Bilattices and Paraconsistency

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

Arnon Avron
Department of Computer Science
School of Mathematical Sciences
Tel-Aviv University
Ramat-Aviv 69978, Israel.
Email: aa@math.tau.ac.il

May 12, 1998

Abstract

Bilattices are algebraic structures that were introduced by Ginsberg, and further examined by Fitting, as a general framework for many applications in computer science. In this paper we consider their applicability for computerized reasoning in general, and for reasoning with inconsistent data in particular.

1 Background

A great deal of research has been devoted in the last twenty years for constructing plausible paraconsistent systems. One of the pioneering works towards this purpose was that of Belnap, who introduced his well-known four-valued logic [Be77a, Be77b]. The idea is that in addition to the classical values t, f, two additional truth-values are introduced for intuitively representing incomplete knowledge. One, denoted here by $\bot$, represents lack of knowledge. The other, $\top$, denotes "over"-knowledge (conflicts). These four elements form a structure called *FOUR* (see Figure 1). The basic idea is that this structure is "two-dimensional"; Each "dimension" corresponds to another partial ordering of the truth values. One order, $\leq_t$, is represented in the horizontal axis of Figure 1. It intuitively reflects differences in the "measure of *truth*" that each value represents. The corresponding lattice was originally denoted L4 by Belnap. The vertical axis of Figure 1 represents the other partial order, $\leq_k$, that might be understood as reflecting differences in the amount of *knowledge* or *information* that each truth value exhibits. Belnap denotes the corresponding lattice by A4.

Belnap's logic was generalized by Ginsberg [Gi88], who introduced the notion of *bilattices*, which are algebraic structures that contain *arbitrary* number of truth values simultaneously arranged in two partial orders. These orders are related by a negation operator ($\neg$) that is an involution w.r.t. $\leq_t$ and order preserving w.r.t $\leq_k$. This reflects the intuition that while one expects negation to invert the notion of truth, we know no more and no less about $\neg p$ than we know about p. Formally:

Definition 1.1 [Gi88] A *bilattice* is a structure $\mathcal{B} = (B, \leq_t, \leq_k, \neg)$ such that B is a nonempty set containing at least two elements; $(B, \leq_t)$, $(B, \leq_k)$ are complete lattices; and $\neg$ is a unary operation

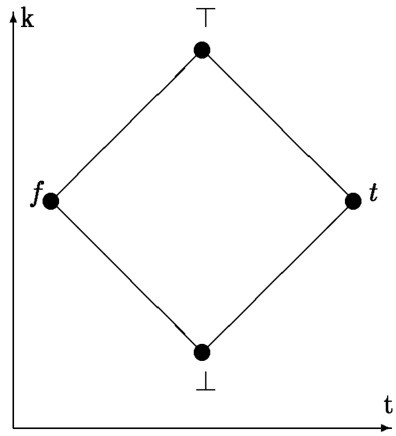

Figure 1: *FOUR*

on B that has the following properties: (1) if $a \leq_t b$, then $\neg a \geq_t \neg b$, (2) if $a \leq_k b$, then $\neg a \leq_k \neg b$, (3) $\neg\neg a = a$.

Following Fitting, we shall use $\wedge$ and $\vee$ for the lattice operations which correspond to $\leq_t$, and $\otimes$, $\oplus$ for those that correspond to $\leq_k$. Also, f and t denote, respectively, the least and the greatest element w.r.t. $\leq_t$, while $\bot$ and $\top$ – the least and the greatest element w.r.t. $\leq_k$. It is easy to see that $t, f, \top$, and $\bot$ are all distinct from each other.

Definition 1.2 A bilattice is called ***distributive*** [Gi88] if all the twelve possible distributive laws concerning $\wedge$, $\vee$, $\otimes$, and $\oplus$ hold. It is called *interlaced* [Fi90a] if each one of $\wedge$, $\vee$, $\otimes$, and $\oplus$ is monotonic with respect to both $\leq_t$ and $\leq_k$.

Proposition 1.3 [Fi90b] Every distributive bilattice is interlaced.

Proposition 1.4 [Fi91] If $\mathcal{B}$ is interlaced, then $t \oplus f = \top$, $t \otimes f = \bot$, $\top \vee \bot = t$, and $\top \wedge \bot = f$.

The next definition describes a general method for constructing distributive and interlaced bilattices:

Definition 1.5 [Gi88] Let $(L, \leq_L)$ be a complete lattice. The structure $L \odot L = (L \times L, \leq_t, \leq_k, \neg)$ is defined as follows:
$(y_1, y_2) \geq_t (x_1, x_2)$ iff $y_1 \geq_L x_1$ and $y_2 \leq_L x_2$,
$(y_1, y_2) \geq_k (x_1, x_2)$ iff $y_1 \geq_L x_1$ and $y_2 \geq_L x_2$,
$\neg(x_1, x_2) = (x_2, x_1)$.

Proposition 1.6 Let $L \odot L$ be the structure defined in 1.5.
a) [Gi88] If L is distributive then so is $L \odot L$.
b) [Fi90a] Every distributive bilattice is isomorphic to $L \odot L$ for some distributive lattice L.
c) [Fi90a] $L \odot L$ is always an interlaced bilattice.
d) [Av96] Every interlaced bilattice is isomorphic to $L \odot L$ for some bounded lattice L.

A truth value $(x,y) \in L \odot L$ may intuitively be understood so that x represents the amount of evidence *for* an assertion, while y represents the amount of evidence *against* it. It is easily verified that $\bot_{L\odot L} = (\inf(L), \inf(L))$; $\top_{L\odot L} = (\sup(L), \sup(L))$; $t_{L\odot L} = (\sup(L), \inf(L))$; and $f_{L\odot L} = (\inf(L), \sup(L))$.

The original motivation of Ginsberg for using bilattices was to provide a uniform approach for a diversity of applications in AI (see [Gi88]). Fitting has further investigated these structures [Fi90a, Fi94] and showed that they are useful for providing semantic to logic programs [Fi90a, Fi91, Fi93].

In [AA94, AA96] we presented a preliminary development of bilattice-based *logics* and corresponding proof systems. These logics turned out to have a proof theory with many desirable properties. In particular they may be used for non-monotonic reasoning and for making efficient inferences from inconsistent data. In the present paper we proceed with this logical approach. We consider bilattice-valued logics that are *preferential* in the sense of Shoham [Sh87, Sh88], i.e.: they are based on the idea that inferences should be taken not according to all models of a given theory, but only w.r.t. a subset of them, determined according to certain preference criteria. Roughly speaking, we use two guidelines for making preferences among models: (a) Prefer models that assume as much consistency as possible, and: (b) Prefer models that assume a minimal amount of knowledge (minimal commitment).

The existence of elements like $\top$ and $\bot$ as well as the idea of ordering data according to degrees of knowledge suggest that bilattices are particularly suitable for being a good semantical tool for constructing paraconsistent logics and for reasoning with uncertainty.

2 Logical bilattices

In order to define bilattice-based consequence relations, we first consider the subset of the *designated* truth values of a bilattice. This set is used for defining validity of formulae.

Definition 2.1 [AA94, AA96]

a) A *bifilter* of a bilattice $\mathcal{B}=(B, \leq_t, \leq_k)$ is a nonempty proper subset $\mathcal{F} \subset B$, such that:

(i) $a \wedge b \in \mathcal{F}$ iff $a \in \mathcal{F}$ and $b \in \mathcal{F}$ (ii) $a \otimes b \in \mathcal{F}$ iff $a \in \mathcal{F}$ and $b \in \mathcal{F}$

b) A bifilter $\mathcal{F}$ is called *prime*, if it also satisfies:

(i) $a \vee b \in \mathcal{F}$ iff $a \in \mathcal{F}$ or $b \in \mathcal{F}$ (ii) $a \oplus b \in \mathcal{F}$ iff $a \in \mathcal{F}$ or $b \in \mathcal{F}$

Note: Obviously, if $a \in \mathcal{F}$ and $b \geq_t a$ or $b \geq_k a$, then $b \in \mathcal{F}$. It immediately follows that $t, \top \in \mathcal{F}$ while $f, \bot \notin \mathcal{F}$.

