theories and foundation theories. Somewhat departing from what is usually considered, we view this dichotomy in the following way. These two classes correspond to two attitudes in front of inconsistent knowledge. One (the coherence theories) insists on revising the knowledge base and restoring consistency. The other (the foundation theories) accepts inconsistency and copes with it. Coherence theories propose to give up some formulas of knowledge base in order to get one or several consistent subbases of Σ and to apply classical entailment on these consistent subbases to deduce plausible conclusions of the knowledge base. Foundation theories proceed differently since they retain all available information but each plausible conclusion inferred from the knowledge base is justified by some strong reason for believing in it. Such reasons are based on the idea of argument that goes back to Toulmin (1956), and is related to previous proposals by Poole (1985), Pollock (1987), and Simari and Loui (1992) which were however suggested in the framework of defeasible reasoning for handling exceptions. There also exist approaches to reasoning with inconsistent knowledge bases which, for instance, identify a consistent part and an inconsistent part in the base as in (Lin, 1994), or which combine the consequences obtained from each source of information as in (Dubois et al., 1992b) rather than combining the knowledge bases attached to each source before the inference process takes place.

Our dichotomy coherence versus foundation is somewhat

different from the one used in the literature (Harman, 1986), (Gärdenfors, 1990), (Rao & Foo, 1989), (Doyle, 1992), (DeVal, 1994). In this paper, we do not assume any particular structure on the beliefs in the knowledge base (contrary for example to RMS defined in (Doyle, 1992)). Nor do we assume any (in)dependence relations between beliefs. Moreover, beliefs in a knowledge base are all "self-justifying", namely all pieces of information are put in the knowledge base as they are and as they come from their sources of information, and we do not add to the knowledge base any derived beliefs.

This paper extends previous results of Benferhat et al. (1993b) and discusses two foundation approaches in greater details. The following points are developed: We survey some coherence theories to the inconsistency handling. We recall four consequence relations which consist in replacing the inconsistent knowledge base by one or several of its consistent subbases. We then recall the so-called *argumentation inference*, proposed in (Benferhat et al., 1993a). We finally propose a new foundation theory. It is based on a so-called *safely supported consequence relation*, which deals with inconsistency in a "local" way. The approach is local in the sense that each formula in the base is also associated with a "level of paraconsistency" which reflects the maximal strength of arguments in favour of the opposite formula. Section 5 is entirely devoted to a comparison of the different consequence relations considered in the paper. Three criteria are used to do this: i) cautiousness; ii) properties; iii) syntax-sensitivity.

2 Background

In this paper, we only consider a finite propositional language. The symbol \vdash represents the classical consequence relation, Greek letters $\alpha, \beta, \delta, \dots$, represent formulas. Let Σ be a multiset of propositional formulas, possibly inconsistent but not deductively closed. When the knowledge base Σ is not deductively closed, we call it a "belief base", following Nebel (1991). In presence of inconsistency, the approaches developed in this paper must be syntactic in nature, since they explicitly use formulas that appear in the knowledge base originally, while two inconsistent knowledge bases over the same language are semantically equivalent (in a trivial way). This paper deals only with stratified knowledge bases which can be viewed as layered knowledge bases of the form $\Sigma = S_1 \cup \dots \cup S_n$, such that formulas in S_i have the same level of priority or certainty and are more reliable than the ones in S_j where $j > i$. This stratification is modelled in possibilistic logic (Dubois et al., 1994) by attaching a weight $a \in [0, 1]$ to each formula with the convention that $(\phi, a_i) \in S_i, \forall i$ and $a_1 > a_2 > \dots > a_n > 0$. From now on, a stratification is used to represent prioritized knowledge bases, the lower is the rank i of a stratum, the higher is the level of certainty a_i of the formulas included in it. The rank i can be viewed as a *level of defeasibility* of the formulas in S_i . The greater i , the more defeasible the formulas in S_i .

Throughout this paper, we denote subbases by capital letters A, B, C, \dots and they are also represented in a stratified way, namely $A = A_1 \cup \dots \cup A_n$ where $\forall j = 1, n, A_j \subseteq S_j$ and

possibly $A_j = \emptyset$. From now on, we denote by $MC(\Sigma)$ the set of all maximally consistent subbases, and by $Free(\Sigma)$ the set of all the formulae, called *free formulas*, which are not involved in any inconsistency of the belief base Σ , namely: $Free(\Sigma) = \{\phi / \nexists A \subseteq \Sigma, \text{ s.t. } \phi \in A \text{ and } A \text{ is minimal inconsistent}\}$. Clearly $Free(\Sigma)$ may be empty. We finish this section by defining the notion of free-consequence:

Def. 1: A formula ϕ is said to be a *free consequence* of Σ , denoted by $\Sigma \vdash_{Free} \phi$, iff ϕ is logically entailed from $Free(\Sigma)$, namely: $\Sigma \vdash_{Free} \phi$ iff $Free(\Sigma) \vdash \phi$.

