Logique & Analyse **145** (1994), 57-94

INCONSISTENCY-ADAPTIVE LOGICS AND THE FOUNDATION OF NON-MONOTONIC LOGICS[1]

Diderik BATENS

Abstract

In this paper I propose the reconstruction of (what I shall call) mixed non-monotonic logics as a combination of a deductive and a preferential component. The first leads from the premises to a possibly inconsistent consequence set; the second weeds out the inconsistencies. Among the candidates for the deductive component inconsistency-adaptive logics prove most suitable. The ensuing preferential component is formulated in terms of models and is itself split into two parts: (i) a transparent, purely logical procedure leads from a set of inconsistent models to a set of associated consistent models and (ii) the choice between the latter relies on the preferences. The real fight between mixed non-monotonic logics should concentrate on this last aspect. The outlined approach has a broader domain of application than mixed non-monotonic logics.

1. *Aim of this paper*

The first person who came across a bird that was not a flyer, was facing an inconsistency and had to decide which 'half' should be retained. At present, the belief that birds are flyers is considered a rule with exceptions. The belief may still lead to inconsistencies, but we have a standardised procedure to remove them. A rule with exceptions is arrived at by (devising and) adopting such a procedure and the procedure determines the functioning of the rule. In the present paper I pursue this line of thought and spell out an approach to rules with exceptions in terms of *two components*: a deductive component that may lead to inconsistencies, and a preferential component that selects one half from each inconsistency.

A number of logics devised to handle rules with exceptions are advertised under the label "non-monotonic logics"; I shall call them "mixed non-

[1] Previous versions of the present paper were presented in lectures in Krakow and Warsaw, and at a conference in Ghent. Special thanks to Andrzej Wronski, Ewa Orłowska, Joke Meheus and Laszlo Polos.

monotonic logics", as both components are blended in them. The label "mixed" separates these logics from non-monotonic logics — logics for which there are Γ, Δ, and A such that $\Gamma \vdash A$ and $G \cup \Delta \nvdash A$ — that are free of any preferential elements and hence logics in the strict sense.

My aim is *not* to offer another type of mixed non-monotonic logic, but rather a way to *reconstruct* existing ones. Such reconstruction has quite striking further advantages, which I discuss in section 2. The separation of both components may facilitate an ordering of the somewhat chaotic domain of mixed non-monotonic logics. It will moreover enable the logicians working in this field to concentrate on the *real* disagreement, viz. aspects of the preferential component. I argue, convincingly I hope, for this point in section 12, after I presented the formal results.

I shall consider several candidates for each component. An interesting candidate for the logical component is a mechanism presented as early as 1964 by Nicholas Rescher, next to (monotonic) paraconsistent logics and (non-monotonic) inconsistency-adaptive logics. The latter will prove by far superior to the alternatives.

The main technical results of the present paper concern the way in which the set of models arrived at by inconsistency-adaptive logics may be turned into a set of selected consistent models that defines the same consequences as mixed non-monotonic logics.

Some parts of the present paper are unavoidably very technical. I have taken great care to explain the results (as well as some unsolved difficulties) for people who prefer to skip the proofs of the theorems. Also, rushy readers might skip the sections on the reconstruction in terms of Rescher's mechanism. Their outcome is rather negative. Nevertheless, it is worthwhile to mention these results, as (i) Rescher's mechanism deserves a special tribute as the first non-monotonic system ever, (ii) several A.I. researchers continue working on this approach, and (iii) the lessons to be drawn from it are quite illuminating (and no technicalities are involved).

2. *Separating the deductive and the preferential component*

As soon as one considers examples that are not extremely simple, mixed non-monotonic logics lead to a variety of consequence sets that depend on the order in which the rules-with-exceptions (henceforth "non-monotonic rules") are applied. If Tweety is a penguin, "Penguins are birds" is believed unexceptionally, and both "Birds are flyers" and "Penguins are non-flyers" are non-monotonic rules, the only way to decide whether Tweety is or is not a flyer, is by relying on an ordering amongst the rules. What this really comes to is that applications of some rules are *preferred* over applications of other rules. In the example, applications of the least general rule will typically be preferred, thus leading to the conclusion that Tweety is a non-flyer. A (somewhat different) relation between non-monotonic rules and preferences is revealed in Yoam Shoham's model-theoretic approach that actually turns around the notion of a preferred model.

There are several annoying things about non-monotonic logics. I mention three. The first is the very fact that the deductive and preferential components are blended together. This makes it rather difficult to study systems of preferences and their general properties. The second nuisance is that the non-monotonic rules should be identified beforehand and that the preferential order should be fixed beforehand. This drastically restricts the domain of application of non-monotonic logics. For example, they do not shed any light on belief revisions that lead to non-monotonic rules by degrading general statements. That is not worse than classical logic, which also is unfit for really interesting thought episodes. Yet, a better situation is easy to imagine. The third nuisance is the rather wild proliferation of non-monotonic systems (some people produce at least three a year) and the fact that considering somewhat more complex examples steadily results in new approaches involving new complications. Even if one does not regard, with Kuhn, such proliferation as typical for dying paradigms, one cannot see it as a mark of great health either.

These difficulties may be overcome by separating the two mechanisms that are blended together within mixed non-monotonic logics. Roughly, the idea is to proceed as follows: given some set of premises Γ, a derivation of A from Γ and a derivation of $\sim A$ from Γ, prefer A (and drop $\sim A$) iff the first derivation does not contain applications of non-monotonic rules whereas the second does (and *vice versa*). Of course, the criterion will in general have to be more complicated.

On this analysis, the separation of the purely deductive component from the preferential one is straightforward. The deductive component should define a possibly inconsistent[2] consequence set, whereas the preferential component should take care of weeding out this set in such a way that the result is consistent. As a side-effect, this reformulation does not require non-monotonic rules to have a specific form; it only requires that suitable preference values be attached to certain statements.

The sketched procedure will not enable one to reconstruct all mixed non-monotonic logics. For example, some may require that each inconsistency be resolved as soon as it occurs. For some sets of premises, however, such logics result in different consequence sets depending on the order in which the derivation proceeds (in other words, they do not define a unique fixed

[2] Γ is *inconsistent* iff $\Gamma \vdash A$ and $\Gamma \vdash \sim A$ for some A. Obviously, consequence sets are inconsistent iff they contain an explicit inconsistency.

point for each set of premises). Although such 'indeterministic' non-monotonic logics may have some suitable applications, they are not within the mainstream of the research tradition and I shall disregard them in the present paper.

The proposed separation of the deductive and the preferential component will enable us to study them independently. We may look for a variety of candidates for the deductive component as well as study the properties of preferential systems. clearly, this separation can but clarify things (and lead to some more order in the present proliferation).

The separation has, however, a further important advantage. I mentioned already that one cannot even start applying some (mixed) non-monotonic logic unless the preferences are given (in suitable form). This eliminates important cases in which a logical analysis of the premises or additional information is required before the preferences are fixed.[3] If both components are separated, however, nothing prevents us from applying the deductive component to analyze the situation, to decide what further information is required, and to fix the preferences on the basis of this analysis or of information obtained. Also, one may decide to eliminate some inconsistencies only. In other words, the approach I propose not only clarifies things in the field of mixed non-monotonic logic, but, being more general, it also opens up new domains of application.

3. *Rescher's 1964 mechanism as a candidate for the deductive component*

The deductive component will, in interesting cases, lead to an inconsistent consequence set. One naturally expects some paraconsistent logic here, but first I shall consider a system that is not a paraconsistent logic in the usual sense, viz. Nicholas Rescher's 1964 mechanism.[4]

Given some possibly inconsistent set Γ of premises, we consider the maximally consistent subsets (m.c.s.) of Γ (the Δ such that (i) $\Delta \subseteq \Gamma$, (ii) Δ is consistent, and (iii) for any $A \in \Gamma - \Delta$, $\Delta \cup \{A\}$ is inconsistent).

[3] Joke Meheus 199+ studies such a case, and supports the claim that mixed non-monotonic logics are useless to understand creative processes that occur in the course of some scientific discovery. Her contribution moreover suggests that applications of mixed non-monotonic logics to discovery, conceptual shift, and the like, is bound to be restricted and very superficial.

[4] In Rescher 1964 the mechanism is applied to counterfactuals; other applications are presented in Rescher 1973 and Rescher and Manor 1970 — the latter deals explicitly with inference from inconsistent premises. As I see it, this mechanism was a central source of inspiration for Rescher and Brandom 1980.

Rescher defines the set of *strong* consequences of Γ as the set of the formulas derivable from *each* m.c.s. of Γ, and the set of *weak* consequences of Γ as the set of the formulas derivable from *some* m.c.s. of Γ. We are clearly interested in the set of weak consequences only, the set of strong consequences being always consistent and too poor to function as the logical component of mixed non-monotonic logics.

It is typical that weak consequence is defined with respect to some set of *consistent sets* of premises. The union of these sets may be inconsistent, but the sets never are. It is instructive to see what this comes to in terms of consistent models (of classical logic). Given a (possibly inconsistent) set Γ and a model *M*, *define* S^M as the set of formulas true in *M* and $S^{M\Gamma}$ as $\Gamma \cap S^M$. Next, *define* a maximal model of Γ as a model *M* such that, for all models *M'*, if $S^{M\Gamma} \subseteq S^{M'\Gamma}$, then $S^{M\Gamma} = S^{M'\Gamma}$ — clearly, a maximal model is a model of an m.c.s. *A* is a weak consequence of Γ iff there is a maximal model *M* of Γ such that, for all *M'*, if $S^{M\Gamma} = S^{M'\Gamma}$, then $A \in S^{M'}$. In other words, the set of weak (respectively strong) consequences[5] is a union (respectively intersection) of intersections of sets S^M. In a sense, a preferential semantics is already present here: among all (classical) models, we prefer those that contain a maximal number of members of Γ (and combine them in a certain way); if Γ is consistent, we obtain classical semantic consequence: we prefer the models in which all members of Γ are true.

Rescher's 1964 mechanism is a special case as it was supplemented from the very beginning with a preferential component. I shall, however, first comment on its suitability as a deductive component, and postpone to section 8 the discussion of the preferential component built into the mechanism by Rescher.

It is well known that Rescher's mechanism is extremely dependent on the formulation of the premises: if two premises are replaced by their conjunction, the resulting set of m.c.s. may change drastically and so may the sets of weak and strong consequences. Such procedure is quite sensible if the formulation of the premises has some specific rationale, for example if each premise describes the conjunction of information obtained from some source. Mixed non-monotonic logics are sometimes applied to such situations. Moreover, some mixed non-monotonic logics are themselves highly dependent on the formulation of the premises, viz. of the rules.

5 If the set of weak consequences is inconsistent, it has no models and its maximal models are identical to those of the set of premises. In Rescher and Brandom 1980, 'disjunction-worlds' (models in which are true all formulas true in two or more classical models) constitute models for inconsistent sets (but not for explicit contradictions such as $p \& \sim p$).

The upshot is that we should take Rescher's 1964 formalism as a serious candidate for functioning as the logical component of some non-monotonic logics. Its deductive properties are transparent and may be studied without having to deal in the same breath with the preferential component.