Example 2.2 Ginsberg's *DEFAULT* (Figure 2, right) and Belnap's *FOUR* are bilattices that contain exactly one bifilter, $\{\top, t\}$, which is prime in both. *NINE* (Figure 2, left), on the other hand, contains two bifilters: $\{b \mid b \geq_k t\}$, as well as $\{b \mid b \geq_k dt\}$; both are prime.

The following propositions generalize the cases of *FOUR* and *NINE*:

Proposition 2.3 Let $\mathcal{B}=(B, \leq_t, \leq_k)$ be an interlaced bilattice.

a) A subset $\mathcal{F}$ of B is a (prime) bifilter iff it is a (prime) filter relative to $\leq_t$, and $\top \in \mathcal{F}$.

b) A subset $\mathcal{F}$ of B is a (prime) bifilter iff it is a (prime) filter relative to $\leq_k$, and $t \in \mathcal{F}$.

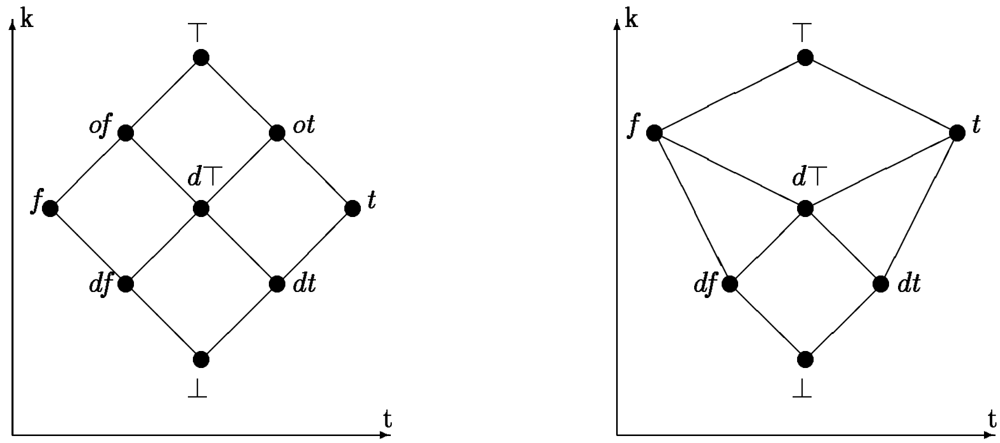

Figure 2: *NINE* and *DEFAULT*

Proof: Assume that $\mathcal{B}$ is interlaced.
a) The condition is obviously necessary. For the converse it suffices to show that: (i) if $a \in \mathcal{F}$ and $b \in \mathcal{F}$ then $a \otimes b \in \mathcal{F}$, (ii) if $a \in \mathcal{F}$ and $b \geq_k a$ then $b \in \mathcal{F}$, and (iii) if $\mathcal{F}$ is prime relative to $\leq_t$ then $a \oplus b \in \mathcal{F}$ iff either $a \in \mathcal{F}$ or $b \in \mathcal{F}$. Now, (i) and (iii) follow, respectively, from the facts that in interlaced bilattices $a \otimes b \geq_t a \wedge b$ and $a \vee b \geq_t a \oplus b$. For (ii) we note that $a \leq_k b$ is equivalent to $a \leq_k b \leq_k \top$. Since $\mathcal{B}$ is interlaced, it follows that $a \wedge (a \wedge \top) \leq_k b \wedge (a \wedge \top) \leq_k \top \wedge (a \wedge \top)$. Thus $a \wedge \top \leq_k b \wedge (a \wedge \top) \leq_k a \wedge \top$, and so $b \wedge (a \wedge \top) = a \wedge \top$. Hence $b \geq_t a \wedge \top$. Since $a \in \mathcal{F}$, $\top \in \mathcal{F}$ and $\mathcal{F}$ is a filter w.r.t. $\leq_t$, necessarily $b \in \mathcal{F}$ as well.
b) The proof is dual to that of part (a). □

Notation 2.4 $\mathcal{F}_k(a) = \{b \mid b \geq_k a\}$, $\mathcal{F}_t(a) = \{b \mid b \geq_t a\}$.

Proposition 2.5 Let $\mathcal{B} = (B, \leq_t, \leq_k)$ be an interlaced bilattice.
a) $\mathcal{F}_k(a)$ is a bifilter of $\mathcal{B}$ if $\bot \neq a \leq_k t$, iff $a >_t \bot$. Moreover, in this case $\mathcal{F}_k(a) = \mathcal{F}_t(a \wedge \top)$.
b) $\mathcal{F}_t(a)$ is a bifilter of $\mathcal{B}$ if $f \neq a \leq_t \top$, iff $a >_k f$. Moreover, in this case $\mathcal{F}_t(a) = \mathcal{F}_k(a \otimes t)$.

Proof:
a) If $a \neq \bot$ then the set $\{b \mid b \geq_k a\}$ is obviously a filter relative to $\leq_k$. By Proposition 2.3(b) it follows, therefore, that it is a bifilter if $\bot \neq a \leq_k t$. In other words, $a \neq \bot$ and $\bot \leq_k a \leq_k t$. Since $\mathcal{B}$ is interlaced this means that $a \neq \bot$ and $\bot \wedge \bot \leq_k a \wedge \bot \leq_k t \wedge \bot = \bot$, and so $a \neq \bot$ and $a \wedge \bot = \bot$. It follows that $a \neq \bot$ and $a \geq_t \bot$, and so $a >_t \bot$. For the other part of the proposition, recall that in the proof of Proposition 2.3(a) it is shown that if $b \geq_k a$ then $b \geq_t a \wedge \top$. Thus $\mathcal{F}_k(a) \subseteq \mathcal{F}_t(a \wedge \top)$. On the other hand, if $a >_t \bot$ then $a \vee \top \geq_t \bot \vee \top = t$ (Proposition 1.4), and so $a \vee \top = t$. It follows that if $b \geq_t a \wedge \top$ then $a \wedge \top \leq_t b \leq_t a \vee \top$, and so $a \otimes (a \wedge \top) \leq_t a \otimes b \leq_t a \otimes (a \vee \top)$. But $a \leq_k \top$ implies that $a = a \wedge a \leq_k a \wedge \top$, and so $a \otimes (a \wedge \top) = a$. Similarly $a \otimes (a \vee \top) = a$. Hence $a \leq_t a \otimes b \leq_t a$, and so $a \otimes b = a$, which means that $a \leq_k b$. Thus, $\mathcal{F}_t(a \wedge \top) \subseteq \mathcal{F}_k(a)$ and so $\mathcal{F}_t(a \wedge \top) = \mathcal{F}_k(a)$.
b) The proof is dual to that of part (a). □

Proposition 2.6 Let $\mathcal{B} = (B, \leq_t, \leq_k)$ be an interlaced bilattice. If $\mathcal{F}$ is a bifilter in $\mathcal{B}$, then $\inf_k \mathcal{F} \in \mathcal{F}$ iff $\inf_t \mathcal{F} \in \mathcal{F}$. Moreover, in such a case $\inf_t \mathcal{F} = \top \wedge \inf_k \mathcal{F}$ and $\inf_k \mathcal{F} = t \otimes \inf_t \mathcal{F}$.

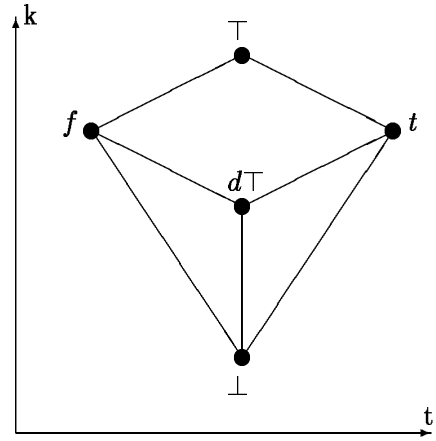

Figure 3: *FIVE*

Proof: Follows from Proposition 2.5. □

Definition 2.7 [AA94, AA96] A *logical bilattice* is a pair $(\mathcal{B},\mathcal{F})$, in which $\mathcal{B}$ is a bilattice and $\mathcal{F}$ is a prime bifilter of $\mathcal{B}$.

