Four-Valued Logics for Reasoning with Uncertainty in Prioritized Data

Ofer Arieli
Department of Computer Science
School of Mathematical Sciences
Tel-Aviv University
Tel-Aviv 69978, Israel.
Email: ofera@math.tau.ac.il

Abstract: We present an approach for drawing plausible conclusions from inconsistent and incomplete knowledge-bases, which may also be prioritized. Our method is based on a four-valued semantics that is particularly suitable for reasoning with uncertainty. Our inference mechanism is closely related to some other well-known formalisms for handling inconsistent data, such as reasoning with maximal consistent subsets and possibilistic logic. It is shown that the formalism presented here is nonmonotonic, paraconsistent, and is capable of managing ranked data without having the "drowning problem".

1 Introduction

The purpose of this work is to present plausible methods for reasoning with uncertainty. The approaches considered here are based on a four-valued semantics, which seems to be particularly suitable for our goal, since it contains truth values that can intuitively be understood as representing incomplete and inconsistent information.

Given a knowledge-base *KB*, the idea is to construct a set of "worlds", which are possible representations of the information provided in *KB*. Each one of these worlds is consistent, and preserves the semantics of *KB* in the sense that it has a model in which all the assertions of *KB* are true. As such, a possible world description of *KB* might not be a maximal consistent subset of *KB*, and so we trade maximal size considerations with an obligation to closely reflect the semantics of *KB*. A formula is accepted as a consequence of *KB* when it can be inferred from all the possible world descriptions of *KB*.

In the second part of the paper we extend the method mentioned above to cases in which the knowledge-bases under consideration are prioritized. It is shown that like many other formalisms for reasoning with inconsistent and prioritized data (e.g., [1, 8, 9, 12, 17, 18]) our method is also nonmonotonic and paraconsistent [11]. In addition, unlike some other formalisms like system Z [17], possibilistic logic [12], and default reasoning with conditional objects [10], the present approach is capable of managing ranked data without having the "drowning problem" [9, 10].

2 Background

2.1 The algebraic structure

The approach that we consider here is based on Belnap's well-known algebraic structure, *FOUR*, presented in [6, 7]. This structure contains four elements: the two classical values, t and f, and two other values, $\bot$ and $\top$, that respectively denote lack of information and "too much" information, i.e.: contradictions.

The main idea is to arrange the elements of *FOUR* in two partial orders: One, $\leq_t$, is usually understood as reflecting differences in the "measure of truth" that each value represents. According to this order f is the minimal element, t is the maximal one, and $\bot, \top$ are two intermediate values that are incomparable. $(\{t, f, \top, \bot\}, \leq_t)$ is a distributive lattice with an order reversing involution $\neg$, for which $\neg\top = \top$ and $\neg\bot = \bot$. The meet and the join of this lattice are denoted by $\wedge$ and $\vee$, respectively. The other partial order, $\leq_k$, is intuitively understood as reflecting differences in the amount of *knowledge* that each truth value exhibits. In this partial order $\bot$ is the minimal element, $\top$ is the maximal element, and t, f are incomparable.[1]

2.2 Syntax and semantics

The language we treat here is the standard propositional one. Atomic formulae are denoted by p, q, and complex formulae are denoted by ψ, ϕ. Given a set S of formulae, we shall write $\mathcal{A}(S)$ to denote the set of the atomic formulae that occur in S. $\mathcal{L}(S)$ denotes the set of the literals that occur in S.

The various semantic notions are defined on *FOUR* as natural generalizations of similar classical notions: A *valuation* ν is a function that assigns a truth value from *FOUR* to each atomic formula. Any valuation is extended to complex formulae in the obvious way. The set of the four-valued valuations is denoted by $\mathcal{V}$. We will sometimes write $\psi : b \in \nu$ instead of $\nu(\psi) = b$. A valuation ν *satisfies* ψ iff $\nu(\psi) \in \{t, \top\}$. t and $\top$ are called the *designated* elements of *FOUR*. A valuation that satisfies every formula in a given set S of formulae is a *model* of S. A model of S will usually be denoted by M or N. The set of all the models of S is denoted by $mod(S)$.

The formulae that will be considered here are clauses, i.e.: disjunctions of literals. A useful property of clauses is given in the following lemma:

Lemma 2.1 Let ψ be a clause and ν a valuation. Then $\nu(\psi) \in \{t, \top\}$ iff $\exists l \in \mathcal{L}(\psi)$ s.t. $\nu(l) \in \{t, \top\}$.
Proof: By an induction on the structure of ψ. □

A set of clauses is called a *knowledge-base*, and is denoted by *KB*. As the following lemma shows, representing formulae in a clause form does not reduce the generality.

Lemma 2.2 [3] For every formula ψ there is a finite set S of clauses such that for every valuation ν, $\nu(\psi) \in \{\top, t\}$ iff $\nu(\phi) \in \{\top, t\}$ for every $\phi \in S$.

[1] $(\{t, f, \top, \bot\}, \leq_k)$ is also a lattice, and so a $\leq_k$-meet and a $\leq_k$-join operations might be defined on *FOUR* as well (see, e.g., [2, 4, 6, 7, 13, 14, 16]).