4. *Monotonic paraconsistent logics as candidates for the deductive component*

A second set of candidates for the logical component are (monotonic) paraconsistent logics. A host of these have been developed during the last twenty years (including relevant logics). As mixed non-monotonic logics are extensions of classical logic — henceforth **CL** — and as the only 'problem' for the deductive component is the occurrence of inconsistencies, I shall restrict my attention to the predicative versions of some of the logics of Batens 1980. The *weakest* logic discussed there contains the full positive fragment of classical logic as well as the negation-completeness 'half', $A \vee \sim A$, of the meaning of negation; its paraconsistent extensions contain such theorems as $\sim(A\&B) \supset (\sim A \vee \sim B)$ and $(A\&B) \supset (\sim(A\&B) \supset C)$. I list the syntax and semantics of the minimal logic **PIL** — I shall also need it for the inconsistency-adaptive logic discussed in section 5. The definition of terms, formulas and wffs (not containing any free variables) is as for **CL**, and functions are disregarded — including them is straightforward. As usual, $A(x)$is a formula in which x occurs free, and $A(a)$ is obtained from $A(x)$ by replacing every free occurrence of x in A by a. Also, the α and β should be interpreted in such a way that all (main) formulas are wffs.

Syntax

MP	From A and $A \supset B$ to derive B
$A\supset 1$	$A \supset (B \supset A)$
$A\supset 2$	$((A \supset B) \supset A) \supset A$
$A\supset 3$	$(A \supset (B \supset C)) \supset ((A \supset B) \supset (A \supset C))$
$A\&1$	$(A\&B) \supset A$
$A\&2$	$(A\&B) \supset B$
$A\&3$	$A \supset (B \supset (A\&B))$
$A\vee 1$	$A \supset (A \vee B)$
$A\vee 2$	$B \supset (A \vee B)$
$A\vee 3$	$(A \supset C) \supset ((B \supset C) \supset ((A \vee B) \supset C))$
$A\equiv 1$	$(A \equiv B) \supset (A \supset B)$
$A\equiv 2$	$(A \equiv B) \supset (B \supset A)$
$A\equiv 3$	$(A \supset B) \supset ((B \supset A) \supset (A \equiv B))$
$A{\sim}1$	$(A \supset \sim A) \supset \sim A$ (alternatively: $A \vee \sim A$)

$R\forall$ To derive $\vdash A \supset (\forall\alpha)B(\alpha)$ from $\vdash A \supset B(\beta)$, provided β does not occur in either A or $B(\alpha)$.

$A\forall$ $(\forall\alpha)A(\alpha) \supset A(\beta)$

$R\exists$ To derive $\vdash(\exists\alpha)A(\alpha) \supset B$ from $\vdash A(\beta) \supset B$, provided β does not occur in either $A(\alpha)$ or B.

$A\exists$ $A(\beta) \supset (\exists\alpha)A(\alpha)$

$A{=}1$ $\alpha = \alpha$

$A{=}2$ $\alpha = \beta \supset (A \supset B)$ where B is obtained by replacing in A an occurrence of α that occurs outside the scope of a negation by β[6]

Semantics

Let S be the set of sentential letters, P^r the set of letters for predicates of rank r, C and V the set of letters for individual constants and variables respectively, F the set of (open and closed) formulas, and N the set of formulas of the form $\sim A$.[7] A model is a couple $M = \langle D, v\rangle$ in which D is a set and v is an assignment function defined by:

$C1.1$ $v : \mathsf{S} \to \{0, 1\}$

$C1.2$ $v : \mathsf{C} \cup \mathsf{V} \to D$ is such that $D = \{v(\alpha) \mid \alpha \in \mathsf{C} \cup \mathsf{V}\}$[8]

$C1.3$ $v : \mathsf{P}^r \to \mathcal{P}(D^r)$ (the power set of the r-th Cartesian product of D)

$C1.4$ $v : \mathsf{N} \to \{0, 1\}$

The valuation function v_M determined by the model M is defined as follows:

$C2.1$ $v_M : \mathsf{F} \to \{0, 1\}$

$C2.2$ where $A \in \mathsf{S}$, $v_M(A) = v(A)$

6 It is a general characteristic of PIL that expressions that have the same extension cannot in general be replaced by each other within the scope of a negation; e.g., $\vdash_{\text{PIL}} p \equiv (p \& p)$ but $\nvdash_{\text{PIL}} \sim(p \& q) \equiv \sim((p \& q) \& q)$

7 This is not the place to discuss the extent to which negation is analyzed in this type of models. In some extensions of PIL the definition of N may be drastically simplified and the semantics may be further analyzed along the lines of section 11 of Batens 1986b. Other ways of analysis are in terms of positive and negative extensions of predicates, etc.

8 The requirement, which is obviously much weaker than ω-completeness, restricts the semantics to models with countable domain D, but greatly facilitates both some other clauses and the proofs.

$C2.3 \quad v_M(\pi^r\alpha_1 \ldots \alpha_r) = 1$ iff $\langle v(\alpha_1), \ \ldots, \ v(\alpha_r)\rangle \in v(\pi^r)$

$C2.4 \quad v_M(\alpha = \beta) = 1$ iff $v(\alpha) = v(\beta)$

$C2.5 \quad v_M(\sim A) = 1$ iff $v_M(A) = 0$ or $v(\sim A) = 1$

$C2.6 \quad v_M(A \supset B) = 1$ iff $v_M(A) = 0$ or $v_M(B) = 1$

$C2.7 \quad v_M(A \& B) = 1$ iff $v_M(A) = 1$ and $v_M(B) = 1$

$C2.8 \quad v_M(A \vee B) = 1$ iff $v_M(A) = 1$ or $v_M(B) = 1$

$C2.9 \quad v_M(A \equiv B) = 1$ iff $v_M(A) = v_M(B)$

$C2.10 \quad v_M((\forall\alpha)A(\alpha)) = 1$ iff $v_M(A(\beta)) = 1$ for all $\beta \in \mathsf{C} \cup \mathsf{V}$

$C2.11 \quad v_M((\exists\alpha)A(\alpha)) = 1$ iff $v_M(A(\beta)) = 1$ for at least one $\beta \in \mathsf{C} \cup \mathsf{V}$.

Derivability, theoremhood, truth in a model, semantic consequence and validity are defined as usual.

Clause $C1.4$ makes sure that v_M is fully determined by $M = \langle D, \ v\rangle$ Some care is at hand here. It would be unacceptable that the clauses characterizing the assignment function were dependent on each other. $C1.4$ should be taken literally, the value of $v(\sim A)$ being completely independent of the value of $v_M(A)$. As a result, $v(\sim A) = 1$ may but need not make the model inconsistent. Also, if v and v' differ from each other in $C1.4$, they may still determine formula-equivalent models, where M and M' are *formula-equivalent models* iff $v_M(A) = v_{M'}(A)$ for all (closed and open) formulas A.

Whereas the standard CL-semantics rules out the trivial model (in which all formulas are true), the present PIL-semantics does not. However, it is easy to show that it does not make any difference for the semantic consequence relation whether we include the trivial model or not.

An important similarity between CL-models and PIL-models is that, except for the trivial PIL-model, they are maximally non-trivial — this is proved (and commented upon) as Theorem 1 in Batens 199+. There are also important differences. PIL being much weaker then CL, PIL-models are by no means CL-models in which some more formulas may be true, thus making the set of true formulas inconsistent. Indeed, given that $v_M(\sim p)$ may be true, the truth of $v_M(p)$ does not entail the truth of $v_M(\sim\sim p)$. Another and more important difference is that an ω-complete CL-model is fully characterized by the set of sentential letters, primitive predicative wffs, and identities that are true in it — let us call these primitive wffs, and, when open formulas are included, *primitive expressions*. Such characterization will not do for PIL-models. Yet, one may recur to similar types of characterization.

Let there be an ordering over the members of C and of V, and let "$\alpha < \beta$" express either that $\alpha \in C$ whereas $\beta \in V$, or else that α precedes β in the ordering. The simplest (redundant) characterization of an ω-complete PIL-model M is a quartuple $\langle S^\circ, I^\circ, E^\circ, N^\circ \rangle$ in which S°, I°, E°, and N° are respectively the members of S, the identities (that are wffs), the primitive predicative wffs, and the formulas of the form $\sim A$ that are true in M. To obtain a non-redundant characterization, define C^m, the minimization of C with respect to M, as the set of those $\alpha \in C$ for which there is no $\beta \in C$ such that $\beta < \alpha$ and $v_M(\alpha = \beta) = 1$; next define a non-redundant characterization of an ω-complete PIL-model M, by requiring that $\alpha = \beta \in I^\circ$ only if $\alpha \in C^m$ and $\beta \notin C^m$, and that no elements of $C - C^m$ occur in E°. Finally, one may recursively define a N-*minimal characterization* of a PIL-model M by moreover restricting N° as follows. Where the complexity of a wff is the number of occurrences of connectives and quantifiers,[9] and the complexity of $\sim A$ is n, $\sim A \in N^\circ$ only if $\sim$ A is not a semantic consequence of $S^\circ \cup I^\circ \cup E^\circ$ together with the members of N° the complexity of which is smaller than n. Remark that these characterizations are not essentially more complex than those of PC-models. The only difference is N-minimality. Likewise, the proof of the following Theorem proceeds basically as for PC and is left to the reader.

Theorem 1. Any ω-complete PIL-model has a unique N-minimal characterization, and all ω-complete PIL-models that have the same N-minimal characterization are wff-equivalent (for all wffs, $v_M(A)$ is identical for all M).

Corollary 1. The N-minimal characterization of an ω-complete PIL-model M completely determines the value that v_M assigns to wffs.

N-minimal characterizations are often useful, but have the disadvantage to be restricted to ω-complete PIL-models. For some applications, it is handy to have a tool that works for other kinds of models as well. With the present PIL-semantics this is easily accomplished, viz. by considering formulas instead of wffs and members of $C \cup V$ instead of members of C. This time I shall only consider one kind of characterization, viz. the one that will prove most useful later. Let the *canonical characterization* of a PIL-model M be the quartuple $\langle S^\circ, I^\circ, E^\circ, N^\circ \rangle$ in which S° is as before, I°, E°, and N° contain closed as well as open formulas, and N° is minimized as in the definition of the N-minimal characterization of an ω-complete PIL-model

[9] For some paraconsistent extensions of PIL — see below in the text — it is handier to define the complexity of a formula in a different way.

M. clearly, the canonical characterization of a model M is redundant. Moreover, it is obvious that it is possible to turn any canonical characterization into a model by (i) identifying O^m as the set of those $\alpha \in C \cup V$ for which there is no β such that $\beta < \alpha$ and $\alpha = \beta \in I^\circ$, (ii) setting $D = O^m$, and (iii) defining v from S°, I°, E°, and N° in the obvious way. Given this, it is obvious that:

Theorem 2. Any PIL-model has a unique canonical characterization, and two PIL-models have the same canonical characterization iff they are formula-equivalent.

Corollary 2. The canonical characterization of a PIL-model M completely determines v_M.

Finally, it is useful to also have the notion of N-minimality at the semantic level. Define: a model is N-*minimal* iff, for *any* A, $v(\sim A) = 0$ whenever $v_M(A) = 0$. In other words, in a N-minimal model it holds that, for *any* A, $v(\sim A) = 1$ iff $v_M(A) = v_M(\sim A) = 1$. This leads to a useful theorem, proved in Batens 199+, which we shall need in section 10:

Theorem 3. For any PIL-model M, there is a formula-equivalent N-minimal PIL-model M'.