The free inference reduces to the classical inference when Σ is consistent, and is very conservative otherwise, since it corresponds to a maximal revision of Σ , deleting all formulas involved in a conflict.

3 Coherence-Based Approaches to Inconsistency

Coherence approaches can be described in two steps: i) give up some formulas of belief base in order to restore its consistency, the result of this operation is one or several consistent subbases of Σ , and ii) apply classical entailment on these consistent subbases to deduce plausible conclusions from the belief base. We investigate two classes of coherence theories: coherence theories based on the choice of one consistent subbase (not necessarily maximal), and coherence theories based on the selection of several maximal consistent subbases. From a pragmatic point of view, the first class is very interesting (with computational complexity close to one of the classical logic), while selecting several maximal consistent subbases is computationally very difficult (see (Nebel, 1994), (Cayrol & Lagasque-Schiex, 1994) for a discussion of complexity results of inconsistency handling approaches). However, from the minimal change point of view, the second class seems more satisfactory since it keeps as many formulas as possible while the first class selects one consistent subbase which is often not maximal.

3.1 Approaches Based on the Selection of One Consistent Subbase

We start with the possibilistic approach. See (Dubois et al., 1994) for a complete exposition of possibilistic logic. The possibilistic treatment of inconsistency is based on the selection of only one consistent subbase of Σ (in general not maximal), denoted by $\pi(\Sigma)$, induced by the levels of priority and defined in this way: $\pi(\Sigma) = S_1 \cup \dots \cup S_i$, where $i = \max\{j \text{ s.t. } S_1 \cup \dots \cup S_j \text{ is consistent}\}$. If S_1 is inconsistent then $\pi(\Sigma) = \emptyset$. If Σ is consistent, $\pi(\Sigma) = \Sigma$. The basic intuition in the possibilistic approach is to only take into account the first i consistent strata which are the most important ones in terms of certainty. The remaining subbase $\Sigma - \pi(\Sigma)$ is simply inhibited. It is clear that the computational complexity of the possibilistic approach is very attractive since it needs at most $\log(n)$ satisfiability (SAT) tests. However, this approach is very drastic ("liberal") and the amount of the formulas given up may be important. Nebel(1994), elaborating on a suggestion made in (Dubois & Prade, 1991), has proposed a less liberal way to select one consistent subbase. The idea is to consider each

stratum as composed of one element obtained by the conjunction of the formulas inside this stratum. Such belief bases are called "*unambiguous*" or "*linear ordered*". When inconsistency occurs, we give up the whole stratum concerned by the inconsistency, but we continue to add strata with lower certainty levels if consistency is preserved. More formally, the selected subbase is denoted by $lo(\Sigma)$ and is computed in the following way: $lo(\Sigma) := \emptyset$

for $i := 1$ to n do $lo(\Sigma) := lo(\Sigma) \cup S_i$ if consistent
 $:= lo(\Sigma)$ otherwise.

The inference relations for these two approaches are:

Def. 2: A formula ϕ is said to be a *possibilistic consequence* (resp. a *lo-consequence*) of Σ , denoted by $\Sigma \vdash \pi\phi$ (resp. $\Sigma \vdash_{lo}\phi$), iff ϕ is logically entailed by $\pi(\Sigma)$ (resp. $lo(\Sigma)$), namely, iff $\pi(\Sigma) \vdash \phi$ (resp. $lo(\Sigma) \vdash \phi$).

3.2 Approaches Based on a Selection of Maximal Consistent Subbases

Probably, one of the best known approaches to reasoning with inconsistency is the one proposed by Rescher and Manor (1970) and based on the *universal* consequence relation: first compute the set of maximal consistent subsets of the belief base; then a formula is accepted as a consequence when it can be classically inferred from all the maximal consistent subsets of propositions. However universal consequence relation does not take advantage of the layered structure of the belief base, and therefore the cardinality of $MC(\Sigma)$ which increases exponentially with the number of conflicts in the base may be very high. One may think of selecting a non-empty subset of $MC(\Sigma)$, called preferred subbases of Σ , which represents maximal consistent subbases that keep as many formulas of Σ as possible. There exist two criteria to define such preferred subbases of Σ : set-inclusion or cardinality (Benferhat et al., 1993b).

Def. 3: A consistent subbase $A = A_1 \cup \dots \cup A_n$ is an *inclusion-preferred subbase* of Σ (Incl for set-inclusion) iff it does not exist a subbase $B = B_1 \cup \dots \cup B_n$ of Σ such that:
 $\exists i \leq n$, where $A_i \subset B_i$ and for $j < i$ we have $B_j = A_j$.

Def 4: A consistent subbase $A = A_1 \cup \dots \cup A_n$ is a *cardinality preferred subbase* (or a *lex-preferred subbase*) of Σ iff it does not exist a subbase $B = B_1 \cup \dots \cup B_n$ such that:
 $\exists i \leq n$, where $|B_i| > |A_i|$ and for $j < i$ we have $|B_j| = |A_j|$ where $|A|$ is the cardinality of A .