2.3 Measurement of consistency

Notation 2.3 $\mathit{Inc}(\nu) = \{p \mid \nu(p) = \top\}$.

Definition 2.4 Let $\nu_1, \nu_2 \in \mathcal{V}$.
a) ν_1 is *more consistent* than ν_2 iff $\mathit{Inc}(\nu_1) \subset \mathit{Inc}(\nu_2)$.
b) $M \in \mathit{mod}(KB)$ is a *most consistent model* of KB (mcm, for short) if there is no other model of KB which is more consistent than M. The set of the most consistent models of KB is denoted by $mcm(KB)$.

Definition 2.5 A valuation ν is *consistent* if $\mathit{Inc}(\nu) = \emptyset$. A knowledge-base KB is consistent if it has a consistent model.

Proposition 2.6 [3, 5] KB is consistent iff it is classically consistent.

2.4 The basic consequence relation

Definition 2.7 [2] $KB \models_4 \psi$ if every four-valued mcm of KB satisfies ψ.[2]

Example 2.8 $\neg p, p \vee q \models_4 q$ (the only mcm here is $M(p) = f$, $M(q) = t$), while $p, \neg p, p \vee q \not\models_4 q$ (a counter-mcm: $N(p) = \top$, $N(q) = \bot$). This example shows, in particular, that $\models_4$ is nonmonotonic and paraconsistent. It also demonstrates the usefulness of considering only the mcms of a given knowledge-base rather than all it's models; In the latter case $\{\neg p, p \vee q\}$ does *not* entail q, and so the Disjunctive Syllogism is always violated (even in cases in which the relevant formulae are not involved in any conflict).

Denote by $\models_2$ the classical consequence relation. Unlike $\models_2$, in the standard propositional language there are no tautologies w.r.t. $\models_4$. This follows from the fact that if $\forall p \in \mathcal{A}(\psi)$ $\nu(p) = \bot$ then $\nu(\psi) = \bot$ as well. However, it is sometimes possible to draw the same conclusions when using either $\models_2$ or $\models_4$:

Lemma 2.9 [4] Let KB be a consistent knowledge-base, and let ψ be a clause that does not contain any atomic formula and its negation.[3] Then $KB \models_2 \psi$ iff $KB \models_4 \psi$.

Corollary 2.10 If KB is a consistent knowledge-base and ψ is a formula in a conjunctive normal form that none of its conjuncts is a classical tautology, then $KB \models_2 \psi$ iff $KB \models_4 \psi$.

Corollary 2.11 Suppose that $KB \cup \{\phi\}$ is a consistent knowledge-base and ψ is a clause that does not contain any atomic formula and its negation. Then $KB \models_4 \psi$ implies that $KB, \phi \models_4 \psi$.
Proof: By Lemma 2.9, and since $\models_2$ is monotonic. $\square$

[2]This relation is denoted by $\models_{con}$ in [2, 5] and $\models^4_{\mathcal{I}}$ in [4].
[3]I.e., ψ is not a classical tautology.

3 Reasoning with inconsistent data

3.1 World settings and the inference relation

Definition 3.1 A subset $S \subseteq KB$ is *consistent in the context of* KB if S has a consistent model N, and there is a (not necessarily consistent) model M of KB s.t. $\forall p \in \mathcal{A}(S)\ M(p) = N(p)$.

Definition 3.2 [1, 3, 5] A *possible world setting* of a knowledge-base KB is a nonempty maximal subset of KB that is consistent in the context of KB.[4] The set of all the possible world settings of KB is denoted by $\mathcal{W}(KB)$.

Note that the elements of $\mathcal{W}(KB)$ are not necessarily *maximal* consistent subsets of KB. This is so since they should preserve the semantics of the knowledge-base, while the maximal consistent subsets might not do so. For example, the simplest inconsistent knowledge-base $KB = \{p, \neg p\}$ contains two maximal consistent subsets $\{p\}$ and $\{\neg p\}$, but neither of them truly reflects the intended meaning of KB. Moreover, each one of them even *contradicts an explicit assertion* of KB (see also Section 3.3).

Definition 3.3 [3, 5] The set that is *associated with* a valuation ν is defined as follows: $S_\nu(KB) = \{\psi \in KB \mid \nu(\psi) = t,\ \mathcal{A}(\psi) \cap \mathit{Inc}(\nu) = \emptyset\}$.

Example 3.4 Consider the knowledge-base $KB = \{p,\ q,\ h,\ \neg p \vee \neg q\}$. Then $\mathcal{W}(KB) = \{S_1, S_2\}$, where $S_1 = \{p, h\}$ and $S_2 = \{q, h\}$. These sets are associated with the (most consistent) models $\{p : t,\ q : \top,\ h : t\}$ and $\{p : \top,\ q : t,\ h : t\}$, respectively. Note that S_1 is no longer a possible world setting of $KB' = KB \cup \{\neg p\}$, since there is no *consistent* model M of S_1 and some model M' of KB' s.t. $M(s) = M'(s)$ for every $s \in \mathcal{A}(KB')$.