There is a whole family of logics between PIL and CL. For example, one may add such semantic clauses as $v_M(\sim(A \& B)) = v_M(\sim A \vee \sim B)$ (and the corresponding axioms). Some of these logics are maximally paraconsistent, viz. any extension of them leads to either CL or the trivial logic (in which every wff is a theorem). For more information on the propositional logics, see Batens 1980.

Consider some set of premises Γ in which the non-monotonic rules have been turned into universally quantified statements. If $Cn_{\mathrm{CL}}(\Gamma)$ is inconsistent and hence trivial, $Cn_{\mathrm{PIL}}(\Gamma)$ too will be inconsistent but not trivial (except for some border cases). Hence, PIL seems a sensible candidate for the deductive component of mixed non-monotonic logics. However, I shall show in section 7 that an inconsistency-adaptive logic based on PIL is more suitable in this respect.

In view of Theorem 3, it is safe to concentrate on N-minimal PIL-models. It is useful for future reference to consider a CL-model as a PIL-model that is formula-equivalent to a N-minimal PIL-model in which $v(\sim A) = 0$ for all A. It is indeed obvious that the description of any such model may be transformed into a description that agrees with the (rather, any) standard CL-semantics.

5. *Adaptive logics*

Consider a theory $\langle \Gamma, \mathsf{L} \rangle$, where Γ is the set of axioms and **L** the underlying logic. **L** will contain several presuppositions about the domain described. For example, **CL** presupposes that the domain (as approached by observational, operational, or other criteria) is consistent (that the criteria do not, for some A, lead to both A and $\sim A$). Γ may violate some of these presuppositions, in which case we shall say that Γ has *abnormal* properties (with respect to the intended underlying logic). For example, the **CL**-consequences of Γ may turn out to be inconsistent or to assert incompleteness (by way of non-logical theorems of the form $\sim(A \vee \sim A)$). If the abnormal properties cannot be readily removed, or if we have to reason about $\langle \Gamma, \mathsf{CL} \rangle$ in order to improve this theory, then neither **CL** nor a monotonic weakening of **CL** will do.

Here *adaptive* logics come in. They localize the abnormal properties of the theory, safeguard the theory for triviality by preventing specific rules of **L** from being applied to abnormal consequences of Γ, but behave exactly like **L** in all other cases.

The easiest way to understand how all this proceeds, is to realize that an adaptive logic 'oscillates' between the original logic **L** and a fragment **Lf** of **L** that differs from **L** in not sanctioning the abnormal properties involved. If the abnormal property displayed by Γ is inconsistency, **Lf** will allow for inconsistencies (will not lead from inconsistency to triviality); if the abnormal property is (negation-)incompleteness **Lf** will allow for incompleteness (by not having such theorems as $A \vee \sim A$ or such rules as $A \supset B, \ \sim A \supset B / B$). That the adaptive logic **La** will oscillate between **L** and **Lf** may now be characterized intuitively, but somewhat inaccurately, by saying that **La** allows for the application of all rules of **L**, except for applications to consequences of Γ for which it is derivable from Γ that they display abnormal properties. This formulation is inaccurate because the "derivable" is not specified. The correct specification is somewhat complicated, but if one studies the specification and its rationale, the outcome appears extremely intuitive — see Batens 1986a, 1989, and 199+ for the inconsistency adaptive case, other cases being analogous.

I hope the previous paragraph clarifies that an adaptive logic *localizes* the abnormal properties. At the syntactic level, a rule operates on finite sets of consequences of Γ (as for other logics); if a rule presupposes that an abnormal property is not involved, then it will be applicable or not applicable according as it is or is not derivable from Γ that the wffs included in the set have the abnormal property. In other words, it prevents abnormal properties of specific wffs to turn the consequence set of Γ into the trivial set, but does not restrict the rules of **L** in as far as they are applied to consequences of Γ that have not the abnormal property. If applied to a normal theory, nothing has to be restricted and the adaptive logic **La** leads to exactly the same set of consequences as **L** itself.

Another way to look upon adaptive logics is to say that they minimize (the consequences of) abnormality. **L** presupposes normality. **Lf** gives up some form of normality, thus heavily restricting on the set of consequences of Γ **La** takes into account that Γ is abnormal at specific points, but goes on presupposing normality elsewhere, thus leading to a set of consequences that is a real subset of the **L**-consequence set iff the latter is trivial, but is in general a real superset of the **Lf**-consequence set.

It should be stressed that adaptive logics do not require any inventiveness (or even any intervention) on the part of whoever applies them: applying the adaptive logic leads to correct (although not necessarily interesting) results. Also, adaptive logics (at least, those I have in mind here) have a nice and intuitive semantics that is directed precisely at minimizing abnormality.

Adaptive logics are non-monotonic: if $\Gamma \cup \Delta$ is more abnormal than Γ, some B derivable from the latter need not be derivable from the former. Some adaptive logics (e.g., the one discussed in section 6) are decidable at the propositional level and exactly as undecidable as classical logic at the predicative level.

To end this section, I record some facts. Adaptive logics differ from mixed non-monotonic logics because of two (related) properties: (i) they do not involve any non-logical preferences and (ii) they do not rule out the abnormal properties. In this sense, they form the purely logical basis for some non-monotonic logics: they localize the problems but do not resolve them.

In his (1991) Graham Priest invokes adaptive logics to an end that is completely different from the one I originally intended, but proves very interesting from his philosophical stand. Priest is a dialetheist who believes in the existence of a true logic, which he takes to be paraconsistent. He agrees, however, that in many situations we are justified in presupposing consistency. He goes on to show that, if his preferred paraconsistent logic **LP** (from his 1987) is turned into an adaptive logic $\mathbf{LP}^m$ by assuming consistency until and unless shown otherwise, then $\mathbf{LP}^m$ recaptures all classical reasoning where it is sensible (according to his so qualified dialetheist view).

6. *The inconsistency adaptive logic* APIL2

The predicative inconsistency-adaptive logics **APIL1** and **APIL2** are described in detail in Batens 199+. Both have a *dynamic* proof procedure:

some formulas may only be added to the proof under certain conditions and should be deleted (or marked "out") if these criteria are no longer fulfilled.

Although this dynamics depends on the formulas that occur in the proof (and the way in which some of them have been derived), it is possible to define, for both **APIL1** and **APIL2**, a static notion of final derivability. It may be shown that, if A is finally derivable from Γ, then A is derivable in the extension of any proof from Γ. The set of formulas that are finally derivable from Γ may be shown to be identical to the set of semantic consequences of Γ.

In the present paper, I shall restrict my attention to **APIL2**, and more specifically to its semantics. But first we need some definitions.

Where $(A\&{\sim}A)$ is a formula in which the variables $\alpha_1, \ldots, \alpha_k (k \geq 0)$ occur free (in that order), let $\exists(A\&{\sim}A)$ be $(\exists\alpha_1)\ldots(\exists\alpha_k)(A\&{\sim}A)$. Let $DEK(A_1, \ldots, A_n)$ refer to $\exists(A_1\&{\sim}A_1)\vee\ldots\vee\exists(A_n\&{\sim}A_n)$ — a disjunction of one or more (where necessary) existentially quantified contradictions. I shall say that $A_1, \ldots, A_n$ are the *factors* of $DEK(A_1, \ldots, A_n)$. As permutations of the factors and of the quantifiers in "$\exists$" result in equivalent formulas, I shall from now on use sets to refer to any such permutation. Remark that $DEK(\Sigma\cup\{Px\})$ is **PIL**-equivalent to $DEK(\Sigma\cup\{Py\})$ and is **PIL**-derivable from $DEK(\Sigma\cup\{Pa\})$, but that neither Pa nor Py is a factor of $DEK(\Sigma\cup\{Px\})$. For the sake of generality, $DEK(\emptyset)\vee A$ will be A.

The *semantics* for **APIL2** is both enlightening and intuitive. First remark that all **CL**-models are **PIL**-models. If Γ is consistent, then its consistent **PIL**-models are its **CL**-models, but it will have inconsistent models as well (except for the border case in which Γ is maximally non-trivial). If Γ is inconsistent, it will have inconsistent **PIL**-models only (and no **CL**-models). For any **PIL**-model M, let $EK(M)=\{A|v_M(\exists(A\&{\sim}A))=1\}$. The set of the **APIL2**-models of Γ is defined by $\{M|$ there is no M' such that, $EK(M')\subset EK(M)\}$, in other words, the set of **APIL2**-models of Γ contains those **PIL**-models of Γ that are not more inconsistent than is required by Γ. I summarize the situation in Figure 1: in both halves of the figure, the two larger ellipses represent all **PIL**-models and all **CL**-models respectively; the circle represents the **PIL**-models of Γ; the smallest field marked Γ represents the **APIL2**-models of Γ. If Γ is inconsistent, its **APIL2**-models form in general a real subset of its **PIL**-models, if Γ is consistent, its **APIL2**-models coincide with its **CL**-models. For all three logics, $\Gamma \models A$ is defined as usual (there is no model in which all members of Γ are true and in which A is false). In other words, Γ has in general more **APIL2**-consequences than it has **PIL**-consequences and if Γ is consistent, its **APIL2**-consequences are identical to its **CL**-consequences.

It seems worth mentioning that this semantics too is a form of preferential semantics: among the PIL-models of Γ, we prefer those that are as consistent as possible with respect to Γ. Clearly, APIL2 is non-monotonic, but it is not a mixed non-monotonic logic (as it does not involve any extra-logical preferences).

APIL2 will be our third candidate for the deductive component of a mixed non-monotonic logic. To every paraconsistent extension of PIL corresponds an inconsistency-adaptive logic. For reasons already mentioned in Batens 1989, it does not seem likely that these other inconsistency-adaptive logics will be preferable to APIL2, except perhaps for some special situations.

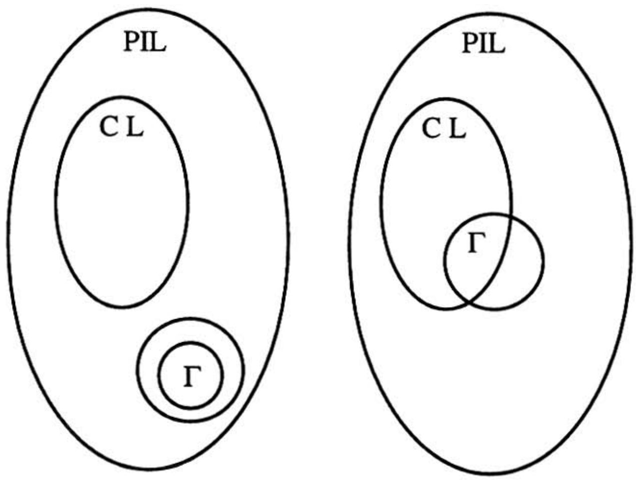

Figure 1

There are many brands of paraconsistent logics besides PIL (based on weaker positive logics) and most of them allow one to define an inconsistency-adaptive logic. I cannot study such systems in the present paper.[10] Also, some logics that are adaptive with respect to other forms of abnormality seem to some extent promising; I hope to publish some results on such logics soon.

[10] A nice example is the adaptive logic presented in Priest 1987 and 1991.

