AUTOMATED REASONING PROJECT

TECHNICAL REPORT TR-ARP-3/88 7 March, 1988

CONSISTENCY BY DEFAULT

Graham Priest

TECHNICAL REPORT SERIES

Research School of Social Sciences
Australian National University

TECHNICAL REPORT TR-ARP-3/88 7 March, 1988

CONSISTENCY BY DEFAULT

Graham Priest

University of Western Australia

CONSISTENCY BY DEFAULT

1 Introduction: Consistency by Default

*Paraconsistent logics provide the vehicle for non-trivial reasoning in inconsistent situations. Such are important for many reasons. (See, for example, Priest and Routley [1983], ch 4, ch 14 of reprint.) I shall not survey them here. However, an important theme in the following investigations is the mechanisation of inference procedures. Hence, let me give one example which is pertinent to that. Automated reasoning systems may well end up in inconsistency. This may be so since the information on which our theorem prover works is inconsistent (see Belnap [1977]); or it may be so due to the very nature of the inferential machinery. Thus, the reasoning essential to any cognitive AI system is itself wont to produce contradictions and paradoxes (see, e.g., Asher and Kamp [1986], Priest [198+]). Since consistency is undecidable, theorem provers must learn to live with it; thus, they must use a paraconsistent logic.

Despite the fact that inconsistency is unavoidable, we have a right to expect consistency most of the time, since consistency is the norm. I shall not defend this position here. (It is defended in detail in Priest [1987], 8.4.) Let me just point out that widespread inconsistencies in rules, codes of law, and our thought about the cognitive states of ourselves and others, would make these entire "language games" a mockery. Hence consistency is a natural default assumption: the situation is assumed to be innocent of contradiction until proven guilty.

Another claim that informs the following investigations is that in consistent situations, classical logic is correct. I should say immediately that I make this claim only with respect to the extensional connectives $\wedge$, $\vee$, and $\neg$, and the quantifiers $\forall$ and $\exists$. Even so, the claim is not uncontentious. I shall not defend it at great length here. (It is defended in Priest [1987], esp ch 8.) Perhaps the major challenge to it comes from the following thought. Just as there are inconsistent situations that pose a threat to classical principles, so there are incomplete situations, where we have neither of α and $\neg\alpha$, which equally pose a challenge. The only point that I shall make about this here is that it is not at all clear that such situations, if and when they exist, pose a challenge to classical logic. For example, suppose that a data base is incomplete since it has been told neither α nor $\neg\alpha$. We would not, for that reason, want to be informed that $\alpha\vee\neg\alpha$ is not true; presumably it is: our information is just incomplete.

There are more radical challenges to the claims of the previous two paragraphs. However, since this is not a philosophical paper, I shall not go into them here. At least I have laid my assumptions on the table. And given these two claims, that consistency is a correct default assumption, and that classical logic is correct in consistent situations, it follows that:

Maxim M

Classical inferences should be allowed until and unless inconsistency interferes with them.

The content of Maxim M is not completely transparent. However, it obviously implies that an adequate inference engine should allow us to infer all the classical consequences of a (classically) consistent set of sentences. Rather more follows, however. Granted that q should follow from {p, ¬pvq}, it should equally follow from {p, ¬pvq, r∧¬r} since the contradiction r∧¬r has nothing to do with the matter. This is clear enough in this case, but what does 'until and unless inconsistency interferes with the inference' mean in general? To give a precise answer to this is the point of this paper.

To answer the question we need to fix, first, on a system of paraconsistent logic to be used. In what follows I shall work with the logic LP. (See Priest [1980], appendix, or Priest [1987], ch 5.) I think that this is the correct paraconsistent logic for the extensional connectives, though, again, I shall not defend this here. It is certainly simple, and almost all of what follows can be generalised to an arbitrary paraconsistent logic provided only that its extensional theorems coincide with classical theorems. This is true of most relevant paraconsistent logics (for a survey of which see Routley *et al* [1982]); and provides enough justification for using LP in these investigations.

In what follows, some notation and one particular result concerning LP will be useful. The notation is as follows: I will use α, β, γ, ... for formulas; p, q, r, ... for atomic formulas; Σ, Γ, ... for sets of formulas. α! will mean α∧¬α. (! will always have narrow scope.) Apart from LP, I will refer to four other logics: classical logic, CL, and three extensions of LP: LP^+, LP^* and LP^m. These will get abbreviated to +, * and m respectively. If X is any of the above systems $\models^X$ will denote its consequence relation and Σ^X will be the set of consequences of the set Σ under $\models^X$. The theorem is:

Theorem 0

$\Sigma \models^{CL} \alpha$ iff $\Sigma \models^{LP} \alpha \vee \beta!$

Proof

See Priest [1987], p 149.