As the following proposition shows, there is a strong connection between the possible world settings of a knowledge-base KB and its mcms. In particular, every possible world setting of KB is associated with some mcm of KB:

Proposition 3.5 [3, 5]
a) For every $S \in \mathcal{W}(KB)$ there is an mcm M of KB s.t. $S = S_M(KB)$.
b) For every mcm M of KB there is an $S \in \mathcal{W}(KB)$ s.t. $S_M(KB) \subseteq S$.

Corollary 3.6 Let $W(KB) = \{S_M(KB) \mid M \in mcm(KB)\}$. Then $\mathcal{W}(KB) = \{S \in W(KB) \mid \neg\exists T \in W(KB) \text{ s.t. } S \subset T\}$.

Definition 3.7 $KB \models_{\mathcal{W}} \psi$ if $\forall S \in \mathcal{W}(KB)\ S \models_4 \psi$.

Example 3.8 Consider the knowledge-base of Example 3.4. Then $KB \not\models_{\mathcal{W}} p$, $KB \not\models_{\mathcal{W}} q$, and $KB \models_{\mathcal{W}} h$. This might be explained by the fact that unlike p, q, the assertion h is not involved in any conflict in KB, and so it is a "reliable" conclusion of KB (see also Proposition 3.12 below).

[4] In [1, 3, 5] these sets are called *recovered knowledge-bases*.

3.2 Basic properties of $\models_{\mathcal{W}}$

Proposition 3.9 If KB is consistent then $KB \models_{\mathcal{W}} \psi$ iff $KB \models_4 \psi$.
Proof: Follows from the fact that in this case $\mathcal{W}(KB) = \{KB\}$. $\square$

Proposition 3.10 $\models_{\mathcal{W}}$ is nonmonotonic and paraconsistent.
Proof: For instance, $p, q \models_{\mathcal{W}} q$, but $p, q, \neg q \not\models_{\mathcal{W}} q$. $\square$

The proof of the last proposition also shows that $\models_{\mathcal{W}}$ is not reflexive.[5] However, in many reasoning systems (especially those for drawing nontrivial conclusions from inconsistent data) reflexivity is not valid in general (see, e.g., [15, 19]). Still, as it clearly follows from Proposition 3.12 below, $\models_{\mathcal{W}}$ *is* reflexive w.r.t. premises that are true in every possible world setting:

Definition 3.11 $Con(KB) = \bigcap\{S \mid S \in \mathcal{W}(KB)\}$.

Proposition 3.12 Let ψ be a clause that does not contain any atomic formula and its negation. If $Con(KB) \models_4 \psi$ then $KB \models_{\mathcal{W}} \psi$.
Proof: Note, first, that since there are no tautologies w.r.t. $\models_4$ in the propositional language, the condition of the proposition assures that $Con(KB) \neq \emptyset$. Now, since $Con(KB) \subseteq S$ for every $S \in \mathcal{W}(KB)$, then by Corollary 2.11 $\forall S \in \mathcal{W}(KB)\ S \models_4 \psi$. Thus $KB \models_{\mathcal{W}} \psi$. $\square$

The converse of the last proposition is not true: In Examples 3.4 and 3.8, for instance, $Con(KB) = \{h\}$ and so although $Con(KB) \not\models_4 p \vee q$, still $KB \models_{\mathcal{W}} p \vee q$.

Proposition 3.13 If $\mathcal{W}(KB) \neq \emptyset$ and $KB \models_{\mathcal{W}} \psi$, then $KB \not\models_{\mathcal{W}} \neg\psi$.
Proof: Every element in $\mathcal{W}(KB)$ is consistent, and so if S is a possible wold setting s.t. $S \models_4 \psi$, then $S \not\models_4 \neg\psi$. Hence $KB \not\models_{\mathcal{W}} \neg\psi$. $\square$

3.3 Inference with maximal consistent sets

A famous approach for reasoning with uncertainty accepts formulae provided that they classically follow from all the maximal consistent subsets of a given knowledge-base. Denote by $\models_{\mathcal{MC}}$ the corresponding consequence relation. Then $\models_{\mathcal{W}}$ is usually at least as cautious as $\models_{\mathcal{MC}}$:

Proposition 3.14 If $Con(KB) \neq \emptyset$ and ψ is a clause that does not contain an atomic formula and its negation, then $KB \models_{\mathcal{W}} \psi$ implies that $KB \models_{\mathcal{MC}} \psi$.
Proof: Denote by $\mathcal{MC}(KB)$ the set of the maximal consistent subsets of KB. If $KB \not\models_{\mathcal{MC}} \psi$ then $\exists T \in \mathcal{MC}(KB)$ s.t. $T \not\models_2 \psi$. By Lemma 2.9, then, $T \not\models_4 \psi$ as well. Since T is a maximal consistent subset of KB, and $Con(KB)$ is an intersection of consistent subsets of KB, then $Con(KB) \subseteq T$. But $Con(KB) \neq \emptyset$, thus there is a nonempty subset of T that is consistent in the context of KB, and so there is a set which is maximal among the subsets of T that are consistent in the context of KB. Denote this set by S. Since $T \not\models_4 \psi$ then by Corollary 2.11, $S \not\models_4 \psi$ either. To conclude it is left to show, therefore, that $S \in \mathcal{W}(KB)$. Indeed, otherwise there is a set $S' \in \mathcal{W}(KB)$ s.t. $S \subset S'$. Thus $\exists \phi \in S' \setminus S$ s.t.