7. *Why adaptive logics are more promising than (monotonic) paraconsistent logics.*

Consider the following set of premises — the extended Tweety example:

(7.1) $(\forall x)(Bx \supset Fx)$, $(\forall x)(Px \supset Bx)$, $(\forall x)(Px \supset \sim Fx)$, $Pa \& Bb$

with the first premise preferred less than the other three. The PIL-consequence set of (7.1) contains the following atoms:

(7.2) Pa, Ba, Fa, $\sim Fa$, Bb, Fb

Precisely these very same atoms are derivable from (7.1) by (full classical) positive logic. The APIL2-consequence set contains the atoms:

(7.3) Pa, Ba, Fa, $\sim Fa$, $\sim Pb$, Bb, Fb, $\sim P\alpha$ (for all other constants α)

In contradistinction to the PIL-consequence set, the APIL2-consequence set contains $(\forall x)(x \neq a \supset \sim Px)$. This may seem a bit drastic from an intuitive point of view, but it is exactly what one would obtain by the usual circumscription procedure. Indeed, circumscription presupposes that anything that is not bound to be abnormal in view of the premises, is normal; as all penguins are abnormal birds (with respect to being a flyer), all birds not given to be non-flyers are supposed to be flyers (and hence non-penguins).

A *first* central difference is that $\sim Pb \vee (Fb \& \sim Fb)$ is true in all PIL-models of (7.1), whereas neither disjunct is. As APIL2 minimizes inconsistencies, $\sim Pb$ is true in *all* APIL2-models of (7.1). A *second* central difference is that no APIL2-model of (7.1) contains any inconsistency except for Fa and $\sim Fa$, whereas *any* inconsistency will be true in some PIL-model of (7.1). So, APIL2 has the advantage that the preferential component will only have to eliminate inconsistencies that are in a natural sense unavoidable in view of the premises. A different way to put this is by saying that APIL2-models have the important property to be not more inconsistent than is required by the premises. As explained in section 6, if M is an APIL2-model of some set of premises, and $EK(M) \subset EK(M')$, then M' is not an APIL2-model of the same set of premises.[11]

[11] In other words, APIL2 minimizes the inconsistencies in view of the premises; see also Theorem 14 of Batens 199+.

8. *Eliminating inconsistencies: Rescher 1964 style*

The reader will remember from section 3 that Rescher's mechanism proceeds in terms of consistent models only. So does his preferential component. If preferences are attached to the *premises*, one may define a preference order on the m.c.s., e.g. by assigning them the preference of their least preferred member. Next, some criterion may be applied to select the set of preferred m.c.s.; this criterion may be absolute (e.g., those at least preferred to degree n) or relative (e.g., the most preferred ones). According to another procedure, a m.c.s. is preferred iff it contains all members of the premises that have a preference ranking not worse than some limit. The sentences true in all preferred m.c.s. are the preferred consequences of the premises.

It is obvious that this procedure may be modified by defining a preference ordering on the maximal models (as defined in section 3). Also, the preference ordering defined on the models should *not* necessarily be derived from a preference ordering on the premises. E.g., with respect to abnormality predicates (as used in circumscription approaches) the preference ordering of the models might be based on the set of abnormal entities that occur according to the models.

Let us briefly see what comes out of Rescher's mechanism if we apply it to a simple example. Consider again the premises (7.1) with the first premise preferred less than the other three. The maximal models are those in which exactly three out of the four premises are true — $\sim(\exists x)Px$ is true in the models in which the first three premises are true. The set of weak consequences of (7.1) contains the following atoms:

(8.1) $Pa, \sim Pa, Ba, \sim Ba, Fa, \sim Fa, \sim Pb, Bb, Fb, \sim P\alpha$ (for all constants α)

Let the set of preferred models be those in which (only) the first premise is false. The set of all preferred consequences is the CL-consequence set of the (redundant) consistent set containing:

(8.2) $Pa, Ba, \sim Fa, Bb, (\forall x)(Px \supset Bx), (\forall x)(Px \supset \sim Fx)$

This is clearly inadequate, as b should come out a flyer. That is the whole point of mixed non-monotonic logics.

At first sight, this inadequacy seems easily repaired by proceeding as follows. Whenever a universally quantified premise (or an equivalent formula) occurs among the premises, we add all its instances to the set of premises. Next, we consider the maximal models of this so modified set. The aim of this modification is to allow for the possibility of leaving out a universally quantified formula together with its problematic instances —

$(\forall x)(Bx \supset Fx)$ and $Ba \supset Fa$ in the previous example — but to leave in all its other instances — $Bb \supset Fb$, $Bc \supset Fc$, ... in the previous example. The atoms among the weak consequences of (7.1) still are as in (8.1), but the set of atoms true in all preferred models now contains:

(8.3) $Pa, Ba, \sim Fa, \sim Pb, Bb, Fb, \sim P\alpha$ (for any constant $\alpha \neq a$)

However attractive this result may seem, the underlying procedure is not very promising. As the essential aspect of Rescher's mechanism consists in selecting consistent sets of *premises*, any previous logical operation on premises ruins the enterprise. Indeed, according to both the letter and the spirit of this mechanism, the instances of the universally quantified formulas should be added *in conjunction* with the universally quantified statement itself. In that case, however, we are back at the original version, represented by (8.1)-(8.2) — *all* instances of $(\forall x)(Bx \supset Fx)$ are removed together. Even worse, there seems to be no sensible formal criterion to add the instances as separate premises. Suppose indeed that the premises are given thus:

(8.4) $Pa\&(\forall x)(Bx \supset Fx)$, $(\forall x)(Px \supset Bx)$, $(\forall x)(Px \supset \sim Fx)$, Bb

Adding the premises

(8.5) $Pa\&(Ba \supset Fa)$, $Pa\&(Bb \supset Fb)$, ...

would involve a terrible bias in favour of Pa, whereas adding the premises

(8.6) $(Ba \supset Fa)$, $(Bb \supset Fb)$, ...

would come to actually disconnecting $(\forall x)(Bx \supset Fx)$ from Pa, which again runs counter to the letter and spirit of Rescher's mechanism.

Rescher's mechanism does not seem very promising for the present purposes. (This should by no means be read as an objection to Rescher: he never meant the procedure to serve that purpose in the first place.) In general, mixed non-monotonic logics that proceed in terms of subsets of the set of premises — a few complications have been proposed after Rescher 1964 — are themselves not very promising. They will adequately handle a few special cases, but run into trouble as soon as one considers arbitrary sets of premises. Here are some premises that may replace one or more of those in (7.1): $(\forall x)(Bb\&(Bx \supset Fx))$, $(\forall x)((Px \supset Bx)\&(Bx \supset Fx))$, $\sim(\exists x)(Pa \supset (Bx \supset \sim Fx))$, $\sim(Pa\&(\exists x)(Bx\& \sim Fx))$ It is possible that such mixed non-monotonic logics may be reconstructed in terms of Rescher-like mechanisms. However, more sensible mixed non-monotonic logics, such as those based on circumscription, cannot be so reconstructed.

9. Weeding out inconsistencies: first results

From now on I suppose that the logical component may lead to inconsistent models. We have to apply the preferential component in such a way that the result is a set of consistent models. In the present section, I disregard the question whether PIL or APIL2 has been used as the logical component and concentrate merely on the elimination of inconsistencies. Also, I shall suppose that some preference ordering is given, and not discuss its possible origin or foundation.

The most obvious approach seems to proceed as follows: first select all consistent models M such that the formulas true in M are true together in some (non-trivial[12]) PIL-model of the premises, and next choose between these consistent models by means of the preferences. This, however, doesn't work because the first step will lead to the selection of the empty set of models, as the following theorem shows.

Theorem 4. If Γ is inconsistent and M is a non-trivial PIL-model of Γ, then there is no CL-model M' such that all formulas true in M' are true in M.

Proof. Let M be a non-trivial PIL-model of an inconsistent set Γ and suppose that M' is a CL-model. As Γ is inconsistent, there is an A such that $v_M(A) = v_M({\sim} A) = 1$. Forcibly either $v_{M'}(A) = 0$ or $= v_{M'}({\sim} A) = 0$; suppose $v_{M'}(A) = 0$. It follows that $v_{M'}(A \supset B) = 1$ for all B. If all formulas true in M' are true in M, then $v_M(A \supset B) = 1$ for all B. But then, as $v_M(A) = 1$, $v_M(B) = 1$ for all B. Hence M is the trivial model. ■

So, we have to look for a procedure that relies upon the preferences and enables us to weed out the inconsistencies by *transforming* the PIL-models into CL-models. Considering N-minimal characterizations of the PIL-models offers some help here. A PIL-model is inconsistent iff N° is not empty. So, we should transform the PIL-models in such a way that N° is empty in the N-minimal characterizations of the resulting models.

Remark that there is no reason why the same inconsistencies would occur in all the PIL-models of some set of premises Γ. Indeed, if $\Gamma \vdash$

[12] If we do not rule out the trivial model, then obviously all consistent models qualify, whence the effect of the logical component reduces to nil.

$DEK(A, B)$ whereas $\Gamma \nvdash DEK(A)$ and $\Gamma \nvdash DEK(B)$, then $DEK(A)$ will be true in some of the models and false in others. Moreover, inconsistencies in PIL-models may be related to each other in two different ways. Each of these need special attention.

The first form of relation between inconsistencies obtains if the elimination (in a specific way) of one inconsistency involves or presupposes the elimination of another inconsistency. Suppose that A, $\sim A$, B, and $\sim(A\&B)$, and hence also $A\&B$, are true in some model M. If the inconsistency consisting of A and $\sim A$ is removed by transforming M to M' in such a way that A is false in M', then $A\&B$ is false in M' as well, and hence the inconsistency consisting of $A\&B$ and $\sim(A\&B)$ is removed at once. On the other hand, if the inconsistency consisting of $A\&B$ and $\sim(A\&B)$ is removed by transforming M into M' in which $A\&B$ is false, then either A or B is false in M'. If A is false in M', the inconsistency consisting of A and $\sim A$ is removed at once.

The second form of relation between inconsistencies obtains if the elimination (in a specific way) of one inconsistency entails the introduction of another inconsistency. Suppose that A, $\sim A$, $\sim B$, and $\sim(A \supset B)$ are true in some model M, whereas B is false in M. If the inconsistency consisting of A and $\sim A$ is removed by transforming M into a model M' in which A is false, then $A \supset B$ is true in M'; if the other sentences true in M are true in M', then both $A \supset B$ and $\sim(A \supset B)$ will be true in M'. This problem seems rather difficult. Should it be allowed that the removal of one inconsistency results in the introduction, at least temporarily, of another inconsistency? And will it be possible to find some procedure that leads to the elimination of all inconsistencies, and does so in a deterministic way?

I think this is the right time for an extremely important remark. We have to operate on models in order to show that the preferential component serves its purpose. But this does not mean that the preferences would be determined by the logical structure of the models. Quite to the contrary, preferences will be determined by non-logical features. Some statements true in a model will be observational data, others will be laws of nature, still others will derive from the combinations of observations and unquestioned laws. All of these will obtain the maximal preference value in the usual examples to which mixed non-monotonic logics are applied. In other words, we should not try to find a general solution on the basis of formal properties for the last example. If $\sim A$ and $\sim B$ are empirical data, and $\sim(A \supset B)$ is just a side effect of the fact that A is true in M whereas B is false in M (hence $A \supset B$ is false in M), then we should decide to make $\sim(A \supset B)$ false in M'. But if $\sim(A \supset B)$ has a higher preference than $\sim A$ and $\sim B$, then the inconsistency should not be removed by making A false in the first place.