From now on, we denote by $Incl(\Sigma)$ and $Lex(\Sigma)$ the set of inclusion-preferred subbases and Lex-preferred subbases of Σ . Inclusion-preferred subbases have been proposed by (Brewka, 1989) under the name "*preferred sub-theories*" and have also been independently introduced by (Dubois et al., 1992a) in the setting of possibilistic logic under the name of *strongly maximal consistent subbases*. Baral et al. (1992) have also used a similar approach to combine belief bases. The definition of $Lex(\Sigma)$ has been proposed in another form in (Dubois et al., 1992a) and also independently in (Lehmann, 1993). The idea of selecting a subset of the set of maximally consistent subbases of Σ using a cardinality criterion was used independently in diagnostic problems (De Kleer, 1990;

Lang, 1994). It corresponds to the property of parsimony advocated in (Reggia et al., 1985). Once $Incl(\Sigma)$ and $Lex(\Sigma)$ are computed, we define the nonmonotonic consequence relation in the following way:

Def. 5: A formula ϕ is said to be a *Incl-consequence* (resp. *Lex-consequence*) of Σ , denoted by $\Sigma \vdash_{Incl}\phi$ (resp. $\Sigma \vdash_{Lex}\phi$), if and only if it is entailed from each element of $Incl(\Sigma)$ (resp. $Lex(\Sigma)$), namely, iff $\forall A \in Incl(\Sigma)$, (resp. $\forall A \in Lex(\Sigma)$) $A \vdash \phi$.

4 Foundation Theories of Inconsistency Handling

This section presents two foundation theories to deal with inconsistency. The first approach, proposed in (Benferhat et al., 1993a), is called argumentation inference and recalled in a concise manner below. The second approach is a new one which treats inconsistency in a "local way". This latter approach is presented and discussed in details in this section. Contrary to the coherence theories, foundation theories do not throw pieces of information away from the inconsistent belief base in order to maintain the whole consistency of the belief base. In the case of multiple sources problems, restoring consistency looks much more debatable, since the goal of retaining all available information is then quite legitimate. However, in foundation approaches each plausible conclusion is supported by an argument which can be seen as a reason (formed from explicit information of the belief base) to believe in it.

4.1 Argumentation Approach

The two foundation approaches are based on the idea of argument:

Def. 6: A consistent subbase A of Σ is said to be an *argument* to a rank i for a formula ϕ if it satisfies the following conditions: (i) $A \vdash \phi$, (ii) $\forall \psi \in A$, $A - \{\psi\} \not\vdash \phi$, and (iii) $i = \max \{j / (\phi, j) \in A\}$

Def. 7: An argument A of rank i which supports ϕ is a *best* argument iff each argument which supports ϕ is of rank $j \geq i$.

An argument A for ϕ is a minimal consistent subbase of Σ which entails logically ϕ . Its rank i is all the smaller as A supports the conclusion ϕ more strongly, first rank arguments being the best. Note that this notion of argument is an extension of the one proposed by Simari and Loui (1992). Elvang-Goransson et al. (1994) have also proposed a formal framework for argumentation. Cayrol (1995) discusses links between nonmonotonic consequence relations making use of the idea of arguments and nonmonotonic coherence-based entailment. The first of the two foundation theories considered here, called the *argumentation inference*, suggests that a conclusion can be inferred from an inconsistent belief base if the latter contains an argument of rank i that supports this conclusion, but there is no argument of rank smaller than or equal to i that supports its negation. More formally:

Def. 8: A formula ϕ is said to be an *argued consequence* of Σ , denoted by $\Sigma \vdash_{\mathcal{A}} \phi$, if and only if: (i) there exists an

argument of rank i for ϕ in Σ , and (ii) arguments for $\neg\phi$ (if any) are of rank $j > i$.

See (Benferhat et al., 1993a) for the properties of this consequence relation.

4.2 Safely Supported Inference

In the definition of argumentation inference, there is no constraint on beliefs used to build arguments in favor of plausible conclusions of the belief base. Namely, one argument may for instance contain pieces of information which are directly involved in the inconsistency of the belief base. Levels of priority or of certainty attached to formulas have only been used to distinguish between strong and less strong arguments in favour of a proposition or of its contrary. However it is possible to go one step further in the use of the certainty or priority levels by i) attaching to each proposition ϕ in the belief base the rank i of the best argument attached to ϕ , ii) the rank j attached to the best argument in favour of $\neg\phi$ if any, and by iii) inferring from weighted premises such as $(\phi \ i \ j)$ by propagating the ranks i and j . It will enable us to distinguish between consequences obtained only from "free" propositions in the belief base Σ (i.e., propositions for which there is no argument in Σ in favour of their negation), and consequences obtained using also propositions which are not free (for which there also exists an argument in favour of their negation). For $\phi \in \Sigma$, we denote by $\text{Def}(\phi)$ the rank of best arguments for ϕ in Σ (including ϕ itself). $\text{Def}(\phi)$ reflects the defeasibility level of ϕ . On the contrary $\text{Def}(\neg\phi)$ expresses our confidence in the belief $\phi \in \Sigma$, since the higher $\text{Def}(\phi)$ the less reasons for doubting ϕ .