Logical bilattices will be our primary semantical tool for defining paraconsistent logics. Not every bilattice can be turned into a logical one. *FIVE* (Figure 3), for instance, has only one filter $\mathcal{F}=\{\top,t\}$, which is not prime: $d\top \vee \bot = t \in \mathcal{F}$, while $d\top \notin \mathcal{F}$, $\bot \notin \mathcal{F}$. However, as Propositions 2.8 and 2.10 below show, logical bilattices are very common, and easily constructed:

Proposition 2.8 Let $L\odot L$ be a bilattice as described in Definition 1.5.
a) $\mathcal{F}$ is a bifilter in $L\odot L$ iff $\mathcal{F}=\mathcal{F}_L\times L$, where $\mathcal{F}_L$ is a filter in L.
b) $\mathcal{F}$ is a prime bifilter in $L\odot L$ iff $\mathcal{F}=\mathcal{F}_L\times L$, where $\mathcal{F}_L$ is a prime filter in L.

Proof:
a) ($\Leftarrow$) Let $\mathcal{F}_L$ be a filter in L and let $\mathcal{F}=\mathcal{F}_L\times L$. Since $\inf(L)\notin\mathcal{F}_L$ and $\sup(L)\in\mathcal{F}_L$, so for every $x\in L$ $(\inf(L),x)\notin\mathcal{F}$ and $(\sup(L),x)\in\mathcal{F}$. Thus $\mathcal{F}$ is a nonempty proper subset of $L\odot L$. Now, $(x_1,x_2)\wedge(y_1,y_2)\in\mathcal{F}$, iff $(x_1\wedge_L y_1, x_2\vee_L y_2)\in\mathcal{F}$, iff $x_1\wedge_L y_1\in\mathcal{F}_L$, iff $x_1\in\mathcal{F}_L$ and $y_1\in\mathcal{F}_L$, iff $(x_1,x_2)\in\mathcal{F}$ and $(y_1,y_2)\in\mathcal{F}$. The proof in the case of $\otimes$ is similar. Therefore $\mathcal{F}$ is a bifilter in $L\odot L$.
($\Rightarrow$) Let $\mathcal{F}$ be a bifilter in $L\odot L$. Denote: $\mathcal{F}_L=\{x \mid \exists y\ (x,y)\in\mathcal{F}\}$. We shall show that $\mathcal{F}=\mathcal{F}_L\times L$. Obviously, $\mathcal{F}\subseteq\mathcal{F}_L\times L$. For the converse, let $(x,l)\in\mathcal{F}_L\times L$. Then there is a $y\in L$ s.t. $(x,y)\in\mathcal{F}$. Now, $(x,l\vee_L y)\geq_k(x,y)\in\mathcal{F}$, and so $(x,l\vee_L y)\in\mathcal{F}$. On the other hand, $(x,l)\geq_t(x,l\vee_L y)\in\mathcal{F}$, and so $(x,l)\in\mathcal{F}$. It follows, therefore, that $\mathcal{F}_L\times L\subseteq\mathcal{F}$. Hence $\mathcal{F}=\mathcal{F}_L\times L$.
b) Suppose first that $\mathcal{F}_L$ is a prime filter in L. Then: $(x_1,x_2)\vee(y_1,y_2)\in\mathcal{F}$, iff $(x_1\vee_L y_1, x_2\wedge_L y_2)\in\mathcal{F}$, iff $x_1\vee_L y_1\in\mathcal{F}_L$, iff $x_1\in\mathcal{F}_L$ or $y_1\in\mathcal{F}_L$, iff $(x_1,x_2)\in\mathcal{F}$ or $(y_1,y_2)\in\mathcal{F}$. The proof in the case of $\oplus$ is similar. For the converse, assume that $\mathcal{F}$ is a prime bifilter in $L\odot L$. By part (a), $\mathcal{F}=\mathcal{F}_L\times L$, where $\mathcal{F}_L$ is a filter in L. We show that $\mathcal{F}_L$ is prime: Assume that $x\vee_L y\in\mathcal{F}_L$ and let z be some element in L. Then $(x\vee_L y,z)\in\mathcal{F}\Rightarrow(x,z)\vee(y,z)\in\mathcal{F}\Rightarrow(x,z)\in\mathcal{F}$ or $(y,z)\in\mathcal{F}\Rightarrow x\in\mathcal{F}_L$ or $y\in\mathcal{F}_L$. □

Corollary 2.9
a) Let $x_0 \in L$, $x_0 \neq \inf(L)$. Denote: $\mathcal{F}(x_0) = \{(y_1, y_2) \mid y_1 \geq_L x_0,\ y_2 \in L\}$, and $\mathcal{F}_L(x_0) = \{y \in L \mid y \geq_L x_0\}$. Then $(L \odot L, \mathcal{F}(x_0))$ is a logical bilattice iff $\mathcal{F}_L(x_0)$ is prime.
b) $(L \odot L, \mathcal{F}(\sup(L)))$ is a logical bilattice iff $\sup(L)$ is join irreducible (i.e.: iff $x \vee_L y = \sup(L)$ implies that $x = \sup(L)$ or $y = \sup(L)$).
c) If the condition of case (b) is met, then $\mathcal{F}(\sup(L))$ is minimal among the (prime) bifilters of $L \odot L$.

Proof: Part (a) immediately follows from Propositions 2.5(a) and 2.8(b), since $\mathcal{F}(x_0) = \mathcal{F}_k(z)$ where $z = (x_0, \inf(L))$. Part (b) follows from (a), since $\mathcal{F}_L(\sup(L)) = \{\sup(L)\}$ is a prime filter in L iff $\sup(L)$ is join irreducible. For part (c) note that $\mathcal{F}(\sup(L)) = \mathcal{F}_k(t_{L \odot L})$. The claim follows therefore from (b) and the fact that every bilattice contains the set $\{b \in \mathcal{B} \mid b \geq_k t_{\mathcal{B}}\}$. □

Proposition 2.10 Every distributive bilattice can be turned into a logical bilattice.

First proof: Let $\mathcal{B}$ be a distributive bilattice. Consider a $\leq_t$-filter $\mathcal{F}'$ in B s.t. $\top \in \mathcal{F}'$ (clearly there is such a filter, e.g.: $\mathcal{F}_t(\top)$). By a famous theorem of lattice theory (see [Bi67]) $\mathcal{F}'$ can be extended to a prime $\leq_t$-filter $\mathcal{F}$. By Proposition 2.3a, $\mathcal{F}$ is a prime bifilter. □

Second proof: By Fitting's theorem mentioned in Proposition 1.6(b), every distributive bilattice is isomorphic to $L \odot L$, where L is a distributive lattice. Let $\mathcal{F}_L$ be any prime filter of L (again, such a filter exists by a theorem of lattice theory). Then $\mathcal{F}_L \times L$ is a prime bifilter by Proposition 2.8. □

Note: Not every logical bilattice needs to be distributive or even interlaced. $(DEFAULT, \{\top, t\})$ is, for example, a logical bilattice although $DEFAULT$ is *not* interlaced.

3 Bilattice-valued logics for paraconsistent reasoning

Given a logical bilattice $(\mathcal{B}, \mathcal{F})$, the standard notions of valuations, models, etc. are defined in the usual way. Now, let ψ be a formula in the basic language of bilattices $(\{\neg, \vee, \wedge, \otimes, \oplus\})$, and suppose that ν is a valuation that assigns $\bot$ to every atomic formula. Then $\nu(\psi) = \bot$ as well, and so there are no tautologies in the basic language of bilattices. Thus, e.g., excluded middle is not a valid rule, and this implies that the definition of the material implication $p \mapsto q$ as $\neg p \vee q$ is not adequate for representing entailments. We use therefore instead another connective, denoted $\supset$ ([AA96]), which does function as an implication (see Proposition 3.2 below). It is defined as follows: $a \supset b = t$ if a is not designated, otherwise: $a \supset b = b$.

In the rest of this section we briefly consider several families of plausible logics, the semantics of which is based on logical bilattices. As we shall see, all the consequence relations involved are paraconsistent. Some of these logics may be viewed as generalizations of other well known paraconsistent logics, such as D'ottaviano J_3 [Do85] (see also [Ro89, Av91] and chapter IX of [Ep90]), Belnap's four-valued logic [Be77a, Be77b], and Priest's LPm [Pr89, Pr91].