There is another, quite puzzling problem with respect to PIL-models of inconsistent sets of premises. Most such models contain inconsistencies that are not even related to the premises (some PIL-models of $p\&\sim p$ contain $r\&\sim r$, etc.) and it is by no means clear how these should be eliminated. The fact that we have preferences about such statements as "All birds fly." does by no means offer us a way to handle "The moon both is and is not made of blue cheese." As we saw in section 7, this is precisely the (second and major) reason why APIL2 proves superior as the logical component: the APIL2-models are minimally inconsistent with respect to the premises. This warrants that all inconsistencies that occur in them are directly related to the premises.

10. *Weeding out inconsistencies from* APIL2*-models: defining a set of consistent models*

If Γ is inconsistent, so are all APIL2-models of Γ. The problem we are facing is to transform the set of these models into a set of consistent models, taking into account the preferences. This procedure should be deterministic: it should not depend on the order in which inconsistencies are removed. To guarantee so turns out somewhat tricky, as the reader may expect. For this reason, I proceed in two steps. The first step leads from the set of APIL2-models of Γ to a set of consistent models, which I shall call *the consistent models associated with* Γ; this step depends on purely logical considerations, not on the preferences. In the second step we select, on the basis of the preferences, a subset of the consistent models associated with Γ; the members of this subset will be called *selected consistent models associated with* Γ. The first step contains several complications and will be described in the present section. The second step is postponed to the next section.

Each APIL2-model M of Γ will be transformed into one or more consistent models associated with M. The set of consistent models associated with Γ will consist of the consistent models associated with an APIL2-model of Γ. The transformation itself is defined by a set of sets of *requirements*. Each set of requirements determines one or more consistent models arrived at by minimally transforming — see below — the original model. Needless to say, there will be no transformation at all if the APIL2-models of Γ are consistent. In this case, there will be only one set of requirements and it will be empty.

The sets of requirements

Where M is an APIL2-model of Γ, consider the canonical characterization of M, hereafter called CC_M. $\sim A \in \mathrm{N}^{\circ}$ iff $V_M(A) = V_M(\sim A) = 1$. So, in order to remove the inconsistencies from M it is necessary to require either "A is false" or "$\sim A$ is false" for any $\sim A \in \mathsf{N}^\circ$. Let us call a set of requirements *coherent* iff it agrees with at least one CL-model.[13]

The *initial* sets of requirements, IR_1, IR_2, ..., are the coherent sets that contain, for each $\sim A \in \mathsf{N}^\circ$, either "$A$ is false" or "$\sim A$ is false" — as a border case, the (single) initial set of requirements is empty if N° is empty. Each of the IR_i is extended and possibly 'split' into several alternative sets according to the rules of Table 1 — it is always supposed that $\alpha, \beta \in \mathsf{C} \cup \mathsf{V}$. Remark that the only formulas for which there is no rule are primitive expressions (in which members of V may occur free). Some applications of the rules of Table 1 may lead to incoherent sets of requirements. Incoherent sets are discarded at once. The extension comes to an end if the correct rule has been applied to all requirements for which there is a rule in Table 1. At that stage, the sets of requirements as well as the set of these sets will be called "*finished*". Let $FR_{i,1}$, $FR_{i,2}$, ... be the finished sets deriving from IR_i

All $FR_{i,j}$ are (coherent and) well-defined but cannot in general be effectively written down. As we shall see in section 11, it is easy enough to spell out the effects of the procedure I am describing for the intended examples.

By inspecting the rules in Table 1, one readily sees that:

Lemma 1. The order in which the rules from Table 1 are applied has no effect on the finished set of sets of requirements.

Lemma 2. Applying the rules of Table 1 to any initial set of requirements IR_i results in at least one finished set of requirements $FR_{i,j}$.

Proof. There is a CL-model (i.e. PIL-model with $v(\sim A) = 0$ for all A), say M, that meets all the requirements in IR_i. By inspecting Table 1, it is easily seen that, for any rule, (i) if the rule does not require the set to be split, then the added requirements are met by M, and (ii) if the rule requires the set to be split, then the requirements added to one of the subsets is met by M. Whence the Lemma follows by induction. ■

Canonical characterizations of consistent models associated with Γ

Having started from an APIL2-model M, how do we get to the consistent models associated with M from the finished set of sets of requirements? We shall do so by first defining the canonical characterizations of those models

[13] A set may be incoherent because it contains, for some CL-theorem A, the requirement "A is false"; or because it contains, for some A, requirements to the effect that A and $\sim A$ have the same truth-value; etc.

given:	add:
$(\forall\alpha)A(\alpha)$ is true	for all β, add "$A(\beta)$ is true"
$(\forall\alpha)A(\alpha)$ is false	split the set into a copy *for each* β and add in each copy "$A(\beta)$ is false" (for that β)
$(\exists\alpha)A(\alpha)$ is true	split the set into a copy *for each* β and add in each copy "$A(\beta)$ is true" (for that β)
$(\exists\alpha)A(\alpha)$ is false	for all β, add "$A(\beta)$ is false"
$A \supset B$ is true	split the set into two copies, add "*A* is false" to the first and "*B* is true" to the second
$A \supset B$ is false	"*A* is true" and "*B* is false"
A&B is true	"*A* is true" and "*B* is true"
A&B is false	split the set into two copies, add "*A* is false" to the first and "*B* is false" to the second
$A \vee B$ is true	split the set into two copies, add "*A* is true" to the first and "*B* is true" to the second
$A \vee B$ is false	"*A* is false" and "*B* is false"
$A \equiv B$ is true	split the set into two copies, add "*A* is true" and "*B* is true" to the first and "*A* is false" and "*B* is false" to the second
$A \equiv B$ is false	split the set into two copies, add "*A* is true" and "*B* is false" to the first and "*A* is false" and "*B* is true" to the second
~*A* is true	"*A* is false"
~*A* is false	"*A* is true"

Table 1

from the set of requirements and CC_M. The idea is this: each finished set of requirements $FR_{i,j}$ defines a set of primitive expressions that have to be true in the consistent model; the latter set will be combined with as much primitive expressions from CC_M as is possible in view of (i) the requirements in $FR_{i,j}$ and (ii) the definition of a CL-model.

For each finished set of requirements $FR_{i,j}$, we proceed as follows. Let R be the set of the primitive expressions, that are required to be true by $FR_{i,j}$. Some subsets of $S° \cup I° \cup E° \cup N° \cup R$ are compatible with some CL-model and fulfil all requirements in $FR_{i,j}$. Let $MSS_{i,j,1}$, $MSS_{i,j,2}$, ... be the *maximal* subsets[14] of $S° \cup I° \cup E° \cup N° \cup R$ that are compatible with some CL-model and fulfil all requirements in $FR_{i,j}$.

No members of N° occur in any $MSS_{i,j,k}$, and there is a unique maximal subset of $S° \cup E°$ that is compatible with R. The fact that there may be several $MSS_{i,j,k}$ is caused solely by the members of I°. Here is an example: if $a = b$, $\sim a = b$, $a = c$ and $b = c$ are true in M, and $FR_{i,j}$ contains the requirement "$a = b$ is false", then there will be $MSS_{i,j,k}$ of which $a = c$ is and $b = c$ is not a member, and there will be $MSS_{i,j,k}$ of which $b = c$ is and $a = c$ is not a member.

Let $CMSS_{i,j,k}$ be the closure of $MSS_{i,j,k}$ under the rule $\alpha = \beta$, $A(\alpha)/A(\beta)$. Summarizing: for any $FR_{i,j}$, we defined the sets $CMSS_{i,j,1}$, $CMSS_{i,j,2}$, ...

Lemma 3. Each $CMSS_{i,j,k}$ is the union of the elements of the canonical characterization of some CL-model.

Proof. Consider some $CMSS_{i,j,k}$. By separating its elements according to their form, we obtain a quartuple $\langle S°, I°, E°, N° \rangle$. To see that this is a canonical characterization of some CL-model, remark that CL-models impose no conditions on S°; that N° is empty; and that the only condition on I° and E° is that they should be closed under the semantic consequence relation. But the rule $\alpha = \beta$, $A(\alpha)/A(\beta)$ takes care of this condition. ■

In view of Lemma 3, I shall call the $CMSS_{i,j,k}$ canonical characterizations (of CL-models). Actually, they are *nearly* canonical characterizations: we only have to separate their elements as indicated in the proof.

Lemma 4. For each finished set of requirements $FR_{i,j}$ there is at least one canonical characterization $CMSS_{i,j,k}$.

[14] It is obvious that some subsets are maximal in this respect. This entails that they contain all members of R, that they do not contain any A for which "A is false" is required, and that they contain no formula of the form $\sim A$.

Proof. Let $\langle S°, I°, E°, N° \rangle$ be CC_M and let R be defined from $FR_{i,j}$ as before. R is a subset of $S° \cup I° \cup E° \cup N° \cup R$ that is compatible with some CL-model and meets all the requirements of $FR_{i,j}$. Hence, even if, for any A in $S° \cup I° \cup E° \cup N°$, $R \cup \{A\}$ either is incompatible with some CL-model or conflicts with the requirements in $FR_{i,j}$, there is a $MSS_{i,j,k}$ (viz. R) and hence also a canonical characterization $CMSS_{i,j,k}$. ■

The consistent models associated with Γ

To obtain the consistent models associated with *M*, we turn any canonical characterizations of consistent models associated with Γ into a CL-model which we shall call *typical* for this canonical characterization. The (completely standard) procedure to do so is as follows:

(i) We define $D = \{\alpha \mid \alpha \in C \cup V$ and there is no β such that $\beta = \alpha \in I°$ and $\beta < \alpha\}$.

(ii) If $\alpha \in D$ then $v(\alpha) = \alpha$; if $\alpha \notin D$ and $\alpha = \beta \in I°$, then $v(\alpha) = \beta$.

(iii) $v(\pi^r)$ is the set of all r-tuples $\langle \alpha_1, \ldots, \alpha_r \rangle$ such that $\alpha_1, \ldots, \alpha_r \in D$ and $\pi^r \alpha_1 \ldots \alpha_r \notin E°$.

(iv) Where $A \in S$, $v(A) = 1$ iff $A \in S°$.

(v) for all A, $v(\sim A) = 0$.

I leave it to the reader to show that the resulting model and more specifically that C1.2 is well defined.

Let me quickly recapitulate. We start from the APIL2-models of Γ. For each such model we define an initial set of requirements. Each initial set of requirements is turned into one or more finished sets of requirements. Each of these finished sets leads to one or more canonical characterizations of CL-models. Each of these canonical characterizations may be turned into a CL-model. The thus obtained CL-models are the consistent models associated with Γ.

In section 11, I schematically illustrate the procedure by means of some concrete examples and discuss the way in which we arrive at the selected consistent models associated with Γ. For now, it is important to show that the present procedure leads to the desired results.

Let a *coherent selection from* a set of contradictions $\{A_1 \& \sim A_1, A_2 \& \sim A_2, \ldots\}$ be any set $\{\pm A_1, \pm A_2, \ldots\}$ (in which each "±" is either a negation or nothing) that has a CL-model — it has such a model iff it is not itself inconsistent.

Theorem 5. There is at least one coherent selection from each (non-empty) set of contradictions.

Proof. First, there obviously is a coherent selection from the subset $\{A_1 \& {\sim} A_1\}$ in view of C2.5. Next, if $\{\pm A_1, \ldots, \pm A_n\}$ has a CL-model, say M, then $v_M(A_{n+1}) = 1$ or $v_M({\sim} A_{n+1}) = 1$; consequently, $\{\pm A_1, \ldots, \pm A_n, A_{n+1}\}$ or $\{\pm A_1, \ldots, \pm A_n, {\sim} A_{n+1}\}$ has a CL-model. The Theorem follows from these by induction. ■

As, moreover, consistent PIL-models have the empty set as their initial set of requirements:

Corollary 3. For each PIL-model M, there is at least one initial set of requirements.