Def. 9: Let ϕ be a formula of Σ . Then ϕ is said to be *paraconsistent* iff there also exists an argument for $\neg\phi$ in Σ . We define the consistency rank of ϕ , denoted by $\text{Cons}(\phi)$, as the maximum of the ranks corresponding to the best arguments in favour of ϕ and in favour of $\neg\phi$. If ϕ is free then by convention $\text{Cons}(\phi) = \infty$.

Note that the higher the consistency rank, the less paraconsistent is ϕ . Moreover $\text{Cons}(\phi) = \text{Cons}(\neg\phi)$. We now introduce the notion of defeated formula:

Def. 10: Let ϕ be a formula of Σ . ϕ is said to be *defeated* iff $\text{Def}(\neg\phi) \leq \text{Def}(\phi) = \text{Cons}(\phi)$.

Classically and roughly speaking, knowledge about ϕ is paraconsistent if there exist reasons to state both ϕ and $\neg\phi$. It corresponds to the situation where we have conflicting information about ϕ . It is why we speak here of paraconsistent information, although the approach presented in the following departs from usual paraconsistent logics (following Da Costa (1963)). More formally, let $\Sigma_j = S_1 \cup \dots \cup S_j$ be the subbase of Σ composed of the first i strata, and $\text{Free}(\Sigma_j)$ denotes its free part. It is clear that $\text{Free}(\Sigma_j)$ is different from $(\text{Free}(\Sigma))_j$, and more precisely we have the following relation: $(\text{Free}(\Sigma))_i \subseteq \text{Free}(\Sigma_i)$.

Def. 11: A formula ϕ is said to be a *safely supported consequence* of Σ , denoted by $\Sigma \vdash_{\text{SS}} \phi$, iff there exists a rank i such that $\Sigma_i \vdash_{\text{Free}} \phi$.

Denote again by $\text{Def}(\phi)$ be the smallest rank i such that

$\Sigma_i \vdash_{\text{Free}} \phi$ for $\phi \in \Sigma$. Notice that if for a given rank $k > \text{Def}(\phi)$ we have $\Sigma_k \not\vdash_{\text{Free}} \phi$, then there is no longer a proof of ϕ in Σ_k made of free formulas only and that at least one of the formulas, say ψ , used in the free proof of ϕ from $\Sigma_{\text{Def}(\phi)}$ is paraconsistent. However, this does not mean that there is an argument for $\neg\phi$ in Σ_k , although there is an argument for ϕ in Σ_k obviously. Indeed, consider the following counter-example: $\Sigma = S_1 = \{\psi, \neg\psi, \neg\psi \vee \phi\}$. It is clear that there is an argument for ϕ in the belief base Σ , and ϕ is not a free consequence of Σ , but there is no argument which supports $\neg\phi$.

Proposition 1: Let ψ be a safely supported consequence of Σ , then there exists an argument, say A , for ψ in Σ such that none of the formulas of A is defeated.

See proof in (Benferhat et al., 1995). The previous proposition means that the safely supported inference is based on undefeated arguments, and therefore the conclusions produced by this consequence relation are safe. The converse of Proposition 1 is not true. Indeed, consider the following counter-example $\Sigma = \{S_1 = \{\psi\}, S_2 = \{\neg\psi\}, S_3 = \{\neg\psi \vee \phi\}\}$, where the formula ψ is attacked, $\neg\psi$ is defeated and $\neg\psi \vee \phi$ is free. It is clear that ϕ is not a safely supported conclusion even if in the belief base Σ we have an argument for ϕ composed of the two undefeated formulas $\neg\psi \vee \phi$ and ψ .

Let us now evaluate our confidence in ϕ , namely to what extent the conclusion ϕ is safe. This safety depends on our confidence in the formulas of the belief base which are involved in inferring the conclusion ϕ . The safety rank of an argument A is denoted by $\text{Safe}(A)$ and computed in the following way: $\text{Safe}(A) = \text{Min}\{\text{Def}(\neg\psi) / \psi \in A\}$. It is clear that the best argument for ϕ in $\Sigma_{\text{Def}(\phi)}$ is the safest one, i.e.,

$$\text{confidence}(\phi) = \max\{\text{Safe}(A) / A \subseteq \Sigma_{\text{Def}(\phi)}\}.$$

The following proposition shows that if a formula ϕ is a safely supported consequence of Σ if and only if its defeasibility is strictly lower than the safety of its arguments:

Proposition 2: Let ϕ be a safely supported consequence of Σ . Then there exists an argument, say A , of rank i , which supports ϕ and where $i < \text{Safe}(A)$. The converse is also true.