3.1 The basic consequence relation

We start with the simplest consequence relation which naturally corresponds to logical bilattices:

Definition 3.1 Let $(\mathcal{B},\mathcal{F})$ be a logical bilattice and suppose that Γ and Δ are two sets of formulae. $\Gamma \models^{\mathcal{B},\mathcal{F}} \Delta$ if every model of Γ in $(\mathcal{B},\mathcal{F})$ is a model of some formula of Δ.

The main properties of $\models^{\mathcal{B},\mathcal{F}}$ are summarized in the following proposition:

Proposition 3.2 [AA96] $\models^{\mathcal{B},\mathcal{F}}$ is monotonic, compact, and paraconsistent. Is has a cut free, sound and complete Gentzen-type proof system (*GBL*; see Figure 4), and the deduction theorem is valid for it w.r.t. $\supset$.

$\models^{\mathcal{B},\mathcal{F}}$ is therefore a consequence relation in the standard sense of Tarski and Scott. Note that the $\{\wedge,\vee,\neg\}$-fragment of $\models^{\mathcal{B},\mathcal{F}}$ in case that $\mathcal{B}=FOUR$ and $\mathcal{F}=\{t,\top\}$ is identical to the set of "first degree entailments" in relevance logic (see [AB75, Du86]).

As the following proposition shows, $\models^{\mathcal{B},\mathcal{F}}$ has a strong connection to Belnap's four-valued logic. In what follows we shall denote $\langle FOUR\rangle = (FOUR, \{t,\top\})$, and write "4" whenever $\langle FOUR\rangle$ should appear as a superscript.

Theorem 3.3 [AA96, AA97a] $\Gamma \models^{\mathcal{B},\mathcal{F}} \Delta$ iff $\Gamma \models^{4} \Delta$.

Despite the nice properties of $\models^{\mathcal{B},\mathcal{F}}$, it appears that it has several drawbacks. One of which is that $\models^{\mathcal{B},\mathcal{F}}$ is strictly weaker than classical logic, even for consistent theories (e.g., $\not\models^{\mathcal{B},\mathcal{F}} p \vee \neg p$). Also, it completely invalidates some intuitively justified inference rules, like the Disjunctive Syllogism: From $\neg p$ and $p \vee q$ one can *never* infer q by using $\models^{\mathcal{B},\mathcal{F}}$.

3.2 The logics $\models_k^{\mathcal{B},\mathcal{F}}$

A natural approach for reducing the set of models which are used for drawing conclusions is to consider only the k-minimal ones. The idea behind this approach is that one should not assume anything that is not really known. Keeping the amount of knowledge as minimal as possible may be taken as a kind of consistency preserving method: As long as one keeps the redundant information as minimal as possible, the tendency of getting into conflicts decreases.

Definition 3.4 Let ν_1, ν_2 be two four-valued valuations, and Γ – a set of formulae.
a) ν_1 is ***k*-*smaller*** than ν_2 ($\nu_1 \leq_k \nu_2$) if for every atomic p, $\nu_1(p) \leq_k \nu_2(p)$.
b) ν is a ***k*-*minimal model*** of Γ if there is no model of Γ which is k-smaller than ν.

Definition 3.5 $\Gamma \models_k^{\mathcal{B},\mathcal{F}} \Delta$ iff every k-minimal model of Γ in $(\mathcal{B},\mathcal{F})$ is a model of some $\delta \in \Delta$.

Note: Obviously, if $\Gamma \models^{\mathcal{B},\mathcal{F}} \Delta$ then $\Gamma \models_k^{\mathcal{B},\mathcal{F}} \Delta$.

Lemma 3.6 Let $\mathcal{B}$ be a finite bilattice. For every model M of Γ there exists a k-minimal model N of Γ s.t. $N \leq_k M$.[1]

[1]Property of this kind is called in another context *smoothness* ([KLM90]), or *stopperedness* ([Ma94]).

Axioms: $\Gamma, \psi \Rightarrow \Delta, \psi$

Rules: Exchange, Contraction, and the following logical rules:

$$[\neg\neg\Rightarrow]\ \frac{\Gamma, \psi \Rightarrow \Delta}{\Gamma, \neg\neg\psi \Rightarrow \Delta} \qquad \frac{\Gamma \Rightarrow \Delta, \psi}{\Gamma \Rightarrow \Delta, \neg\neg\psi}\ [\Rightarrow\neg\neg]$$

$$[\wedge\Rightarrow]\ \frac{\Gamma, \psi, \phi \Rightarrow \Delta}{\Gamma, \psi \wedge \phi \Rightarrow \Delta} \qquad \frac{\Gamma \Rightarrow \Delta, \psi \quad \Gamma \Rightarrow \Delta, \phi}{\Gamma \Rightarrow \Delta, \psi \wedge \phi}\ [\Rightarrow\wedge]$$

$$[\neg\wedge\Rightarrow]\ \frac{\Gamma, \neg\psi \Rightarrow \Delta \quad \Gamma, \neg\phi \Rightarrow \Delta}{\Gamma, \neg(\psi \wedge \phi) \Rightarrow \Delta} \qquad \frac{\Gamma \Rightarrow \Delta, \neg\psi, \neg\phi}{\Gamma \Rightarrow \Delta, \neg(\psi \wedge \phi)}\ [\Rightarrow\neg\wedge]$$

$$[\vee\Rightarrow]\ \frac{\Gamma, \psi \Rightarrow \Delta \quad \Gamma, \phi \Rightarrow \Delta}{\Gamma, \psi \vee \phi \Rightarrow \Delta} \qquad \frac{\Gamma \Rightarrow \Delta, \psi, \phi}{\Gamma \Rightarrow \Delta, \psi \vee \phi}\ [\Rightarrow\vee]$$

$$[\neg\vee\Rightarrow]\ \frac{\Gamma, \neg\psi, \neg\phi \Rightarrow \Delta}{\Gamma, \neg(\psi \vee \phi) \Rightarrow \Delta} \qquad \frac{\Gamma \Rightarrow \Delta, \neg\psi \quad \Gamma \Rightarrow \Delta, \neg\phi}{\Gamma \Rightarrow \Delta, \neg(\psi \vee \phi)}\ [\Rightarrow\neg\vee]$$

$$[\otimes\Rightarrow]\ \frac{\Gamma, \psi, \phi \Rightarrow \Delta}{\Gamma, \psi \otimes \phi \Rightarrow \Delta} \qquad \frac{\Gamma \Rightarrow \Delta, \psi \quad \Gamma \Rightarrow \Delta, \phi}{\Gamma \Rightarrow \Delta, \psi \otimes \phi}\ [\Rightarrow\otimes]$$

$$[\neg\otimes\Rightarrow]\ \frac{\Gamma, \neg\psi, \neg\phi \Rightarrow \Delta}{\Gamma, \neg(\psi \otimes \phi) \Rightarrow \Delta} \qquad \frac{\Gamma \Rightarrow \Delta, \neg\psi \quad \Gamma \Rightarrow \Delta, \neg\phi}{\Gamma \Rightarrow \Delta, \neg(\psi \otimes \phi)}\ [\Rightarrow\neg\otimes]$$

$$[\oplus\Rightarrow]\ \frac{\Gamma, \psi \Rightarrow \Delta \quad \Gamma, \phi \Rightarrow \Delta}{\Gamma, \psi \oplus \phi \Rightarrow \Delta} \qquad \frac{\Gamma \Rightarrow \Delta, \psi, \phi}{\Gamma \Rightarrow \Delta, \psi \oplus \phi}\ [\Rightarrow\oplus]$$

$$[\neg\oplus\Rightarrow]\ \frac{\Gamma, \neg\psi \Rightarrow \Delta \quad \Gamma, \neg\phi \Rightarrow \Delta}{\Gamma, \neg(\psi \oplus \phi) \Rightarrow \Delta} \qquad \frac{\Gamma \Rightarrow \Delta, \neg\psi, \neg\phi}{\Gamma \Rightarrow \Delta, \neg(\psi \oplus \phi)}\ [\Rightarrow\neg\oplus]$$