[5] Hence, in particular, $\models_{\mathcal{W}}$ is not the same as $\models_4$.

$S \cup \{\phi\}$ is consistent in the context of *KB* (since $S \cup \{\phi\} \subset S'$), and so $\phi \notin T$ (otherwise $S \cup \{\phi\}$ would have been a subset of T that is consistent in the context of *KB* and properly contains S – a contradiction to the choice of S). Since T is a maximal subset of *KB* that is classically consistent, necessarily $T \cup \{\phi\}$ is classically inconsistent. Hence $T \models_2 \neg\phi$. By Corollary 2.10 and since De-Morgan's rules are valid in *FOUR*, $T \models_4 \neg\phi$. Now, by Proposition 2.6 S' is in particular classically consistent. So let M be a classical model of S', and let $N \in mod(KB)$ s.t. $\forall p \in \mathcal{A}(S')\ N(p) = M(p)$. Since $\phi \in S'$, $M(\phi) = t$. Thus $N(\phi) = t$ as well. On the other hand, N is also a model of T, and $T \models_4 \neg\phi$, therefore $N(\phi) \in \{f, \top\}$ – a contradiction. □

The following proposition shows that our reasoning process is analogous in spirit to that of $\models_{\mathcal{MC}}$: Instead of making classical conclusions from (all) the maximal consistent subsets, we draw classical conclusions from (all) the possible world settings.

Proposition 3.15 Let ψ be a clause which is not a classical tautology. Then $KB \models_{\mathcal{W}} \psi$ iff ψ classically follows from every possible world setting of *KB*.
Proof: $KB \models_{\mathcal{W}} \psi$ iff $\forall S \in \mathcal{W}(KB)\ S \models_4 \psi$, iff $\forall S \in \mathcal{W}(KB)\ S \models_2 \psi$ (Lemma 2.9). □

4 Prioritized knowledge-bases

4.1 Motivation and basic definitions

In many cases a knowledge-base contains formulae with different importance or certainty. For instance, rules that state default assumptions are usually considered as less reliable than rules without exceptions. Also, inference rules are usually given a lower priority than atomic facts. These kinds of considerations are particularly common when reasoning with inconsistent knowledge-bases; If some formulae are more certain than others, one would probably like to reject the least certain first.

A common approach for making precedences among formulae is to assign them *ranks*. Different ranks reflect differences in the certainty or reliability attached to the assertions (see, e.g., [8, 9, 12, 17, 18]). In what follows we shall we use this additional data for refining the inference mechanism discussed in the previous section.

Definition 4.1 A *ranking* of a knowledge-base *KB* is a function r from the clauses in *KB* to $\{1, 2, \ldots, n\}$.

The ranking function determines a preference relation on the clauses of a knowledge-base. Intuitively, a clause with a lower rank has a higher priority.

Notation 4.2 $KB_i = \{\psi \in KB \mid r(\psi) \leq i\}$.

Definition 4.3 [1]
a) $W_i(KB) = \{S_\nu(KB) \mid \nu \in mcm(KB_i)\}$.
b) $\mathcal{W}_i(KB) = \{S \in W_i(KB) \mid \neg\exists T \in W_i(KB) \text{ s.t. } S \subset T\}$. [6]

Each $\mathcal{W}_i(KB)$ is a set of possible worlds that correspond to the situation described in KB. Following [8] we provide some criteria for choosing the preferred set of worlds:

- *set cardinality:* $\mathcal{W}_i \geq_{\rm sc} \mathcal{W}_j$ iff $\forall S \in \mathcal{W}_i\ \exists T \in \mathcal{W}_j$ s.t. $|T| \leq |S|$.
- *set inclusion:* $\mathcal{W}_i \geq_{\rm si} \mathcal{W}_j$ iff $\forall S \in \mathcal{W}_i\ \exists T \in \mathcal{W}_j$ s.t. $T \subseteq S$.
- *cardinality of consistent consequences:* $\mathcal{W}_i \geq_{\rm cc} \mathcal{W}_j$ iff $\forall S \in \mathcal{W}_i\ \exists T \in \mathcal{W}_j$ s.t. $|\{l \in \mathcal{L}(KB) \mid T \models l, T \not\models \bar{l}\}| \leq |\{l \in \mathcal{L}(KB) \mid S \models l, S \not\models \bar{l}\}|$.[7]
- *inclusion of consistent consequences:* $\mathcal{W}_i \geq_{\rm ci} \mathcal{W}_j$ iff $\forall S \in \mathcal{W}_i\ \exists T \in \mathcal{W}_j$ s.t. $\{l \in \mathcal{L}(KB) \mid T \models l, T \not\models \bar{l}\} \subseteq \{l \in \mathcal{L}(KB) \mid S \models l, S \not\models \bar{l}\}$.