See proof in (Benferhat et al., 1995). We now give a procedure to compute safely supported conclusions of a belief base. Algorithms for computing conclusions obtained using coherence theories are largely described in the literature (Baral et al., 1992; Benferhat et al., 1993b), and a theorem prover has been developed in (Cholvy, 1993). A procedure for computing argued consequences can be found in (Benferhat et al., 1993a). We use an ATMS for computing safely supported inferences. This tool can compute minimal inconsistent sets, and therefore Free formulas. Indeed, links between minimal inconsistent subbases and nogoods can be established in the following way: let Σ be a belief base, and let Σ' be a new belief base obtained from Σ by replacing each formula ϕ_i in Σ by $H_i \vee \phi_i$, where H_i is a hypothesis (all H_i are different). Then the subbase $A = \{\phi_i / i = 1, m\}$ is a minimal inconsistent subbase of Σ iff $H_A = \{H_i / H_i \vee \phi_i \in \Sigma', \phi_i \in A\}$ is a nogood. We denote by $\text{Incons}(K)$ the set of all formulas which belong to at least one minimal inconsistent subbase of K .

```

1 Function SS_consequence (Input:  $\Sigma, \phi$ ): Boolean;
2   Let i:=1; Answer:=false; K:= $\emptyset$ ;
3   While i $\leq$ n and Answer=false do
4     Begin
5       K:=K $\cup$ Si; i:=i+1;
6       Compute Incons(K) using ATMS,
7       If K $\setminus$ Incons(K)  $\models$   $\phi$  then Answer:=true,
8     End {Begin};
9   Return(Answer)
10 End {Procedure}

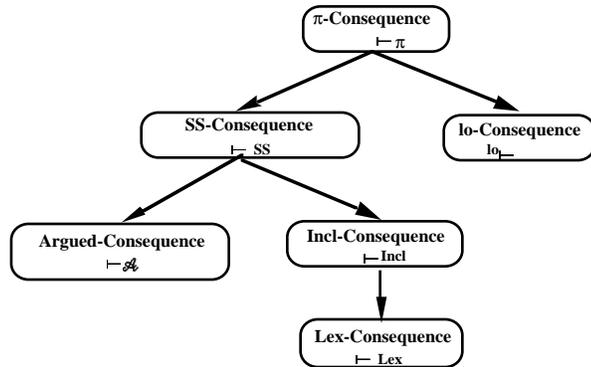
```

Notice that the complexity of the algorithm depends on the step 6. However, this step can be done once and for all.

5 Comparative Study

5.1 Cautiousness

A consequence relation \vdash_1 is said to be at least as cautious than \vdash_2 if and only if every conclusion of Σ obtained using \vdash_1 is also a conclusion using \vdash_2 . The following hierarchy summarizes the cautiousness relation between the different consequence relations studied here, the edges mean the inclusion-set relation between the set of results generated by each consequence relation. The top of the diagram thus corresponds to the most conservative inferences.



Proofs can be found in (Benferhat et al., 1994). Notice that all the consequence relations described above collapse with the possibilistic entailment \vdash_π when Σ is consistent. Moreover, when the base is flat then the safely supported inference is equivalent to the free consequence and Incl-consequence relation becomes more cautious than the argumentation consequence. Besides, when a base Σ contains exactly one formula per stratum then the consequence relations $\vdash_{Lex}, \vdash_{Incl}, \vdash_{SS}$ generate the same set of conclusions. Lastly, observe that if ϕ appears in Σ , and is an argued consequence, ϕ is also a safely supported consequence. But, they differ for other conclusions since the safely supported inference propagates the effects of local inconsistency. From the above figure, one may ask two questions: i) what is the minimal set of conclusions that we are ready to accept, namely that any reasonable consequence relation must contain? and ii) to what extent a consequence relation should be adventurous?

The natural and reasonable answer to the first question is to consider the set of possibilistic conclusions, since the possibilistic entailment takes into account the i most important and consistent strata. However, the possibilistic

way of dealing with inconsistency is not entirely satisfactory, since it suffers from an important drawback named "drowning problem" in (Benferhat et al., 1993), as we can see in the following example: Let Σ be the following stratified belief base: $\Sigma = \{S_1 = \{\neg\alpha \vee \neg\beta\}, S_2 = \{\alpha\}, S_3 = \{\beta\}, S_4 = \{\delta\}\}$. This belief base is inconsistent, and only the subset $A = \{\{\neg\alpha \vee \neg\beta\}, \{\alpha\}\}$ is kept, and therefore δ cannot be deduced despite the fact that δ is outside the conflict. In the example, the lo-consequence solves the problem. Another possibility is to select in Σ the intersection of all inclusion-preferred consistent subbases of Σ , which turns out to be equal to $\Sigma^* = \text{Free}(\Sigma_1) \cup \dots \cup \text{Free}(\Sigma_n)$. Inference from Σ^* gives a more adventurous coherence approach than the possibilistic inference, but it is in general not comparable with the lo-consequence (see Benferhat et al., 1994).