$$[\supset\Rightarrow]\ \frac{\Gamma \Rightarrow \psi, \Delta \quad \Gamma, \phi \Rightarrow \Delta}{\Gamma, \psi \supset \phi \Rightarrow \Delta} \qquad \frac{\Gamma, \psi \Rightarrow \phi, \Delta}{\Gamma \Rightarrow \psi \supset \phi, \Delta}\ [\Rightarrow\supset]$$

$$[\neg\supset\Rightarrow]\ \frac{\Gamma, \psi, \neg\phi \Rightarrow \Delta}{\Gamma, \neg(\psi \supset \phi) \Rightarrow \Delta} \qquad \frac{\Gamma \Rightarrow \psi, \Delta \quad \Gamma \Rightarrow \neg\phi, \Delta}{\Gamma \Rightarrow \neg(\psi \supset \phi), \Delta}\ [\Rightarrow\neg\supset]$$

$$[\neg t\Rightarrow]\ \Gamma, \neg t \Rightarrow \Delta \qquad \Gamma \Rightarrow \Delta, t\ [\Rightarrow t]$$

$$[f\Rightarrow]\ \Gamma, f \Rightarrow \Delta \qquad \Gamma \Rightarrow \Delta, \neg f\ [\Rightarrow \neg f]$$

$$[\bot\Rightarrow]\ \Gamma, \bot \Rightarrow \Delta \qquad \Gamma \Rightarrow \Delta, \top\ [\Rightarrow \top]$$

$$[\neg\bot\Rightarrow]\ \Gamma, \neg\bot \Rightarrow \Delta \qquad \Gamma \Rightarrow \Delta, \neg\top\ [\Rightarrow \neg\top]$$

Figure 4: The system *GBL*

Proof: Suppose that M is some model of Γ, and let $S_M = \{M_i \mid M_i \text{ is a model of } \Gamma,\ M_i \leq_k M\}$. Let $C \subseteq S_M$ be a descending chain w.r.t. $\leq_k$. We shall show that C is bounded in S_M, so by Zorn's lemma S_M has a minimal element, which is the required k-minimal model. Let N be the the following valuation: $N(p) = \min_{\leq_k}\{M_i(p) \mid M_i \in C\}$. N is defined since C is a chain, and $\mathcal{B}$ is finite. Obviously N bounds C. It remains to show that $N \in S_M$. Assume that $\psi \in \Gamma$ and let $\mathcal{A}(\psi) = \{p_1, \ldots, p_n\}$ be the set of the atomic formulae in ψ. Then: $N(p_1) = M_{i_1}(p_1), \ldots, N(p_n) = M_{i_n}(p_n)$. Since C is a chain we may assume, without a loss of generality, that $M_{i_1} \geq_k \ldots \geq_k M_{i_n}$, and so N is the same as M_{i_n} on every atom in $\mathcal{A}(\psi)$. Since M_{i_n} is a model of ψ, so is N. This is true for every $\psi \in \Gamma$ and so $N \in S_M$ as required. $\Box$

As the following proposition shows, it is sometimes sufficient to consider only the k-minimal models of a given theory for making inferences with $\models^{\mathcal{B},\mathcal{F}}$:

Proposition 3.7 Let $\mathcal{B}$ be a finite interlaced bilattice, and $\mathcal{F}$ – a prime bifilter in $\mathcal{B}$. If the formulae of Δ are in the language without $\supset$, then $\Gamma \models^{\mathcal{B},\mathcal{F}} \Delta$ iff $\Gamma \models_k^{\mathcal{B},\mathcal{F}} \Delta$.

Proof: The "only if" direction is trivial. For the other direction, suppose that $\Gamma \models_k^{\mathcal{B},\mathcal{F}} \Delta$, and let M be some model of Γ. By Lemma 3.6 there is a k-minimal model N of Γ s.t. $M \geq_k N$. Thus there is a $\delta \in \Delta$ s.t. $N(\delta) \in \mathcal{F}$. Now, since $\mathcal{B}$ is interlaced, all the operators that correspond to the connectives of Δ are monotone w.r.t. $\leq_k$, and so $M(\delta) \geq_k N(\delta)$. But $\mathcal{F}$ is upwards-closed w.r.t. $\leq_k$, therefore $M(\delta) \in \mathcal{F}$ as well. $\Box$

Corollary 3.8 Let $\mathcal{B}$ be a finite interlaced bilattice. Then in the language without $\supset$, the logics $\models^{\mathcal{B},\mathcal{F}}$ and $\models_k^{\mathcal{B},\mathcal{F}}$ are identical.

Proposition 3.7 shows that in many cases we can limit ourselves to k-minimal models without any loss of generality. This property allows a considerable reduction in the number of models that should be checked.

From Propositions 3.2 and 3.7 it follows that $\models_k^{\mathcal{B},\mathcal{F}}$ is paraconsistent, and is also monotonic w.r.t. conclusions without $\supset$. The last property is no longer true when $\supset$ is allowed in the r.h.s. of $\models_k^{\mathcal{B},\mathcal{F}}$:

Proposition 3.9 $\models_k^{\mathcal{B},\mathcal{F}}$ is in general nonmonotonic.

Proof: Let $(\mathcal{B},\mathcal{F})$ be any logical bilattice in which $b_t = \inf_k\{b \mid b \in \mathcal{F}\} \in \mathcal{F}$.[2,3] Denote: $b_\top = \inf_k\{b \mid b, \neg b \in \mathcal{F}\}$. It is easy to verify that $b_\top, \neg b_\top \in \mathcal{F}$. Now, $q \models_k^{\mathcal{B},\mathcal{F}} \neg q \supset p$, since $M(p) = \bot$, $M(q) = b_t$ is the only k-minimal model of $\{q\}$ in $(\mathcal{B},\mathcal{F})$. On the other hand, $q, \neg q \not\models_k^{\mathcal{B},\mathcal{F}} \neg q \supset p$, since $N(p) = \bot$, $N(q) = b_\top$ is a counter k-minimal model of $\{q, \neg q\}$. $\Box$

Using the example of the last proof, one can easily see that $q \models_k^{\mathcal{B},\mathcal{F}} \neg q \supset p$ and also $\neg q, \neg q \supset p \models_k^{\mathcal{B},\mathcal{F}} p$, but $\neg q, q \not\models_k^{\mathcal{B},\mathcal{F}} p$. It follows that $\models_k^{\mathcal{B},\mathcal{F}}$ is not a consequence relation in the usual sense, since it is not closed under (multiplicative) cut. This is not surprising, since $\models_k^{\mathcal{B},\mathcal{F}}$ is not monotonic,

[2] This is clearly the case whenever B is finite.

[3] See also Proposition 2.6.

and it is usual to require a nonmonotonic relation to be closed only under *cautious* cut and *cautious* monotonicity (see [Ga85, KLM90, Le92, Ma94]):

Proposition 3.10
a) $\models_k^{\mathcal{B},\mathcal{F}}$ preserves cautious cut: If $\Gamma, \psi \models_k^{\mathcal{B},\mathcal{F}} \Delta$ and $\Gamma \models_k^{\mathcal{B},\mathcal{F}} \psi, \Delta$, then $\Gamma \models_k^{\mathcal{B},\mathcal{F}} \Delta$.
b) $\models_k^{\mathcal{B},\mathcal{F}}$ preserves cautious monotonicity: If $\Gamma \models_k^{\mathcal{B},\mathcal{F}} \psi$ and $\Gamma \models_k^{\mathcal{B},\mathcal{F}} \Delta$, then $\Gamma, \psi \models_k^{\mathcal{B},\mathcal{F}} \Delta$.