Definition 4.4 Let $\leq$ be a preference criterion among $\mathcal{W}_i(KB)$, $i = 1, \ldots, n$. The *optimal recovery level* of KB w.r.t. $\leq$ is $i_0 = \max\{i \mid \neg\exists j \neq i \text{ s.t. } \mathcal{W}_j \geq \mathcal{W}_i\}$.

The induced consequence relation is a natural generalization of $\models_{\mathcal{W}}$ (cf. Definition 3.7):

Definition 4.5 Let i_0 be the optimal recovery level of KB w.r.t. $\leq$. Then $KB \models_{\leq\mathcal{W}} \psi$ if $\forall S \in \mathcal{W}_{i_0}(KB)\ S \models_4 \psi$.

4.2 An example – Tweety dilemma

For the following example we first extend the discussion to languages with predicates and variables. It is possible to do so in a straightforward way, provided that each clause that contains variables is considered as universally quantified. Consequently, a knowledge-base containing a non-grounded formula, ψ, will be viewed as representing the corresponding set of ground formulae formed by substituting each variable in ψ with every possible element of the Herbrand universe, U. Formally: $KB^U = \{\rho(\psi) \mid \psi \in KB,\ \rho : var(\psi) \to U\}$.

Example 4.6 Consider the following well-known puzzle:

$$\begin{array}{lll} bird(Tweety), & bird(Fred), & penguin(Tweety), \\ bird(x) \to fly(x), & penguin(x) \to \neg fly(x), & penguin(x) \to bird(x). \end{array}$$

Denote this set by KB, and let T, F, b, p, f abbreviate, respectively, the individuals Tweety, Fred, and the predicates *bird*, *penguin*, and *fly*.[8]

The three mcms of KB are the following:

[6] By Corollary 3.6 the only difference from the non-prioritized case is that here the relevant valuations are mcms of KB_i rather than mcms of KB.

[7] As usual, $\bar{l}$ denotes the complement of l, and $KB \models \psi$ denotes that $\forall M \in mod(KB)\ M(\psi)$ is designated.

[8] Note that the symbol f has double meanings here: abbreviating the predicate fly, and representing the truth value *false*. Each occurrence of f will be understood by the context.

$$M1 = \{b(T):t,\ p(T):t,\ f(T):\top,\ b(F):t,\ p(F):f,\ f(F):t\},$$
$$M2 = \{b(T):t,\ p(T):\top,\ f(T):t,\ b(F):t,\ p(F):f,\ f(F):t\},$$
$$M3 = \{b(T):\top,\ p(T):t,\ f(T):f,\ b(F):t,\ p(F):f,\ f(F):t\}.$$

Each mcm has a set that is associated with it:

$$S_{M1}(KB) = KB^U \setminus \{\psi \in KB^U \mid f(T) \in \mathcal{A}(\psi)\},$$
$$S_{M2}(KB) = KB^U \setminus \{\psi \in KB^U \mid p(T) \in \mathcal{A}(\psi)\},$$
$$S_{M3}(KB) = KB^U \setminus \{\psi \in KB^U \mid b(T) \in \mathcal{A}(\psi)\}.$$

KB is obviously inconsistent, so it is useless as far as classical logic is concerned. Nevertheless, the ambiguous data is related only to the information about Tweety. $\models_{\mathcal{W}}$ "salvages" the consistent part of KB, and allows us to draw nontrivial conclusions about Fred, despite the inconsistency:

$$KB \models_{\mathcal{W}} b(F), \qquad KB \models_{\mathcal{W}} f(F), \qquad KB \models_{\mathcal{W}} \neg p(F),$$
$$KB \not\models_{\mathcal{W}} \neg b(F), \qquad KB \not\models_{\mathcal{W}} \neg f(F), \qquad KB \not\models_{\mathcal{W}} p(F).$$

Since the information about Tweety is inconsistent, $\models_{\mathcal{W}}$ does not allow us to infer nontrivial conclusions about Tweety. We claim, however, that this state of affairs is due to the representation of the problem that does not properly reflect our intuitive understanding of this particular puzzle: The two inference rules that concern with penguins are more specific than the assertion $bird(x) \to fly(x)$. Also, unlike this latter assertion, the former assertions do *not* have exceptions. We claim, therefore, that a more accurate representation of this problem should reflect precedences among the rules of KB. In our case we can do so by using a ranking function r. A reasonable ranking of KB is the following: $r(b(T)) = r(b(F)) = r(p(T)) = 1$, $r(p(x) \to \neg f(x)) = r(p(x) \to b(x)) = 2$, and $r(b(x) \to f(x)) = 3$.