Now let us try to answer the second question, namely to what extent the consequence relations should be adventurous. Safely supported consequences seem cautious and safe since results produced by this consequence relation are based on undefeated arguments. lo-consequence relation may be seen as an adventurous approach since some produced conclusions are debatable when we consider belief bases whose layers can contain more than one formula. Indeed, consider the following example: $\Sigma = \{S_1 = \{\phi\}, S_2 = \{\neg\phi, \psi, \delta\}, S_3 = \{\neg\psi\}\}$. Here, we would like to deduce ψ and δ while lo-consequence relation will produce $\neg\psi$. The remaining consequence relations seem to be also adventurous. Indeed, take the following example $\Sigma = \{S_1 = \{\phi\}, S_2 = \{\neg\phi\}, S_3 = \{\neg\phi \vee \psi\}\}$, where ψ is a plausible consequence of Σ using $\vdash_{Incl}, \vdash_{Lex}, \vdash_{A}$ while ψ is not a safely supported consequence of Σ . Indeed ψ is supported only by $\{\phi, \neg\phi \vee \psi\}$, an argument of rank 3, and $\text{Cons}(\psi) < 3$.

5.2 Properties of Consequence Relations

This section positions the consequence relations presented in this paper inside the general nonmonotonic framework defined by the KLM postulates (System P and Rational inference R) proposed by Lehmann and Magidor (1992). The following array summarizes the properties of the inference relations:

	System P	R
π -consequence	Yes	Yes
lo-consequence	Yes	Yes
Incl-consequence	Yes	No
Lex-consequence	Yes	Yes
Argued-consequence	No	No
SS-consequence	No	No

Proofs can be found in (Benferhat et al., 1992) for the properties of π -consequence relation, in (Nebel, 1994) for lo-consequence relation, and in (Benferhat et al., 1993) for Incl-consequence and Lex-consequence relations.

The argumentation consequence and the safely supported inference do not belong to system P. Indeed, the following counter-examples show that even if ϕ and ψ are safely supported (resp. argued consequences) of a belief base Σ , their conjunction is not necessarily a safely supported consequence (resp. an argued consequence) of Σ (i.e., \vdash_{SS} and \vdash_{A} do not satisfy the "AND" property).

a. Let $\Sigma = \{S_1 = \{\alpha\}, S_2 = \{\neg\alpha \vee \phi\}, S_3 = \{\neg\rho \vee \psi\}, S_4 = \{\rho, \neg\alpha\}\}$. It is clear that ϕ and ψ are both safely supported consequences of Σ (since $\Sigma_2 \vdash_{\text{Free}} \phi$ and $\Sigma_4 \vdash_{\text{Free}} \psi$), while $\phi \wedge \psi$ is not a safely supported consequence of Σ , since there is no $i > 0$ such that $\Sigma_i \vdash_{\text{Free}} \phi \wedge \psi$.

b. Let $\Sigma = \{S_1 = \{\neg\alpha \vee \beta\}, S_2 = \{\alpha \vee \delta, \alpha, \neg\alpha\}\}$. It is clear that β and δ are both argued consequences of Σ , while there is no argument which supports $\beta \wedge \delta$.

The failure of the property of AND must not be seen as a major drawback of $\vdash_{\mathcal{A}}$ and $\vdash_{\mathcal{SS}}$, and should not be a surprise when dealing with multi-source inconsistent information. In some cases the AND property is not welcome, i.e., one should not perform the conjunctions of propositions that are supported by antagonistic views (as in the previous example). The argumentation inference captures the cases when we believe in two mutually consistent properties of some object for conflicting reasons. This situation also happens in numerical settings such as evidence theory (Shafer, 1976) since we may have $\text{Belief}(\phi) > 0$, $\text{Belief}(\psi) > 0$ and $\text{Belief}(\phi \wedge \psi) = 0$ with Shafer belief functions. Besides, the set of argued consequences of Σ can be inconsistent (Benferhat et al., 1993a), while the set of safely supported consequences is always consistent (Benferhat et al., 1995).

Now consider two inconsistency-tolerant consequence relations, the first one is rational while the second is not. Then, is it a sufficient condition to prefer the first one? In our opinion the response is no since properties of rational inference do not always guarantee minimal change between the inconsistent belief base Σ and the selected consistent subsets involved in the adopted consequence relation. This is clear, considering the possibilistic consequence relation which throws away many formulas, not even using a maximal-consistent subbase. It is very cautious but it is rational. The Incl-consequence relation exploits many maximal consistent subbases, is more adventurous, but is not rational. The argumentation inference is even more respectful of Σ , is also adventurous but is not even in system P. In fact, properties advocated by Lehmann and Magidor (1992) describe only the expected behaviour of the consequence relation for handling exceptions, and do not take into consideration a belief base whose inconsistency is due to the presence of several sources of information.