Theorem 6. If M is an inconsistent PIL-model, then any coherent selection from the set of contradictions true in M is true in at least one consistent model associated with M.

Proof. Consider a PIL-model M and any coherent selection S from the set of contradictions true in M. In view of the definition of an initial set of requirements, there is an initial set of requirements IR_i that agrees with S. This means: whenever $v_M(A \& {\sim} A) = 1$, "$\pm A$ is false" $\in IR_i$ iff $\pm A \notin S$.

IR_i results in at least one finished set of requirements $FR_{i,j}$ (from Lemma 2). Each $FR_{i,j}$ results in at least one canonical characterization $CMSS_{i,j,k}$ (from Lemmas 3 and 4). Each $CMSS_{i,j,k}$ has a typical CL-model, which is a consistent model associated with M.

In view of the procedure leading to the consistent models associated with M, each such model fulfils all requirements in the $FR_{i,j}$ from which it is built. Indeed, it fulfils all those requirements that pertain to primitive expressions; and it is easily shown, by an induction on the complexity of formulas, that the rules of Table 1 warrant that all requirements in the set are fulfilled if the requirements on the primitive expressions are fulfilled. ■

Before leaving the matter, I add four comments, some clarifying, other expanding. The first is that the consistent models associated with M need not have the same domain as M and not even a domain of the same cardinality. (But their domain is always countable.) Removing an inconsistency may require that two elements of the domain of M are identified, or that new elements are added.(This formulation is sloppy, but clear enough.) Consider, for example some 'rich finitist arithmetic' as described in Jean Paul van Bendegem 199+ (for L fixed). There is an infinity of ways to make such a system consistent, the results being (consistent) finitist arithmetics with L, L+1, ... distinct numbers, and (infinitistic) classical arithmetic. All but the first system require that the domain be extended with 1, 2, ... up to denumerably many elements; each of these solutions is equally acceptable in view of our present only aim: to define the set of consistent models associated with some set of premises.

The second comment is that, for many purposes (especially computational ones) it is not only advisable but also safe to take *only* N-minimal models into account. First remark:

Theorem 7. For any **APIL2**-model M of Γ, there is a formula-equivalent N-minimal **APIL2**-model M' of Γ.

Proof. From Theorem 3 and the fact that, as M and M' are formula-equivalent, the latter is an **APIL2**-model of Γ iff the former is one. ■

Remember that the first step towards the consistent models associated with M was the canonical characterization of M, viz. CC_M. Hence, Theorem 2 and Theorem 7 warrant that the N-minimal models of a set of premises Γ will give us all consistent models associated with Γ.

Before proceeding to the preferences, it is worthwhile to consider what might happen if one were to stepwise eliminate the inconsistencies. People only interested in the reconstruction of mixed non-monotonic logics based on **CL** may skip the sequel of the present section, as it only concerns other (in my view more important) types of applications that presuppose the elimination of some but not all inconsistencies. The most significant applications are those in which one eliminates inconsistencies one at a time, selects the less inconsistent models in view of the preferences, and possibly continues research to fix further preferences.

A central question, so it seems, is whether the deterministic character of the procedure is safeguarded: if we eliminate inconsistencies one (or some) at a time, gradually passing from inconsistent models to less inconsistent models, do we then arrive at the same set of models as defined by the previous procedure? The answer is: under rather obvious conditions, yes indeed. But it turns out that this answer has no *direct* importance for the applications I have in mind. Let me start with the answer itself.

First, in which way will a single inconsistency (or part of the inconsistencies) be eliminated? We start, as above, with initial sets of requirements. Only, we now define these sets with respect to some subset of N°, rather than with respect to N°, itself. We extend this set by the rules of Table 1 to coherent finished sets of requirements. This time, the constraints on the $MSS_{i,j,k}$ are that they are incompatible with the definition of a **PIL**-model (not a **CL**-model), fulfil all requirements in $FR_{i,j}$, and do not lead to inconsistencies that were false in the original model. From there, we go to the $CMSS_{i,j,k}$ and the *less inconsistent* models associated with the premises.

It turns out to be astonishingly simple to prove that the stepwise elimination of the inconsistencies leads ultimately to the same result as the elimination of all of them at once. However, as I announced, this result is not terribly important. If the reader wants concrete examples of the subsequent arguments, I refer him or her to Meheus 199+ (in this volume) on scientific discovery. *First* argument. In view of the preferences, some consistent models may be ruled out that would turn out preferable if all inconsistencies were eliminated at once and if, quite counterfactually, all preferences were fixed at the outset. Nothing can be done about this: that choices are made on the basis of present insights is unavoidable. A further complication is related to the fact that some inconsistencies are interdependent: if an inconsistency is removed in a specific way (relying on the preferences), another inconsistency may be resolved at once. As a consequence, it will never be studied which preferences should be assigned to each of its 'halves'. *Second* argument. In many cases, it is sensible, in view of the preferences, to eliminate an inconsistency in such a way that new inconsistencies arise; sometimes this will be a reason to eliminate the inconsistency in a different way, but sometimes it simply will not. It is possible that new premises (constraints) are added, e.g., because new experiments are performed or some (non-logical) 'principle' is tried out — as Clausius's precursor of the entropy principle. The requirement that no new inconsistencies should arise is nice from a logician's point of view, but might be quite useless (and mistaken) for a practising scientist. The one and only criterion for the scientist are the *preferences*. For him or her, empirical adequacy and the coherence of the theory are far more important than any formal criterion a logician could imagine.

11. *Weeding out inconsistencies: selecting amongst the consistent models*

Let us return to the preferences to define the set of *selected* consistent models associated with Γ. Here the reader may easily verify the significance of the argument presented in section 7: we need to start from APIL2-models of Γ, not from (all) its PIL-models.

The negative result on the elimination of inconsistencies by Rescher's mechanism (section 8) teaches us a very positive lesson on the preferences: the important question is not how many non-monotonic rules hold true in a model, but how many *instances* of such rules are true in it (and which instances — see below).

It is generally said that it should be left to the user which predicates are circumscribed, but this would not work. There is something wrong with that instruction itself, as it may lead to inconsistencies (see, e.g., Łukasiewicz 1990, p. 238). This, however, is hardly a problem for mixed non-monotonic logic. All the latter requires is that suitably employed abnormality predicates are circumscribed in some order. So, here is my claim: if we transfer this order from the abnormality predicates to the preferences of the instances of the corresponding general statements, we obtain an accurate reconstruction on the approach defended here.

Let us first consider circumscription with one abnormality predicate (or with a set of abnormality predicates that do not conflict with each other in the application). (7.1) forms a simple example that will enable me to illustrate the procedure. I repeat the premises

(7.1) $(\forall x)(Bx \supset Fx)$, $(\forall x)(Px \supset Bx)$, $(\forall x)(Px \supset \sim Fx)$, $Pa \& Bb$

as well as the set of atoms contained in the APIL2-consequence set:

(7.3) Pa, Ba, Fa, $\sim Fa$, Bb, Fb, $\sim P\alpha$ (for all $\alpha \in \mathrm{C} \cup \mathrm{V} - \{a\}$)

In the circumscription approach, the first premise of (7.1) would be rephrased as $(\forall x)((Bx \& \sim Xx) \supset Fx)$, where "$X$" is the abnormality predicate. So, the reconstruction requires two steps. *First,* we attach the lowest preference to $(\forall x)(Bx \supset Fx)$ and its instances, and a higher one to the other premises and their instances. As a result, we have two (in other cases more) sets of (open and closed) formulas: Π_1 containing the formulas with the highest preference, and Π_2 containing the formulas with the lower preference. For each model M and each set Π_i, Let $_M\Pi_i = \{A \mid A \in \Pi_i$ and $v_M(A) = 1\}$. *Next,* starting with Π_1, we eliminate from the set of consistent models associated with (some) Γ those M for which there is an M' (in the set) such that $_M\Pi_1 \subset _{M'}\Pi_1$; we repeat this procedure for Π_2, and so on.

Let us return to (7.1). There is one inconsistency in its APIL2-consequence set. This inconsistency holds true in *all* APIL2-models of (7.1), and no other inconsistency holds true in any APIL2-model of (7.1). So, it is quite obvious that we have two sets of consistent models. Still, let me illustrate schematically the procedure from section 10.

Consider a model $M = \langle D, v\rangle$ with $D = \{o_1, o_2, \ldots, o_{20}\}$, $v(a) = o_1$ (and o_1 is assigned to no other individual constant or variable), $v(b) = o_2$, $v(P) = \{o_1\}$, $v(B) = \{o_1, \ldots, o_{10}\}$, $v(F) = \{o_1, \ldots, O_{15}\}$, $v(\sim Fa) = 1$, and $v(\sim A) = 0$ whenever $A \neq Fa$. M is incompletely described, but in all essential respects, M is representative for all APIL2-models of (7.1). For example, if $v(P)$ would contain more elements, or if $v(\sim A)$ would be 1 for more formulas, or if another individual constant were assigned the value o_1, there would be more inconsistencies in the model than required by

(7.1), and hence M would not be an APIL2-model of (7.1). I leave it to the reader to check that all of (7.1) are true in M.

Given that M is incompletely characterized (for the sake of generality), the description of its canonical characterization is somewhat complex:

CC_M
- S° and I° are fully unspecified
- E° contains Pa, Ba, Fa, Bb, and Fb, does not contain $P\alpha$ for any $\alpha \in C \cup V - \{a\}$, and is unspecified as far as other atoms are concerned
- $N^\circ = \{\sim Fa\}$

The finished sets of requirements are $\{Fa$ is false$\}$ and $\{\sim Fa$ false, Fa is true$\}$. This leads to

$CMSS_{1,1,1}$ (separated into the elements of a canonical characterization)
- S° and I° as in CC_M
- E° contains Pa, Ba, Bb, and Fb, does not contain Fa and does not contain $P\alpha$ for any $\alpha \in C \cup V - \{a\}$, and is identical to CC_M as far as other primitive predicative expressions are concerned
- $N^\circ = \varnothing$

$CMSS_{2,1,1}$ (separated into the elements of a canonical characterization)
- S° and I° as in CC_M
- E° contains Pa, Ba, Fa, Bb, and Fb, does not contain $P\alpha$ for any $\alpha \in C \cup V - \{a\}$, and is identical to CC_M as far as other primitive predicative expressions are concerned
- $N^\circ = \varnothing$

In the corresponding CL-models, say $M_{1,1,1}$ and $M_{2,1,1}$, we have (I mention only a relevant selection):

$M_{1,1,1}$

True: $Pa,\ Ba,\ Bb,\ Fb,\ (\forall x)(Px \supset Bx),\ (\forall x)(Px \supset \sim Fx),$
$B\alpha \supset F\alpha$ for all $\alpha \in C \cup V - \{a\}$

False: $Fa,\ P\alpha$ for all $\alpha \in C \cup V - \{a\} - \{a\},\ Ba \supset Fa$

$M_{2,1,1}$

True: $Pa,\ Ba,\ Fa,\ Bb,\ Fb,\ (\forall x)(Bx \supset Fx),\ (\forall x)(Px \supset Bx),$
$P\alpha \supset \sim F\alpha$ for all $\alpha \in C \cup V - \{a\}$

False: $Pa \supset \sim Fa,\ P\alpha$ for all $\alpha \in C \cup V - \{a\}$

In other words, $M_{1,1,1}$ is just like M except for $v(F)=\{o_2, ..., o_{15}\}$ and $M_{2,1,1}$ is just like M except that $v(\sim A)=0$ for all A. In view of the preferences, $M_{1,1,1}$ is the only selected consistent model associated with (7.1) — it contains all instances of $(\forall x)(Px \supset\sim Fx)$ whereas $M_{2,1,1}$ does not. In other words, a comes out a non-flyer, as it should.