By Proposition 4.11 below, when the preference criterion is either $\leq_{\text{ci}}$ or $\leq_{\text{cc}}$, the optimal recovery level is $i_0 = 2$. In this case,

$$KB_2 = \{b(T),\ b(F),\ p(T),\ p(x) \to \neg f(x),\ p(x) \to b(x)\}.$$

The most consistent models of KB_2 are the following:

$$N1 = \{b(T):t,\ p(T):t,\ f(T):f,\ b(F):t,\ p(F):f,\ f(F):f\},$$
$$N2 = \{b(T):t,\ p(T):t,\ f(T):f,\ b(F):t,\ p(F):f,\ f(F):t\},$$
$$N3 = \{b(T):t,\ p(T):t,\ f(T):f,\ b(F):t,\ p(F):f,\ f(F):\bot\},$$
$$N4 = \{b(T):t,\ p(T):t,\ f(T):f,\ b(F):t,\ p(F):t,\ f(F):f\},$$
$$N5 = \{b(T):t,\ p(T):t,\ f(T):f,\ b(F):t,\ p(F):\bot,\ f(F):f\}.$$

Since $S_{N2} = KB^U \setminus \{b(T) \to f(T)\}$, while $S_{N1} = S_{N3} = S_{N4} = S_{N5} = KB^U \setminus \{b(T) \to f(T),\ b(F) \to f(F)\}$, it follows that $\mathcal{W}_2(KB) = \{S_{N2}(KB)\}$. Thus, according to $\models_{\leq_{cc}\mathcal{W}}$ and $\models_{\leq_{ci}\mathcal{W}}$ one can deduce that Tweety is a bird, a penguin, and cannot fly, while Fred is a bird that can fly and it is not a penguin. The converse assertions are *not deducible*, as expected. It also can be shown that these conclusions are obtained by $\models_{\leq_{sc}\mathcal{W}}$ and by $\models_{\leq_{si}\mathcal{W}}$.

4.3 Basic properties of $\models_{\leq\mathcal{W}}$

First we show that $\models_{\leq\mathcal{W}}$ is an extension of $\models_{\mathcal{W}}$:

Proposition 4.7 When all the clauses in KB have the same priority, then $KB \models_{\leq\mathcal{W}} \psi$ iff $KB \models_{\mathcal{W}} \psi$.
Proof: By Corollary 3.6 and Definition 4.3, since in this case $KB = KB_1$. $\Box$

Some basic properties of $\models_{\mathcal{W}}$ remain valid also in the case of $\models_{\leq\mathcal{W}}$:

Proposition 4.8 If KB is consistent then $KB \models_{\leq\mathcal{W}} \psi$ iff $KB \models_4 \psi$.
Proof: Let n be the maximal rank in KB. If KB is consistent then the optimal recovery level w.r.t. either $\leq_{\rm sc}$, $\leq_{\rm si}$, $\leq_{\rm cc}$, or $\leq_{\rm ci}$ is n, and $KB_n = \{KB\}$. The claim now immediately follows from Definition 4.5. $\Box$

By the last proposition and Lemma 2.9 it follows that if KB is consistent and ψ is a clause that is not a classical tautology, then $KB \models_{\leq\mathcal{W}} \psi$ iff ψ classically follows from KB.

Proposition 4.9 $\models_{\leq\mathcal{W}}$ is nonmonotonic and paraconsistent.
Proof: The same as that of Proposition 3.10, with r s.t. either $r(p) < r(q)$ or $r(\neg q) < r(q)$. $\Box$

Proposition 4.10 Let i_0 be the optimal recovery level of KB w.r.t. $\leq$. If $\mathcal{W}_{i_0} \neq \emptyset$ and $KB \models_{\leq\mathcal{W}} \psi$ then $KB \not\models_{\leq\mathcal{W}} \neg\psi$.
Proof: The same as that of 3.13, replacing $\mathcal{W}(KB)$ with $\mathcal{W}_{i_0}(KB)$. $\Box$

In the rest of this paper, unless otherwise stated, we will use either $\leq_{cc}$ or $\leq_{ci}$ as the preference criterion, and so $\models_{\leq\mathcal{W}}$ will abbreviate either $\models_{\leq_{cc}\mathcal{W}}$ or $\models_{\leq_{ci}\mathcal{W}}$. Also, i_0 will henceforth denote the optimal recovery level w.r.t. either one of these criteria. Finally, in what follows we assume that the set of the assertions with the highest priority (i.e. KB_1) is consistent.

Proposition 4.11 [1] $i_0 = \max\{i \mid KB_i \text{ is consistent}\}$.

Proposition 4.12 Let $S \in \mathcal{W}_{i_0}(KB)$. Then there is a (most) consistent model M of KB_{i_0} s.t. $S = KB_{i_0} \cup S_M(KB \setminus KB_{i_0})$.
Proof: By Definition 4.3, $S = S_M(KB)$ for some $M \in mcm(KB_{i_0})$. Thus, $S = S_M(KB_{i_0}) \cup S_M(KB \setminus KB_{i_0})$. But by Proposition 4.11 KB_{i_0} is consistent, and so $S_M(KB_{i_0}) = KB_{i_0}$. It follows that $S = KB_{i_0} \cup S_M(KB \setminus KB_{i_0})$. $\Box$

Definition 4.13 $Con_{i_0}(KB) = \bigcap\{S \mid S \in \mathcal{W}_{i_0}(KB)\}$.