5.3 Syntax-Sensitivity

Baral et al. (1992) noticed that Incl-preferred consequence relation depends upon the syntax of the belief base. It means that even if $\{\phi \wedge \psi\}$ is logically equivalent to $\{\phi, \psi\}$, this equivalence does not matter when treating inconsistency. Following Nebel (1991), a syntactic consequence relation is a consequence relation which refers explicitly to the syntactic representation of the belief base. It is clear that all the consequence relations described in this paper are syntactic since they explicitly use formulas that appear in the belief base. However, some consequence relations seem to be more dependent on the syntax of the belief base than the others. Besides, from a semantic point of view the Lex-consequence relation may appear as an arbitrary selection from a set of maximal consistent subbases of Σ . Suppose indeed that

$A \in \text{Lex}(\Sigma)$, $B \notin \text{Lex}(\Sigma)$ and B is a maximally consistent subbase and one may define C logically equivalent to B but C is lexicographically preferred to A. To have it, it is enough to duplicate some formulas in B a sufficient number of times (remember that our view of Σ is syntactic in the sense that, for instance, $\Sigma = \{\phi\}$ is not seen as equivalent to $\Sigma' = \{\phi, \phi\}$). Moreover, even if some consequence relations explicitly refer to pieces of information of the belief base they can be syntax-independent, namely if we transform each formula in the belief base into its CNF form we will get the same results. An example of such consequence relation is the π -consequence relation, and it is very artificial to regard the possibilistic entailment as being a syntactic approach.

In this section we provide a formal discussion of the syntax-sensitivity of the consequence relations described above by proposing the following properties:

Redundancy insensitivity (RI): An inference relation \vdash is said to be a RI relation iff: $\forall \phi \in S_i, \Sigma \vdash \psi$ iff $\Sigma \cup \{(\phi \ i)\} \vdash \psi$. (We denote by $\Sigma \cup \{(\phi \ i)\}$ the belief base obtained by adding the formula ϕ to the layer S_i in Σ .)

Local consequence insensitivity (LCI): An inference relation \vdash is said to be a LCI relation iff $\forall \phi$ such that there exists in Σ an argument of rank i for ϕ : $\Sigma \vdash \psi$ iff $\Sigma \cup \{(\phi \ i)\} \vdash \psi$.

Clausal form insensitivity (CFI): Let Σ' be a new belief base obtained by replacing each formula in Σ by its clausal form. Then, an inference relation \vdash is said to be a CFI relation iff: $\Sigma \vdash \psi$ iff $\Sigma' \vdash \psi$.

Possibilistic consequence insensitivity (PCI): An inference relation \vdash is said to be a PCI relation iff $\forall \phi$ such that $\pi(\Sigma) \vdash \phi$ where ϕ has a best argument of rank i , $\Sigma \vdash \psi$ iff $\Sigma \cup \{(\phi \ i)\} \vdash \psi$.

The following array summarizes the syntax-sensitivity of the inference relations:

	RI	LCI	CFI	PCI
π -consequence	Yes	Yes	Yes	Yes
lo-consequence	Yes	No	Yes	Yes
Incl-consequence	Yes	No	No	Yes
Lex-consequence	No	No	No	Yes
Argued-consequence	Yes	No	No	No
SS-consequence	Yes	No	No	No

Notice that all the consequence relations, except the Lex-consequence one, are insensitive when we duplicate formulas in the belief bases. Namely, duplicating formulas in the belief base may cause the deletion of some subbase which belongs to $\text{Lex}(\Sigma)$. The possibilistic consequence relation satisfies all the four above properties, which means that it is entirely independent from the syntactic nature of the belief base, in contrast with the remaining consequence relations which are neither LCI relations nor CFI relations. The failure of these two properties shows how much these consequence relations are syntax-sensitive. This study clarifies the ambiguity between syntactic and model-theoretic approaches to the handling of inconsistency, and can be used as a criterion to select an appropriate approach for a given application; for example if in some application the lack of equivalence between $\{\phi\}$ and $\{\phi, \phi\}$ makes sense then Lex-consequence may be preferred. Future research will pursue the investigation on the meaning of syntax-sensitivity. For

instance one might study if adding to Σ one of its consequences (in the sense of one of the inconsistency tolerant inferences presented here) alters the set of consequences of Σ . This is closely related to the properties of cautious monotony and cut of system P.

6 Conclusion

It does not always make sense to revise an inconsistent knowledge base. In the case of multiple sources of information, revision always comes down to destroying part of the knowledge. This paper suggests that it is not even necessary to restore consistency in order to make sensible inferences from an inconsistent knowledge base. The argumentation inference can derive conclusions with reasons to believe them. It is not conservative, it is reasonably syntax-dependent, and it does not inhibit pieces of information. The safely supported inference proposed here only delivers safe conclusions, while the other argued consequences are more debatable since any three of them can be globally inconsistent (Benferhat et al., 1993), because based on antagonistic arguments. Of course, the AND property is lost, and this is the price paid for living in an inconsistent world. However the set of safely supported consequences is consistent and it is possible to close it deductively. As shown in Benferhat et al. (1994), what is then obtained by this deductive closure is the set of possibilistic consequences of $\Sigma^* = \text{Free}(\Sigma_1) \cup \dots \cup \text{Free}(\Sigma_n)$, which shows an example where a coherence approach is more adventurous than a foundation one.

Acknowledgements: This work is partially supported by the European project ESPRIT-BRA n° 6156 DRUMS-II (Defeasible Reasoning and Uncertainty Management Systems).