Would this work equally well for other APIL2-models of (7.1)? The only possible difference with the (incompletely described) model M is that an alternative model would map some members of V on $o_1 = v(a)$.[15] In this case, the procedure leads to the somewhat unexpected result that these members of V will be mapped on an element of the domain that is different from $v(a)$. This, however, has not the slightest effect on the truth of any *wff* in the selected models, and hence the wffs true in all selected consistent models associated with (7.1) will be the CL-consequences of $\{Pa, Ba, \sim Fa, Bb, Fb, (\forall x)(Px \supset Bx), (\forall x)(Px \supset\sim Fx), (\forall x)(x \neq a \supset(Bx \supset Fx))\}$ — $(\forall x)(Px \supset x = a)$ is among these.

It is easy to see that this result is, the formulation left aside, identical to the one we obtain by circumscribing X. Only, the selection of models proceeds in a different way. The first premise will be formulated as $(\forall x)((Bx \& \sim Xx) \supset Fx)$. All premises will be true in all models (so nothing corresponds to the second set above), but models will differ with respect to the number of abnormal entities. If we minimize the abnormal entities by circumscription, we obtain exactly the set of models we selected above (with $v(X)=v(P)$).

Let us now consider circumscription with two interfering abnormality predicates, say X_1 and X_2. I first write the extended set of premises with the abnormality predicates:

(11.1) $(\forall x)((Bx \& \sim X_1x) \supset Fx)$, $(\forall x)(Px \supset Bx)$, $(\forall x)((Px \& \sim X_2x) \supset\sim Fx)$, Pa, Bb, Pc, Fc

Here we allow "Penguins don't fly" to have exceptions, and fortunately so, as c is a flying penguin. All penguins now are abnormal: non-flying penguins are X_1-abnormal (abnormal as birds with respect to flying), whereas the flying ones are X_2-abnormal (abnormal as penguins with respect to flying). A suitable circumscription will require that we first minimize X_2-abnormality, and only then minimize X_1-abnormality.

According to the reconstruction, the premises are:

15 Where M' is such a model, there will be open formulas A such that $v_{M'}(A \& \sim A) = 1$ and $v_M(A \& \sim A) = 0$. However, for all open and closed formulas A, $v_{M'}(\exists(A \& \sim A)) = v_M(\exists(A \& \sim A))$, and hence M' is an APIL2-model of (7.1).

(11.2) $(\forall x)(Bx \supset Fx)$, $(\forall x)(Px \supset Bx)$, $(\forall x)(Px \supset \sim Fx)$, Pa, Bb, Pc, Fc

The preferences will obviously be assigned thus:

$P1$ Pa, Bb, Pc, Fc and $(\forall x)(Px \supset Bx)$
$P2$ instances of $(\forall x)(Px \supset \sim Fx)$
$P3$ instances of $(\forall x)(Bx \supset Fx)$

with $P1 > P2 > P3$.
The **APIL2**-consequence set contains the following atoms:

(11.3) Pa, Ba, Fa, $\sim Fa$, $\sim Pb$, Bb, Fb, Pc, Bc, Fc, $\sim Fc$, $\sim P\alpha$ (for all $\alpha \in \mathrm{C} \cup \mathrm{V} - \{a\}$.

As before, inconsistencies will be minimal in **APIL2**-models and we only have to eliminate the two inconsistencies that occur in (11.3). This will give us four (sets of) models associated with an **APIL2**-model of (11.2), for which I list only the atoms and the significant (instances of) generalizations:

M_1
True: Pa, Ba, Fa, Bb, Fb, Pc, Bc, $(\forall x)(Px \supset Bx)$, $P\alpha \supset \sim F\alpha$ for all $\alpha \in \mathrm{C} \cup \mathrm{V} - \{a\}$, $B\alpha \supset F\alpha$ for all $\alpha \in \mathrm{C} \cup \mathrm{V} - \{c\}$
False: $P\alpha$ for all $\alpha \in \mathrm{C} \cup \mathrm{V} - \{a, c\}$, Fc, $Pa \supset \sim Fa$, $Bc \supset Fc$

M_2
True: Pa, Ba, Bb, Fb, Pc, Bc, $(\forall x)(Px \supset Bx)$, $(\forall x)(Px \supset \sim Fx)$, $B\alpha \supset F\alpha$ for all $\alpha \in \mathrm{C} \cup \mathrm{V} - \{a, c\}$
False: $P\alpha$ for all $\alpha \in \mathrm{C} \cup \mathrm{V} - \{a, c\}$, Fa, Fc, $Bc \supset Fc$, $Ba \supset Fa$

M_3
True: Pa, Ba, Fa, Bb, Fb, Pc, Bc, Fc, $(\forall x)(Px \supset Bx)$, $P\alpha \supset \sim F\alpha$ for all $\alpha \in \mathrm{C} \cup \mathrm{V} - \{a, c\}$, $(\forall x)(Bx \supset Fx)$
False: $P\alpha$ for all $\alpha \in \mathrm{C} \cup \mathrm{V} - \{a, c\}$, $Pa \supset \sim Fa$, $Pc \supset \sim Fc$

M_4
True: Pa, Ba, Bb, Fb, Pc, Bc, Fc, $(\forall x)(Px \supset Bx)$, $P\alpha \supset \sim F\alpha$ for all $\alpha \in \mathrm{C} \cup \mathrm{V} - \{c\}$, $B\alpha \supset F\alpha$ for all $\alpha \in \mathrm{C} \cup \mathrm{V} - \{a\}$

False: $P\alpha$ for all $\alpha \in C \cup V - \{a, c\}$, Fa, $Pc \supset \sim Fc$, $Ba \supset Fa$

First, M_1 and M_2 are eliminated because a wff with the highest preference is false in them, viz. Fc. When we come to wffs with preference $P2$, M_3 is eliminated in favour of M_4 because $\{Pa \supset \sim Fa\} \subset \{Pa \supset \sim Fa, Pc \supset \sim Fc\}$.

Remark that the above argument is representative for all consistent models associated with (11.2); for any **APIL2**-model of (11.2), we obtain a set of models that correspond to M_1 - M_4 and the one corresponding to M_4 will be selected. The wffs true in all selected consistent models associated with (11.2) are the **CL**-consequences of $\{Pa, Ba, \sim Fa, Bb, Fb, Pc, Bc, Fc, (\forall x)(Px \supset Bx), (\forall x)(x \neq c \supset (Px \supset \sim Fx)), (\forall x)(x \neq a \supset (Bx \supset Fx))\}$ — remark that $(\forall x)(Px \supset (x = a \vee x = c))$ is a **CL**-consequence of this set.

If the last example is reformulated with one abnormality predicate only, circumscription does not enable us to derive either Fa or $\sim Fa$. The same result obtains under the present reconstruction. Remark that we have to posit $P2 = P3$. In this case, the set of *selected* consistent models associated with (11.2) contains M_3 as well as M_4. Fc is true in each of them, but neither Fa nor $\sim Fa$ is.

A proof that the reconstruction is accurate in general has to be postponed to another paper. Still, I hope it will be quite clear to logicians that followed the argument that such proof is within reach (and that the reconstruction of some other mixed non-monotonic logics is within reach as well). Let me adduce some evidence for this claim .

Let the premises Γ be given as the union of n+1 non-overlapping sets $\Gamma = \Gamma_0 \cup \Gamma_1 \cup \ldots \cup \Gamma_n$. The members of Γ_0 do not contain any abnormality predicates. The members of the other n sets contain wffs of the form $(\forall \alpha)((A \& \sim X_{\ldots} \alpha) \supset B)$.[16] Γ_1 contains the wffs in which occur the abnormality predicates that are first circumscribed, say $X_{1,1}$, $X_{1,2}$, etc. ; Γ_2 the wffs in which occur the abnormality predicates that are next circumscribed, say $X_{2,1}$, $X_{2,2}$, etc.; and so on.[17] For the reconstruction, we start from $\Gamma_0 \cup \Gamma_{1*} \cup \ldots \cup \Gamma_{n*}$ where Γ_{i*} contains $(\forall \alpha)(A \supset B)$ for each $(\forall \alpha)((A \& \sim X_j \alpha) \supset B) \in \Gamma_{i*}$. The members of Γ_0 obtain the highest preference $P0$, those of Γ_{1*} the next highest preference $P1$, etc.

[16] Whether α does or does not occur free in A or B is immaterial. This restriction on Γ_i is justified in view of the application of circumscription as a mixed non-monotonic logic. Circumscription itself is a much more general mechanism, and a rather wild one, viz. a mechanism with rather unpredictable effects.

[17] I disregard formulas in which occur several abnormality predicates.

Circumscription considers **CL**-models only. Let Δ be the set of wffs true in all models of Γ in which the abnormality predicates have been circumscribed in the correct order. Let $\Delta_0 = Cn_{\mathrm{CL}}(\Gamma_0)$ obviously $\Delta_0 \subseteq \Delta$. Let $\Xi_i \vee A$ abbreviate any formula of the form $(X_{i,j+1}\alpha_{j+1} \vee \ldots \vee X_{i,j+n}\ \alpha_{j+n}) \vee A (n \geq 1)$. $A \in \Delta_{i+1}$ iff there is Ξ_i such that $\Xi_i \vee A \in Cn_{\mathrm{CL}}(\Gamma_0 \cup \ldots \cup \Gamma_{i+1})$ and there is no Ξ_i such that $\Xi_i \vee {\sim} A \in Cn_{\mathrm{CL}}(\Gamma_0 \cup \ldots \cup \Gamma_{i+1})$. Let $\Delta = \Delta_n$. Let me somewhat explain the recursive definition. First remark that each Δ_i is consistent and that $\Delta_0 \subseteq \Delta_1 \subseteq \ldots \subseteq \Delta_n$. Next, suppose that $\Xi_{i+1} \vee A \in Cn_{\mathrm{CL}}(\Gamma_0 \cup \ldots \cup \Gamma_{i+1})$. If $A \notin \Delta_0 \cup \ldots \cup \Delta_i$ and ${\sim} A \in \Delta_0 \cup \ldots \cup \Delta_i$ then there is a Ξ_i such that $\Xi_i \vee {\sim} A \in Cn_{\mathrm{CL}}(\Gamma_0 \cup \ldots \cup \Gamma_{i+1})$ and hence $A \notin \Delta_{i+1}$ and ${\sim} A \in \Delta_{i+1}$. If there is no Ξ_{i+1} such that $\Xi_{i+1} \vee {\sim} A \in Cn_{\mathrm{CL}}(\Gamma_0 \cup \ldots \cup \Gamma_{i+1})$ — hence ${\sim} A \notin \Delta_0 \cup \ldots \cup \Delta_i$ — then $A \in \Delta_{i+1}$ and ${\sim} A \notin \Delta_{i+1}$. If A, ${\sim} A \notin \Delta_0 \cup \ldots \cup \Delta_i$ and there is Ξ_{i+1} such that $\Xi_{i+1} \vee {\sim} A \in Cn_{\mathrm{CL}}(\Gamma_0 \cup \ldots \cup \Gamma_{i+1})$ then A, ${\sim} A \notin \Delta_{i+1}$.