Proposition 4.14 Let ψ be a clause that is not a classical tautology. Then:
a) If $KB_{i_0} \models_4 \psi$ then $KB \models_{\leq\mathcal{W}} \psi$.
b) If $Con_{i_0}(KB) \models_4 \psi$ then $KB \models_{\leq\mathcal{W}} \psi$.
Proof: First note that since there are no tautologies w.r.t. $\models_4$, the conditions of parts (a) and (b) assure (respectively) that $KB_{i_0} \neq \emptyset$ and $Con_{i_0}(KB) \neq \emptyset$. Now, part (a) follows from the fact that $\forall S \in \mathcal{W}_{i_0}(KB)$ $KB_{i_0} \subseteq S$ (Proposition 4.12). Thus, since $KB_{i_0} \models_4 \psi$, by Corollary 2.11 $\forall S \in \mathcal{W}_{i_0}(KB)$ $S \models_4 \psi$. The proof of (b) is similar, and follows from the fact that $\forall S \in \mathcal{W}_{i_0}(KB)$ $Con_{i_0}(KB) \subseteq S$. $\Box$

By Proposition 4.14 it follows that $\models_{\leq_{cc}\mathcal{W}}$ and $\models_{\leq_{ci}\mathcal{W}}$ preserve the semantics of the clauses with the i_0-highest priorities (see also Corollary 4.16 below). In addition, it is possible to deduce conclusions that are based on assertions with lower priorities than the optimal recovery level, provided that they are not involved in any conflict. In Example 4.6, for instance, $b(x) \rightarrow f(x)$ cannot be inferred in general, since it is causes conflicts when $x = Tweety$. However, the instance $b(F) \rightarrow f(F)$ *is* deducible, since it does not harm the consistency of any possible world setting. In particular this shows that $\models_{\leq\mathcal{W}}$ does not suffer from the so called "drowning effect" (see Section 5.2 below).

Corollary 4.15 Suppose that $\psi \in Con_{i_0}(KB)$. Then $KB \not\models_{\leq\mathcal{W}} \neg\psi$.
Proof: Follows from Propositions 4.14(b) and 4.10. □

Corollary 4.16 If $\psi \in KB_{i_0}$ then $KB \not\models_{\leq\mathcal{W}} \neg\psi$.
Proof: By Corollary 4.15 and the fact that $KB_{i_0} \subseteq Con_{i_0}(KB)$ (see 4.12). □

5 Related systems

5.1 Coherent approaches for restoring consistency

The formalism introduced here takes into account *every* possible world that plausibly represents the intended meaning of KB. In [1, 3, 5], on the other hand, only *one* such world is chosen. This world is considered as the "recovered" version of the "polluted" knowledge-base, and the rest of the data is discarded. Since the recovered knowledge-base is (classically) consistent, it is possible to draw nontrivial classical conclusions from it.[9]

The main difference between the approach of [1, 3, 5] and the present one concerns with the way of treating contradictory data: Here we accept inconsistency and try to cope with it. This approach allows us to make nontrivial conclusions from an inconsistent theory without throwing pieces of information away. The approach of [1, 3, 5] revises inconsistent information and restores consistency. Thus, contradictory data is considered useless, and only a consistent part of the original information is used for making inferences. Such methods are sometimes called *coherent*. See [9] for a survey on coherent techniques for reasoning with prioritized knowledge-bases.

5.2 The possibilistic approach

In [10, 12] Benferhat et al. present a well-known approach for reasoning with inconsistency in prioritized knowledge-bases, called *possibilistic logic*. Briefly, the idea is to consider a consistent subset $\pi(KB)$ of KB, so that in terms of Notation 4.2 $\pi(KB) = KB_i$, where i is the maximal index for which KB_i is classically consistent. A formula ψ is a possibilistic consequence of KB ($KB \models_\pi \psi$) if it classically follows from $\pi(KB)$.

[9]In particular, every conclusion that is deducible by $\models_{\mathcal{W}}$ is also valid according the approach taken in [3, 5], and every conclusion that is drawn by $\models_{\leq\mathcal{W}}$ is valid according to the approach of [1]; The opposite directions are obviously not true.

Proposition 5.1 Suppose that KB_1 is consistent and ψ is a clause which is not a classical tautology. If $KB \models_{\pi} \psi$ then $KB \models_{\leq \mathcal{W}} \psi$.[10]

Proof: If $KB \models_{\pi} \psi$ then $\pi(KB) \models_2 \psi$. But by Propositions 2.6 and 4.11 $\pi(KB) = KB_{i_0}$, so $KB_{i_0} \models_2 \psi$. Since ψ is not a classical tautology and KB_{i_0} is consistent (Proposition 4.11), then by Lemma 2.9 $KB_{i_0} \models_4 \psi$. Hence, by Proposition 4.14(a), $KB \models_{\leq \mathcal{W}} \psi$. $\Box$

The converse of the last proposition is not true. This follows from the fact that unlike the case of $\models_{\leq \mathcal{W}}$, the possibilistic consequence relation has the so called "drowning problem" [9, 10]: Formulae with ranks that are greater than the inconsistency level are inhibited even if they are not involved in any conflict. We demonstrate this phenomenon in the following example:

Example 5.2 Consider again the knowledge-base of Example 3.4, and suppose that $r(\neg p \vee \neg q) = 1$, $r(p) = 2$, $r(q) = 3$, $r(h) = 4$. Then $\pi(KB) = \{\neg p \vee \neg q,\ p\}$, and so $KB \not\models_{\pi} h$, even though h is not involved in the inconsistency. As Proposition 4.14(b) shows, this is not the case with $\models_{\leq \mathcal{W}}$: Since $h \in Con_2(KB)$, then $Con_2(KB) \models_4 h$, and so $KB \models_{\leq \mathcal{W}} h$.