References

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

S. Benferhat, D. Dubois, H. Prade (1993a) Argumentative inference in uncertain and inconsistent knowledge bases. *Proc. UAI'93*, 411-419.

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

S. Benferhat, D. Dubois, H. Prade (1994) Some syntactic approaches to the handling of inconsistent knowledge bases: A comparative study. Research Report IRIT/94-55-R (revised version), IRIT, Toulouse.

S. Benferhat, D. Dubois, H. Prade (1995) A local approach to reasoning under inconsistency in stratified knowledge bases. *Proc. 3rd ECSQARU Europ. Conf.*, Fribourg, Switzerland.

G. Brewka (1989) Preferred subtheories: An extended logical framework for default reasoning. *Proc. IJCAI'89*, 1043-1048.

C. Cayrol, (1995) On the relation between argumentation and non-monotonic coherence-based entailment. In this proceedings.

C. Cayrol, M.C. Lagasquie-Schiex (1993) Comparaison de relations d'inférence nonmonotone: Etude de complexité. Research Report IRIT/93-23-R, IRIT, Toulouse.

L. Cholvy (1993) Proving theorems in a multi-source environment. *Proc. IJCAI'93*, 66-71.

N.C.A. Da Costa (1963) Calculs propositionnels pour les systèmes formels inconsistants. *Compte Rendu Acad. des Sciences (Paris)*, 257:3790-3792.

J. De Kleer (1990) Using crude probability estimates to guide diagnosis. *Artificial Intelligence*, 45:381-391.

A. Del Val (1994) On the relation between the coherence and foundation theories of belief revision. *Proc. AAAI'94*, 909-914.

J. Doyle (1992) Reason maintenance and belief revision: foundations vs coherence theories. In: *Belief Revision* (P. Gärdenfors, ed.), Cambridge Univ. Press, 29-51.

D. Dubois, J. Lang, H. Prade (1992a) Inconsistency in possibilistic knowledge bases — To live or not live with it. In: *Fuzzy Logic for the Management of Uncertainty* (L.A. Zadeh, J. Kacprzyk, eds.), Wiley, 335-351.

D. Dubois, J. Lang, H. Prade (1992b) Dealing with multi-source information in possibilistic logic. *Proc. ECAI'92*, 38-42.

D. Dubois, J. Lang, H. Prade (1994) Possibilistic logic. In: *Handbook of Logic in Artificial Intelligence and Logic Programming, Vol. 3* (D.M. Gabbay et al., eds.), Oxford Univ. Press, 439-513.

D. Dubois, H. Prade (1991) Epistemic entrenchment and possibilistic logic. *Artificial Intelligence*, 50:223-239.

M. Elvang-Goransson, P. Krause, J. Fox (1993) Dialectic reasoning with inconsistent information. *Proc. UAI'93*, 114-121.

R. Fagin, J.D. Ullman, M.Y. Vardi (1983) On the semantics of updates in database. *Proc. 2nd ACM SIGACT-SIGMOD Symp. on the Principles of Databases Systems*, Atlanta, 352-365.

P. Gärdenfors (1988) *Knowledge in Flux — Modeling the Dynamic of Epistemic States*. MIT Press, Cambridge, MA.

P. Gärdenfors (1990) The dynamics of belief systems: Foundations v. coherence theories. *Revue Internationale de Philosophie*, 172:24-46.

G. Harman (1986) *Change in View, Principles of Reasoning*. MIT Press, Cambridge, MA.

J. Lang (1994) Syntax-based default reasoning as probabilistic model-based diagnosis. *Proc. UAI'94*, 391-298.

D. Lehmann (1993) Another perspective on default reasoning. Technical Report, Hebrew University, Jerusalem.

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

J. Lin (1994) Consistent belief reasoning in the presence of inconsistency. *Proc. TARK'94*.

B. Nebel (1991) Belief revision and default reasoning: Syntax-based approaches. *Proc. KR'91*, 417-428.

B. Nebel (1994) Base revision operator ans schemes: Semantics representation and complexity. *Proc. ECAI'94*, 341-345.

J.L. Pollock (1987) Defeasible reasoning. *Cognitive Science*, 11:481-518.

D. Poole (1985) On the comparison of theories: preferring the most specific explanation. *Proc. IJCAI'85*, 144-147.

A.S. Rao, N.Y. Foo (1989) Formal theories of belief revision. *Proc. KR'89*, 369-380.

J.A. Reggia, D.S. Nau, P.Y. Wang, H. Peng (1985) A formal model of diagnostic inference. *Information Sciences*, 37:227-285.

N. Rescher, R. Manor (1970) On inference from inconsistent premises. *Theory and Decision*, 1:179-219.

G.R. Simari, R.P. Loui (1992) A mathematical treatment of defeasible reasoning and its implementation. *Artificial Intelligence*, 53:125-157.

G. Shafer (1976) *A Mathematical Theory of Evidence*. Princeton Univ. Press, Princeton, NJ.

S. Toulmin (1956) *The Uses of Argument*. Cambridge Univ. Press, Cambridge, UK.