Proof:
a) Suppose that M is a k-minimal model of Γ, but $M(\delta) \notin \mathcal{F}$ for every $\delta \in \Delta$. Since $\Gamma \models_k^{\mathcal{B},\mathcal{F}} \psi, \Delta$, then $M(\psi) \in \mathcal{F}$, and so M is a model of $\{\Gamma, \psi\}$. Moreover, M must be a k-minimal model of $\{\Gamma, \psi\}$, since any other model of this set which is strictly smaller than M w.r.t. $\leq_k$ must be a model of Γ, which is k-smaller than M. Now, $\Gamma, \psi \models_k^{\mathcal{B},\mathcal{F}} \Delta$, thus $M(\delta) \in \mathcal{F}$ for some $\delta \in \Delta$ – a contradiction.
b) Assume that $\Gamma \models_k^{\mathcal{B},\mathcal{F}} \psi$, and $\Gamma \models_k^{\mathcal{B},\mathcal{F}} \Delta$. Let M be a k-minimal model of $\Gamma \cup \{\psi\}$. In particular, M is a model of Γ. Moreover, it must be an k-minimal model of Γ as well, since otherwise there would be a model N of Γ that is strictly k-smaller than M. Since $\Gamma \models_k^{\mathcal{B},\mathcal{F}} \psi$, this N would have been an k-minimal model $\Gamma \cup \{\psi\}$ and therefore $N <_k^{\mathcal{B},\mathcal{F}} M$ w.r.t. $\Gamma \cup \{\psi\}$ — a contradiction. Therefore, M is a k-minimal model of Γ. Now, since $\Gamma \models_k^{\mathcal{B},\mathcal{F}} \Delta$, M is a model of some $\delta \in \Delta$. Hence $\Gamma, \psi \models_k^{\mathcal{B},\mathcal{F}} \Delta$.
□

3.3 The logics $\models_{\mathcal{I}}^{\mathcal{B},\mathcal{F}}$

The motivation behind the last family of bilattice-based consequence relations that we consider here is perhaps the closest in spirit to the original idea of paraconsistent reasoning: We allow a nontrivial reasoning in the presence of inconsistency, while still trying to minimize the amount of contradictions. This approach reflects the intuition that while one has to deal with conflicts in a nontrivial way, contradictory data corresponds to inadequate information about the real world, and therefore should be minimized.

Definition 3.11 [AA94, AA96] Let $(\mathcal{B}, \mathcal{F})$ be a logical bilattice. A subset $\mathcal{I}$ of B is called an *inconsistency set* of $\mathcal{B}$ if it has the following properties:
a) $b \in \mathcal{I}$ iff $\neg b \in \mathcal{I}$.
b) $b \in \mathcal{F} \cap \mathcal{I}$ iff $b \in \mathcal{F}$ and $\neg b \in \mathcal{F}$.

Note: It is easy to see that if $\mathcal{I}$ is an inconsistency set then $t, f \notin \mathcal{I}$ and $\top \in \mathcal{I}$.

Example 3.12 In $\langle FOUR \rangle$ there are two inconsistency sets: $\mathcal{I}_1 = \{\top\}$ and $\mathcal{I}_2 = \{\top, \bot\}$. The use of $\mathcal{I}_1$ means preference of consistent values, while the use of $\mathcal{I}_2$ means preference of classical values.

Notation 3.13 $I(\nu, \mathcal{I}) = \{p \mid p$ is atomic and $\nu(p) \in \mathcal{I}\}$.

Intuitively, $\mathcal{I}$ is a set of inconsistent values of $(\mathcal{B}, \mathcal{F})$, and $I(\nu, \mathcal{I})$ corresponds to the inconsistent assignments of ν w.r.t. $\mathcal{I}$.

Definition 3.14 Let $(\mathcal{B}, \mathcal{F})$ be a logical bilattice and $\mathcal{I}$ – an inconsistency set of B.
a) ν_1 is *more consistent* than ν_2 w.r.t. $\mathcal{I}$ ($\nu_1 >_{\mathcal{I}}^{\mathcal{B},\mathcal{F}} \nu_2$) if $I(\nu_1, \mathcal{I}) \subset I(\nu_2, \mathcal{I})$.
b) ν is a *most consistent* model of Γ w.r.t. $\mathcal{I}$ ($\mathcal{I}$-mcm, for short), if there is no model of Γ which is more consistent than ν.

Definition 3.15 [AA94, AA96] $\Gamma \models_{\mathcal{I}}^{\mathcal{B},\mathcal{F}} \Delta$ if every $\mathcal{I}$-mcm of Γ in $(\mathcal{B},\mathcal{F})$ is a model of some formula of Δ.

Example 3.16 (Tweety Dilemma) Let $(\mathcal{B},\mathcal{F}) = \langle FOUR \rangle$. Using the notations of example 3.12, let $\mathcal{I}_1 = \{\top\}$ and $\mathcal{I}_2 = \{\top, \bot\}$. Consider the following set of assertions, Γ:

$bird(Tweety) \mapsto fly(Tweety)$[4]

$penguin(Tweety) \supset bird(Tweety)$

$penguin(Tweety) \supset \neg fly(Tweety)$

$bird(Tweety)$

Unlike the other formulae of Γ, the first assertion is an instance of a rule that *has exceptions*. Thus it is formulated with a weaker "implication" connective.

The 18 models of Γ in $\langle FOUR \rangle$ are given in Figure 5. Two of these models, M17 and M18, are

Model No.	$bird(Tweety)$	$fly(Tweety)$	$penguin(Tweety)$
M1 – M8	$\top$	$\top, f$	$\top, t, f, \bot$
M9 – M12	$\top$	$t, \bot$	$f, \bot$
M13 – M16	t	$\top$	$\top, t, f, \bot$
M17 – M18	t	t	$f, \bot$

Figure 5: The models of Γ (Example 3.16)

the $\mathcal{I}_1$-mcms of Γ. M17 – the only classical model of Γ – is also the only $\mathcal{I}_2$-mcm of Γ. Thus, when using $\models^4_{\mathcal{I}_1}$ one can infer that $bird(Tweety)$ (but $\neg bird(Tweety)$ is not true), and $fly(Tweety)$ (while $\neg fly(Tweety)$ is not true). Also, nothing is yet known about Tweety being a penguin. According to $\models^4_{\mathcal{I}_2}$ one can infer $bird(Tweety)$, $fly(Tweety)$, and $\neg penguin(Tweety)$. The inverse assertions are not true, as expected. Note that $fly(Tweety)$ is *not* a consequence of $\models^4_k$ (and so it is not a consequence of $\models^4$ as well. M12 is a counter-model for both cases). One might view this fact as an evidance that $\models^4_k$ is "over-cautious".

Suppose now that a new data arrives: $penguin(Tweety)$. The models of the modified knowledge-base, $\Gamma' = \Gamma \cup \{penguin(Tweety)\}$, are listed in Figure 6. This time the $\mathcal{I}_1$-mcms and the $\mathcal{I}_2$-mcms of Γ' coincide. They are denoted by M4 and M6. It follows that according to the new

Model No.	$bird(Tweety)$	$fly(Tweety)$	$penguin(Tweety)$
M1 – M2	$\top$	$\top$	$\top, t$
M3 – M4	$\top$	f	$\top, t$
M5 – M6	t	$\top$	$\top, t$

Figure 6: The models of Γ' (Example 3.16)

information one should change his belief and infer new conclusions: $bird(Tweety)$, $penguin(Tweety)$, and $\neg fly(Tweety)$. Although Γ' is classically *inconsistent*, the complements of these assertions *cannot* be inferred by $\models^4_{\mathcal{I}_j}$ $(j = 1, 2)$, as indeed one expects.

[4]Recall that $\mapsto$ denotes the material implication.

Note: There is a slight (but significant) difference between the definition of $\models^{\mathcal{B},\mathcal{F}}_{\mathcal{I}}$ and the definition of the paraconsistent relation $\models^{\mathcal{B},\mathcal{F}}_{con(\mathcal{I})}$ (abbreviation: $\models_{con}$), considered in [AA94, AA96]. Here we consider the inconsistent assignments of a given valuation w.r.t. *all* the atomic formulae. In [AA94, AA96], on the other hand, only the assignments on the atomic formulae that appear in the language of the premises are relevant for making preferences among valuations. In other words, if $\mathcal{A}(\Gamma)$ is the set of atomic formulae in the language of the premises Γ, then the relevant set of assignments according to [AA94, AA96] is $I(\nu,\Gamma,\mathcal{I}) = \{p \in \mathcal{A}(\Gamma) \mid \nu(p) \in \mathcal{I}\}$ (cf. Definition 3.13). Obviously, $\models^{\mathcal{B},\mathcal{F}}_{\mathcal{I}}$ is not the same logic as $\models^{\mathcal{B},\mathcal{F}}_{con(\mathcal{I})}$. For example, $p \models^{4}_{\{\top,\bot\}} q \vee \neg q$, while $p \not\models^{4}_{con(\{\top,\bot\})} q \vee \neg q$ (a counter-model assigns t to p and $\bot$ to q).