6 Conclusion and further work

We have considered a logic for reasoning with incomplete and inconsistent knowledge-bases. The corresponding consequence relation is nonmonotonic, paraconsistent, and allows to draw conclusions that are not involved in any conflict in the knowledge-base. In practice, this means that in case that a small part of a large knowledge-base is contradictory, one would still be able to draw nontrivial conclusions based on the "robust" part of the knowledge-base, and so the inference process will not be totally damaged by the "spoiled" data.

In the second part of the paper we considered cases in which the formulae of the knowledge-bases are ranked. This additional information allowed us to refine the inference procedure so that we also draw conclusions with contradictory data, provided that this data has a sufficiently high priority.

The formalisms presented here are based on a four-valued semantics, which has been shown to be very useful in reasoning with uncertainty [4, 6, 7]. The next natural step is to use more than just four values. This will allow us, e.g., to view truth values as representing probabilities, confidence factors, etc. One possible way of doing so is to use *bilattices* [16], which are algebraic structures, with *arbitrary* number of truth values, that naturally generalize *FOUR* [2, 13, 14]. Such extensions will be considered in a future work.

Acknowledgment

I would like to thank Arnon Avron for helpful discussions on the topics of this paper.

[10] Recall that here $\models_{\leq \mathcal{W}}$ is either $\models_{\leq_{cc} \mathcal{W}}$ or $\models_{\leq_{ci} \mathcal{W}}$.

References

[1] O.Arieli. *A four-valued approach for handling inconsistency in prioritized knowledge-bases.* Proc. FLAIRS'97, pages 92–96, 1997.

[2] O.Arieli, A.Avron. *Reasoning with logical bilattices.* Journal of Logic, Language, and Information 5(1), pages 25–63, 1996.

[3] O.Arieli, A.Avron. *Four-valued diagnoses for stratified knowledge-bases.* Proc. CSL'96. LNCS 1258, pages 1–17, Springer, 1997.

[4] O.Arieli, A.Avron. *The value of the four values.* Artificial Intelligence 102(1), pages 97–141, 1998.

[5] O.Arieli, A.Avron. *A model-theoretic approach to recover consistent data from inconsistent knowledge-bases.* To appear in the Journal of Automated Reasoning.

[6] N.D.Belnap. *A useful four-valued logic.* Modern Uses of Multiple-Valued Logic (G.Epstein and J.M.Dunn – eds.), pages 7–37. Reidel Publishing Company, 1977.

[7] N.D.Belnap. *How computer should think.* Contemporary Aspects of Philosophy (G.Ryle – ed.), pages 30–56. Oriel Press, 1977.

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

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

[10] S.Benferhat, D.Dubois, H.Prade. *Nonmonotonic reasoning, conditional objects and possibility theory.* Artificial Intelligence 92, pages 259–276, 1997.

[11] N.C.A.da-Costa. *On the theory of inconsistent formal systems.* Notre Damm Journal of Formal Logic 15, pages 497–510, 1974.

[12] D.Dubios, J.Lang, H.Prade. *Possibilistic logic.* Handbook of Logic in Artificial Intelligence and Logic Programming (D.Gabbay, C.Hogger, J.Robinson – eds.), Oxford Science Publications, pages 439–513, 1994.

[13] M.Fitting. *Bilattices and the semantics of logic programming.* Journal of Logic Programming 11(2), pages 91–116, 1991.

[14] M.Fitting. *Kleene's three-valued logics and their children.* Fundamenta Informaticae 20, pages 113–131, 1994.

[15] D.M.Gabbay. *Theoretical foundation for non-monotonic reasoning in expert systems.* Proc. of the NATO Advanced Study Inst. on Logic and Models of Concurent Systems (K.P.Apt – ed.) Springer-Verlag, pages 439–457, 1985.

[16] M.L.Ginsberg. *Multivalued logics: A uniform approach to reasoning in AI.* Computer Intelligence 4, pages 256–316, 1988.

[17] M.Goldszmidt, J.Pearl. *Qualitative probabilities for default reasoning, belief revision, and causal modeling.* Artificial Intelligence 84, pages 57–112, 1996.

[18] D.Lehmann, M.Magidor. *What does a conditional knowledge base entail?* Artificial Intelligence 55(1), pages 1–60, 1992.

[19] G.Wagner. *Vivid logic.* LNAI 764. Springer-Verlag, 1994.