The present reconstruction starts from **APIL2**-models. The **APIL2**- consequences of $\Gamma_0 \cup \Gamma_{1*} \cup \Gamma_{2*} \cup \ldots \cup \Gamma_{n*}$ may be characterized in two steps. *First* we define, with respect to Γ (the formulas in which the abnormality predicates occur), a sequence of sets $\Theta_0, \Theta_1, \ldots, \Theta_n$ as follows: $\Theta_0 = Cn_{\mathrm{PIL}}(\Gamma_0)$ and $A \in \Theta_{i+1}$ iff there is a Ξ_i such that $\Xi_i \vee A \in Cn_{\mathrm{PIL}}(\Gamma_0 \cup \ldots \cup \Gamma_{i+1})$. Again $\Theta_0 \subseteq \Theta_1 \subseteq \ldots \subseteq \Theta_n$. *Next* we define Θ as $Cn_{\mathrm{APIL2}}(\Theta_n)$. By this closure under **APIL2**, we rule out any further inconsistencies: if $A \in \Theta$ and ${\sim} A \notin \Theta$, then ${\sim} A \supset B \in \Theta$ for all B, and hence $v_M(A) = 0$ for all **PIL**-models of Θ (except for the trivial model, which is not an **APIL2**-model of $\Gamma_0 \cup \Gamma_{1*} \cup \Gamma_{2*} \cup \ldots \cup \Gamma_{n*}$ anyway). Incidentally, the non-trivial **PIL**-models of Θ *are* the **APIL2**-models of $\Gamma_0 \cup \Gamma_{1*} \cup \Gamma_{2*} \cup \ldots \cup \Gamma_{n*}$. Defining the consistent models associated with $\Gamma_0 \cup \Gamma_{1*} \cup \Gamma_{2*} \cup \ldots \cup \Gamma_{n*}$ from the **APIL2**-models of $\Gamma_0 \cup \Gamma_{1*} \cup \Gamma_{2*} \cup \ldots \cup \Gamma_{n*}$, we obtain **CL**-models that maximally agree with some **APIL2**-model of $\Gamma_0 \cup \Gamma_{1*} \cup \Gamma_{2*} \cup \ldots \cup \Gamma_{n*}$ and moreover contain a coherent selection from the inconsistencies in Θ. The subsequent selection of consistent models associated with $\Gamma_0 \cup \Gamma_{1*} \cup \Gamma_{2*} \cup \ldots \cup \Gamma_{n*}$ warrants that (i) all members of Γ_0 are true in all selected consistent models, and (ii) if ${\sim} A \notin Cn_{\mathrm{CL}}(\Theta_{i+1})$ — hence ${\sim} A \notin Cn_{\mathrm{CL}}(\Theta_i)$ — then $A \in (\Theta_{i+1})$ warrants that A is true in all selected consistent models. And as all selected consistent models are **CL**-models, all **CL**-consequences of the wffs true in each of them are true in each of them.

The preceding paragraphs do, as announced, not constitute a proof. Yet, they make it plausible that what is lost by allowing for inconsistencies, is regained by eliminating them according to the preferences.

12. *Concluding remarks and open problems*

What does my proposal come to? As I promised, rules with exceptions are seen as rules that may lead to inconsistencies for which we have a fixed elimination procedure. Given an inconsistent set of premises, paraconsistent logics such as **PIL** lead to an inconsistent but non-trivial consequence set. By moving to **APIL2**, we minimize the inconsistencies with respect to the premises, thus arriving at a richer consequence set than if we apply **PIL**. This is the first selecting move, the purely logical one. It actually is a very important step, and I return on it below.

Our next step was to define, from the set of inconsistent **APIL2**-models of the premises, a set of consistent models. This step too was completely determined on logical grounds; preferences do not play any role in it. Not only the deductive aspect, but also the definition of the consistent models associated with the premises from their **APIL2**-models, is independent *both* of the preferences and of the way in which these are handled: the way in which preferences of sentences are turned into a selection of the consistent models. *Only this very last aspect, settling the preferences and selecting amongst the consistent models in terms of them, is what the discussion between mixed non-monotonic logicians should be about.*

To see the importance of the first and second step, reconsider circumscription for a moment. Circumscription minimizes (the occurrence of) abnormality predicates in a certain order. The minimization of the abnormality predicates itself corresponds to a purely logical step in the reconstruction: the restriction to **APIL2**-models minimizes inconsistencies. The *order* in which abnormality predicates are minimized corresponds to the step where preferences come in: the selection of the (associated) consistent models. The last paragraph of the previous section shows the way in which **APIL2**-models take care of minimizing abnormality predicates, and the role of preferences to determine that order.

The procedure I sketched was somewhat complicated. From a philosophical point of view, this is quite unimportant. The gain is that the two components of mixed non-monotonic logics are separated. In other words, that rules with exceptions are shown to be special cases of facing inconsistencies in our belief system. What is special about them is that we have a ready way out: the preferences are known.

As for the practical side of the enterprise, I hope that the examples show the procedure to be rather expeditious in concrete cases. So, even before tableau-methods and similar tools are articulated, the approach allows one to tackle the ordering of the preferences on statements and the way to derive from them a selection of the consistent models. The solution of this difficulty (the real task for students of mixed non-monotonic logics) is by no means trivial. Anyone who looked at mixed non-monotonic logics from a formal point of view, rather than from a couple of examples, knows that it is easy enough to create problems by the dozen. Consider, for example, the assignment of preferences in terms of the generality of the rules. It is by no means clear whether the notion of generality should be understood extensionally or intensionally, and messing up the rules by means of some elementary logical transformations ruins the whole approach. So, from a practical point of view, the reconstruction presented in this paper offers the possibility to study the preference ordering within a transparent formal framework. This would not only lead to an ordering of the field, but would also enable one to decide which approaches are specific cases of others, and which are simply inadequate as tools of some generality.

It is extremely important that the present reconstruction does *not* require that preferences be attached to premises. If we use some general criterion — such as the generality of the rule — rather than a given set of preferences, we may apply the criterion directly to APIL2-consequences of the premises rather than to the premises themselves. This is why we are able to reconstruct circumscription by the present approach (as opposed to a Rescher-like approach), for circumscription operates on the abnormality predicates, not on premises. With respect to the use of a general criterion for determining the preferences, the present approach is even more promising than circumscription. The same abnormality predicates may occur in general statements of different generality, whereas the present approach does not use any abnormality predicates and hence a general criterion may be applied to statements as they are derived.[18]

A rather different advantage of the present approach is that it opens up a number of perspectives. Recently, non-monotonic logicians finally were sensible enough to start considering cases where not all inconsistencies are eliminated. Within the present framework, this is easy to handle. As the order in which inconsistencies are eliminated does not matter to the procedure sketched above, not eliminating some inconsistencies is fully unproblematic.

[18] This does not entail that the present reconstruction of circumscription-based mixed non-monotonic logic might be inadequate. The reconstruction presupposes indeed that the order in which the abnormality predicates are circumscribed is given. But the present approach will be more friendly towards a general preference criterion than circumscription.

Even more remote from the mixed non-monotonic enterprise, but presumably more important for the philosophy of science, is the possibility to first analyze, in terms of **APIL2**, the inconsistencies involved, and then to fix preferences in view of these results, possibly after gaining more information — and the analysis of the inconsistencies will suggest which information to look for. In this respect I already referred to Meheus 199+. That paper is directed to understanding creative processes. Actually, one such creative process concerns the very development of non-monotonic logics itself. The forerunners in the field, viz. those who approached the matter computationally, had to start analyzing the very simple examples, still well-know today, to find out how they might be dealt with.

To conclude this paper, I mention some open problems. First and foremost, it is quite apparent that the present approach demands and deserves some more work at the meta-theoretic level. To make the reconstruction more practicable, tableau-methods should be worked out and the connected metatheorems proved. And, most importantly, general proofs should be presented to show that the reconstruction of a specific mixed non-monotonic logic is adequate.

Variations on the present enterprise suggest themselves. The first concerns the use of adaptive logics based on other substructural logics. A host of those logics is available. Although I am pretty convinced that **APIL2** is an excellent choice for the present purpose,[19] a logical study based on a single system is manifestly too restricted. Not just other logics, but other types of logics suggest themselves in this respect — see, e.g., Joachim Van Meirvenne 199+, who considers indexed adaptive logics. Such variations would clearly lead to new mixed non-monotonic logics, but they would do so independently of the discussion on the way in which to deal with preferences. So, the field will be extended further, but the discussion of the alternatives will be much more transparent than it is today, thanks to the fact that the present reconstruction separates the different aspects. If such problems were tackled, the study of mixed non-monotonic logics would change rather dramatically and —as I hope to have convinced the reader — quite for the better.

Universiteit Gent

[19] For example, many paraconsistent logics validate Contraposition (and hence Modus Ponens). But it is easily seen that, in the present context, each of these merely involves complications, viz. spreading inconsistencies.

REFERENCES

Batens, Diderik

1980 "Paraconsistent extensional propositional logics." *Logique et Analyse* 90-91, pp. 195-234.

1986a "Dialectical dynamics within formal logics." *Logique et Analyse* 114, pp. 161-173.

1986b "Rich relevant logics based on a simple type of semantics." *Logique et Analyse* 116, pp. 437-457.

1989 "Dynamic dialectical logics." in G. Priest, R. Routley & J. Norman (eds.), *Paraconsistent Logic*. München, Philosophia Verlag, pp. 187-217.

199+ "Inconsistency-Adaptive Logics." in Ewa Orłowska (ed.) *Logic at Work. Essays dedicated to the memory of Helena Rasiowa.* Kluwer (in print).

Łukasiewicz, Witold

1990 *Non-Monotonic Reasoning. Formalization of Commonsense Reasoning*. New York, Ellis Horwood.

Meheus, Joke

1993 "Adaptive logic and discovery: the case of Clausius." *Logique et Analyse*, 143-144, pp. 360-391.

Priest, Graham

1987 *In Contradiction. A Study of the Transconsistent.* Dordrecht, Nijhoff.

1991 "Minimally inconsistent LP." *Studia Logica* 50, pp. 321-331.

Rescher, Nicholas

1964 *Hypothetical Reasoning.* Amsterdam, North-Holland.

1973 *The Coherence Theory of Truth.* Oxford,Clarendon.

Rescher, Nicholas & Ruth Manor

1970 "On inference from inconsistent premises." *Theory and Decision* 1, pp. 179-217.

Rescher, Nicholas & Robert Brandom

1980 *The Logic of Inconsistency. A study in Non-Standard Possible-World Semantics and Ontology.* Oxford, Blackwell.

Shoham, Y.
1988 *Reasoning about Change: Time and Causation from the Standpoint of Artificial Intelligence.* Cambridge, Mass., MIT Press.

Van Bendegem, Jean Paul
1994 "Strict Finitism as a viable alternative in the foundations of mathematics." *Logique et Analyse*, this volume.

Van Meirvenne, Joachim
1994 "Indexed inconsistency-adaptive logic." *Logique et Analyse*, this volume.