Our new definition has several advantages over the previous one. One of which is the fact that by Proposition 3.19(a) below, cautious cut is always sound for $\models^{\mathcal{B},\mathcal{F}}_{\mathcal{I}}$. In the case of $\models_{con}$, however, cautious cut is valid only in the $\{\vee,\wedge,\oplus,\otimes,\neg\}$-fragment of the language, and provided that there is a value $b \in B$ s.t. $b, \neg b \notin \mathcal{F}$ and $b \notin \mathcal{I}$. Therefore, e.g., cautious cut fails in the case of $\models^{4}_{con(\{\top,\bot\})}$. Indeed, $q \models^{4}_{con(\{\top,\bot\})} q \vee p$ and $q, q \vee p \models^{4}_{con(\{\top,\bot\})} p \vee \neg p$, but $q \not\models^{4}_{con(\{\top,\bot\})} p \vee \neg p$. Cautious cut also fails in the case of $\models^{\mathcal{B},\mathcal{F}}_{con(\mathcal{I})}$ whenever $\supset$ appears in the language. For a counter-example note that $q \models^{\mathcal{B},\mathcal{F}}_{con(\mathcal{I})} q \vee p$, and $q,\ q \vee p \models^{\mathcal{B},\mathcal{F}}_{con(\mathcal{I})} (p \supset \neg q) \vee (\neg p \supset \neg q)$, but $q \not\models^{\mathcal{B},\mathcal{F}}_{con(\mathcal{I})} (p \supset \neg q) \vee (\neg p \supset \neg q)$ (consider a valuation M, where $M(q) = t$ and $M(p) = \top$). It is shown in [AA96] that in the case of $\models_{con}$, in order to add $\supset$ to the language without losing cautious cut, one has to add a certain constraint to this rule: Every atomic formulae that appears in the language of the cut formula(e) should also appear in the language of the premises.[5]

Proposition 3.17 For every logical bilattice $(\mathcal{B},\mathcal{F})$ and an inconsistency set $\mathcal{I}$
a) If $\Gamma \models^{\mathcal{B},\mathcal{F}} \Delta$ then $\Gamma \models^{\mathcal{B},\mathcal{F}}_{\mathcal{I}} \Delta$.
b) $\models^{\mathcal{B},\mathcal{F}}_{\mathcal{I}}$ is nonmonotonic.
c) $\models^{\mathcal{B},\mathcal{F}}_{\mathcal{I}}$ is paraconsistent.

Proof: Let $(\mathcal{B},\mathcal{F})$ be an arbitrary logical bilattice, and $\mathcal{I}$ – an inconsistency set in it. Then:
a) Immediately follows from the definitions of $\models^{\mathcal{B},\mathcal{F}}$ and $\models^{\mathcal{B},\mathcal{F}}_{\mathcal{I}}$.
b) Consider, e.g., $\Gamma = \{p, \neg p \vee q\}$. Every $\mathcal{I}$-mcm M of Γ must assign to both p and q consistent values (since the valuation that assigns t to p and f to q is an $\mathcal{I}$-mcm of Γ). Now, since $M(p) \in \mathcal{F}$, it follows that $M(\neg p) \notin \mathcal{F}$ (otherwise $M(p) \in \mathcal{I}$). Thus, in order that $M(\neg p \vee q) \in \mathcal{F}$, necessarily $M(q) \in \mathcal{F}$. Therefore $\Gamma \models^{\mathcal{B},\mathcal{F}}_{\mathcal{I}} q$. On the other hand, let $\Gamma' = \Gamma \cup \{\neg p\}$. Then $\Gamma' \not\models^{\mathcal{B},\mathcal{F}}_{\mathcal{I}} q$ ($N(p) = \top$, $N(q) = f$ is a counter $\mathcal{I}$-mcm of Γ').
c) Using the notations of the part (b), Γ' is an inconsistent theory and still $\Gamma' \not\models^{\mathcal{B},\mathcal{F}}_{\mathcal{I}} q$. $\Box$

Proposition 3.18 [AA96] Let $(\mathcal{B},\mathcal{F})$ be a logical bilattice $(\mathcal{B},\mathcal{F})$ and $\mathcal{I}$ – an inconsistency set in it.
a) If Γ and Δ are in the language of $\{\neg,\wedge,\vee,f,t\}$ and $\Gamma \models^{\mathcal{B},\mathcal{F}}_{\mathcal{I}} \Delta$, then the disjunction of the sentences in Δ classically follows from Γ.
b) Let Γ be a classically *consistent* set in the language of $\{\neg,\wedge,\vee,f,t\}$, and ψ – a clause that does not contain any pair of an atomic formula and its negation. If ψ classically follows from Γ, then $\Gamma \models^{\mathcal{B},\mathcal{F}}_{\mathcal{I}} \psi$.

[5] Cautious cut together with this condition is called there *analytic* cautious cut.

Like in the case of $\models_k^{\mathcal{B},\mathcal{F}}$ we have the following proposition:

Proposition 3.19
a) $\models_{\mathcal{I}}^{\mathcal{B},\mathcal{F}}$ preserves cautious cut: If $\Gamma, \psi \models_{\mathcal{I}}^{\mathcal{B},\mathcal{F}} \Delta$ and $\Gamma \models_{\mathcal{I}}^{\mathcal{B},\mathcal{F}} \psi, \Delta$ then $\Gamma \models_{\mathcal{I}}^{\mathcal{B},\mathcal{F}} \Delta$.
b) $\models_{\mathcal{I}}^{\mathcal{B},\mathcal{F}}$ preserves cautious monotonicity: If $\Gamma \models_{\mathcal{I}}^{\mathcal{B},\mathcal{F}} \psi$ and $\Gamma \models_{\mathcal{I}}^{\mathcal{B},\mathcal{F}} \Delta$, then $\Gamma, \psi \models_{\mathcal{I}}^{\mathcal{B},\mathcal{F}} \Delta$.

Proof: Similar to that of Proposition 3.10. $\square$

Several consequence relations similar to $\models_{\mathcal{I}}^{\mathcal{B},\mathcal{F}}$ are considered in the literature. Priest [Pr89, Pr91] uses a similar consequence relation, $\models_{\rm LPm}^3$, for defining the logic LPm from the three-valued logic LP (also known as Kleene 3-valued logic with middle element designated). It is well known that LP invalidates the Disjunctive Syllogism (i.e., if $\models_{\rm LP}^3$ denotes the consequence relation of LP, then $\psi, \neg\psi \vee \phi \not\models_{\rm LP}^3 \phi$). Priest argues that a consistent theory should preserve classical conclusions. He suggests to resolve this drawback by considering as the relevant models of a set Γ only those that are *minimally inconsistent*. Such models assign $\top$ only to some minimal set of atomic formulae. The consequence relation $\models_{\rm LPm}^3$ of the resulting logic, LPm, is then defined as follows: $\Gamma \models_{\rm LPm}^3 \psi$ iff every minimally inconsistent model of Γ is a model of ψ.

In our terms, Priest considers the inconsistency set $\mathcal{I} = \{b \mid b \in \mathcal{F}, \neg b \in \mathcal{F}\}$. In the 3-valued semantics this is the only inconsistency set, and it consists only of $\top$. In the general (multi-valued) case, however, there are many others. It follows that $\models_{\mathcal{I}}^{\mathcal{B},\mathcal{F}}$ might be viewed as a generalization of LPm. Moreover, in [AA97a] it is shown that a switch to a bilattice-based semantics might improve the inference process of LPm: By using, e.g., four-valued semantics and considering only the $\leq_k$-minimal valuations among the $\mathcal{I}_1$-mcms of a given theory, it is possible to infer the same conclusions as those obtained by $\models_{\rm LPm}^3$. The number of such models is usually smaller (and never bigger) than the number of the LPm-models. This is due to the fact that from every k-minimal $\mathcal{I}_1$-mcm one can construct several LPm-models by changing every $\bot$-assignment to either t or f. To see this in a particular case consider, e.g., the following simple example: Let $\Gamma = \{\neg p \vee q, p \vee q\}$. q follows from Γ according to $\models_{\rm LPm}^3$ and according to $\models_{\mathcal{I}_1}^4$ (and classically as well, of course). Now, Γ has *two* LPm-models: $M_1(p) = t, M_1(q) = t$ and $M_2(p) = f, M_2(q) = t$ (these are also its classical models). On the other hand, there is only *one* k-minimal $\mathcal{I}_1$-mcm of Γ: $N(p) = \bot, N(q) = t$. This single model suffices for inferring that q follows from Γ. Clearly, when the number of the atomic formulae that appear in the language of Γ increases, the amount of the k-minimal $\mathcal{I}_1$-mcms might become considerably smaller than the amount of the LPm-models of Γ.

Kifer and Lozinskii [KL92] also propose a similar relation (denoted there $\approx_\Delta$, where Δ stands for the values that are considered as representing inconsistent knowledge). This relation is considered in the framework of annotated logics ([Su90, KS92, Su94]). See [AA96] for a comparison between $\approx_\Delta$ and $\models_{\mathcal{I}}^{\mathcal{B},\mathcal{F}}$.

4 Conclusion and future work

Bilattices have had an extensive use in recent years, most notably in the area of logic programming. These structures are also useful as a semantic tool for defining multi-valued logics. The resulting consequence relations are strongly related to non-monotonic reasoning, and especially suitable for reasoning in the presence of inconsistency.

Despite all their appealing properties, the logics discussed above still lack *efficient* inference procedures. Among the issues that should be addressed in this context is whether it is possible to construct the subset of the preferred models of a given theory without computing the whole set of its models. Another major challenge is related with the problem of an efficient belief revision, i.e.: reducing the amount of computations needed for revising the set of conclusions when the knowledge-base is altered.

A preliminary method for efficiently constructing (four-valued) mcms is presented in [AA97b]. This approach is applied to knowledge-bases which are of a specific structure (called *stratified* knowledge-bases). Another possible approach for dealing with computational limitations is considered in [Le86, Wa94]. The method proposed there is to restrict the representation language, taking again into account the trade-off between expressiveness and efficiency. In both cases there is still much work to be done in order to obtain reasoning processes that are general enough on the one hand and that are computationally feasible on the other hand.

References

[AB75] A.R.Anderson, N.D.Belnap. *Entailment.* Vol.1, Princton University Press, Prinston N.J., 1975.

[AA94] O.Arieli, A.Avron. *Logical bilattices and inconsistent data.* Proc. 9th IEEE Annual Symp. on Logic in Computer Science, IEEE Press, pages 468–476, 1994.

[AA96] O.Arieli, A.Avron. *Reasoning with logical bilattices.* Journal of Logic, Language and Information, Vol.5, No.1, pages 25–63, 1996.

[AA97a] O.Arieli, A.Avron. *The value of the four values.* Technical Report 319/97, Department of Computer-Science, Tel-Aviv University, June 1997.

[AA97b] O.Arieli, A.Avron. *Four-valued diagnosis for stratified knowledge-bases.* Proc. 10th Ann. Conf. of the European Association for Computer Science Logic (CSL'96) D.Van-Dalen, M.Bezem – Eds. Lecture Notes in Computer Science No. 1258, Springer, pages 1-17, 1997.

[Av91] A.Avron. *Natural 3-valued logics: Characterization and proof theory.* Journal of Symbolic Logic Vol.56, No.1, pages 276–294, 1991.

[Av96] A.Avron. *The structure of interlaced bilattices.* Journal of Mathematical Structures in Computer Science, Vol.6, pages 287–299, 1996.

[Be77a] N.D.Belnap. *A useful four-valued logic.* Modern Uses of Multiple-Valued Logic (G.Epstein, J.M.Dunn – Eds.), Reidel Publishing Company, Boston, pages 7–37, 1977.

[Be77b] N.D.Belnap. *How computer should think.* Contemporary Aspects of Philosophy (G.Ryle – Ed.), Oriel Press, pages 30-56, 1977.

[Bi67] G.Birkhoff. *Lattice theory.* American Mathematical Society Coloquium Publications, Vol.25 (3rd eddition), 1967.

[dC74] N.C.A.da-Costa. *On the thory of inconsistent formal systems.* Notre Dame Journal of Formal Logic, Vol.15, pages 497-510, 1974.

[Do85] I.L.M.D'ottaviano. *The completeness and compactness of a three-valued first-order logic.* Revista Colombiana de Matematicas, Vol. XIX, No. 1–2, pages 31–42, 1985.

[Du86] J.M.Dunn *Relevant logic and entailment.* Handbook of Philosophical Logic, Vol.III (D.Gabbay, F.Guenthner – Eds.), 1986.

[Ep90] R.L.Epstein *The semantic foundation of logic. Vol.I: propositional logics.* Kluwer Academic Publisher, 1990.

[Fi90a] M.Fitting. *Bilattices in logic programming.* The 20th Int. Symp. on Multiple-Valued Logic (G.Epstein – Ed.), IEEE Press, pages 238–246, 1990.

[Fi90b] M.Fitting. *Kleene's logic, generalized.* Journal of Logic and Computation, Vol.1, pages 797–810, 1990.

[Fi91] M.Fitting. *Bilattices and the semantics of logic programming.* Journal of Logic Programming, Vol.11, No.2, pages 91–116, 1991.

[Fi93] M.Fitting. *The family of stable models.* Journal of Logic Programming, Vol.17, pages 197–225, 1993.

[Fi94] M.Fitting. *Kleene's three-valued logics and their children.* Fundamenta Informaticae, Vol.20, pages 113–131, 1994.

[Ga85] 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.

[Gi88] M.L.Ginsberg. *Multivalued logics: A uniform approach to reasoning in AI.* Computer Intelligence, Vol.4, pages 256–316, 1988.

[KL92] M.Kifer, E.L.Lozinskii. *A logic for reasoning with inconsistency.* Journal of Automated Reasoning. Vol.9, No.2, pages 179–215, 1992.

[KLM90] S.Kraus, D.Lehmann, M.Magidor. *Nonmonotonic reasoning, preferential models and cumulative logics.* Journal of Artificial Intelligence, Vol.44, No.1–2, pages 167–207, 1990.

[KS92] M.Kifer, V.S.Subrahmanian. *Theory of generalized annotated programming and it's applications.* Journal of Logic Programming, Vol.12, pages 335–367, 1992.

[Le92] D.Lehmann. *Plausibility logic.* Proc. 5th Ann. Conf. of the European Association for Computer Science Logic (CSL'91), Springer, pages 227–241, 1992.

[Le86] H.J.Levesque. *Making Belivers out of Computers.* Journal of Artificial Intelligence, Vol.30, pages 81-108, 1986.

[Ma94] D.Makinson. *General patterns in nonmonotonic reasoning*. Handbook of Logic in Artificial Intelligence and Logic Programming, Vol.3 (D.Gabbay, C.Hogger, J.Robinson – Eds.) Oxford Science Pub, pages 35–110, 1994.

[Pr89] G.Priest. *Reasoning about truth.* Journal of Artificial Intelligence, Vol.39, pages 231-244, 1989.

[Pr91] G.Priest. *Minimally inconsistent LP.* Studia Logica, Vol.50, pages 321–331, 1991.

[Ro89] L.I.Rozoner. *On interpretation of inconsistent theories.* Information Sciences, Vol.47, pages 243–266, 1989.

[Sh87] Y.Shoham. *A semantical approach to nonmonotonic logics.* Readings in Non-Monotonic Reasoning (M.L.Ginsberg – Ed.), Los-Altos, CA, pages 227–249, 1987.

[Sh88] Y.Shoham. *Reasoning about change: time and causation from the standpoint of artificial intelligence*. MIT Press, 1988.

[Su90] V.S.Subrahmanian. *Paraconsistent disjunctive deductive databases.* The 20th Int. Symp. on Multiple Logic (G.Epstein – Ed.), IEEE press, pages 339–345, 1990.

[Su94] V.S.Subrahmanian. *Amalgamating knowledge bases.* ACM Transactions on Database Systems, Vol.19, No.2, pages 291–331; 1994.

[Wa94] G.Wagner. *Vivid logic: Knowledge-based reasoning with two kinds of negations.* Lecture Notes in AI No.764, Springer-Verlag, 1994.