□

It follows from this result that only a single rule of inference need be added to LP to get classical logic: $\alpha \vee \beta!/\alpha$. I will call this rule 'contradiction suppression', CS (Routley *et al* [1982], ch 2, sec 10). Thus, one way of formulating our problem is this: under what conditions may the rule CS be used?

2 Standard Approaches to Non-Monotonic Logic: LP^+

Observe, first, in case it is not already obvious, that however we solve the problem a non-monotonic logic will result. For p should follow from {pvq!} but not from {pvq!, q!}. Artificial intelligence has studied non-monotonic logics which implement various kinds of default assumptions. By and large, the techniques used there are not particularly appropriate in the present case. McCarthy's technique of circumscription ([1980]) is geared to minimising the number of entities of a certain kind, rather than the number of formulas of a certain kind (viz. contradictions). Another way of implementing default reasoning used by McDermott and Doyle (Mcdermott and Doyle [1980], Mcdermott [1982]) employs a modal possibility operator, M, and rules of the form: $\alpha, M\beta/\gamma$ (γ is inferred from α and the fact that β is possible - in a certain model). This suggests the implementation of CS as a default rule: $\alpha \vee \beta!, M\neg(\beta!)/\alpha$. However, this will not work since $\neg(\beta!)$ is a logical truth, as, therefore is $M\neg(\beta!)$. Classical logic would therefore be the result, and inconsistent theories would collapse into triviality.

Another approach to default reasoning pursued by Reiter ([1980]) gives results similar to the previous approach in the standard case, but gives rather different results in the present one. This approach implements rules of the form: $\vdash\alpha$, $\nvdash\beta$ / $\vdash\gamma$. In this context, the rule CS would become:

$$\vdash\alpha\vee\beta!,\ \nvdash\beta!\ /\ \vdash\alpha \qquad (1)$$

Of course, this rule makes little sense proof-theoretically. However, we can make semantic sense of it by defining inferential extensions (augmentations) of a set which, in a sense, satisfy (1), as follows.

Definition 0

Γ is an *augmentation* of Σ iff:

i) $\Gamma \supseteq \Sigma$
ii) Γ is closed under $\models^{LP}$
iii) if $\alpha\vee\beta! \in \Gamma$ and $\beta! \notin \Gamma$ then $\alpha \in \Gamma$

Each augmentation of Σ specifies a "coherent" set of default consequences of Σ. Notice that such extensions always exist. The set of all formulas is such a set (and is also a perfectly good paraconsistent theory). Unfortunately, however, augmentations are by no means unique. Let Σ be any classically consistent set of sentences; then any classical theory containing Σ satisfies the above conditions, and of these there will, in general, be many. Even if Σ is inconsistent there will, in general, be many augmentations. This follows from the following theorem.

Theorem 1

Let Σ be any set of formulas such that $\Sigma \not\models^{LP} \alpha$ then there is an augmentation, Γ, of Σ such that $\Gamma \not\models^{LP} \alpha$ (and hence $\alpha \notin \Gamma$).

Proof

Given such a Σ we can extend it to a prime, deductively closed set Γ, such that $\Gamma\nvDash^{LP}\alpha$. (See, e.g. Priest [1980], 6.2, or Routley *et al* [1982], p 307 f.) By these properties Γ is an augmentation of Σ.

□

Now, let $\Sigma=\{p!\vee q!\}$. Let $\Sigma_1 = \{p!\vee q!,\ p!\}$; $\Sigma_2 = \{p!\vee q!,\ q!\}$. Note that $\Sigma_1\nvDash^{LP}q!$ and $\Sigma_2\nvDash^{LP}p!$. By the above theorem, there are augmentations of Σ_1 and Σ_2, Γ_1 and Γ_2, respectively, such that $\Gamma_1\nvDash^{LP}q!$ and $\Gamma_2\nvDash^{LP}p!$ Then Γ_1 and Γ_2 are distinct augmentations of Σ.

Given a set of sentences, it might make some sense to choose an augmentation indifferently. This, for example, is what Doyle's Truth Maintenance System does. (See McDermott [1982], pp 43-4.) However, from a logical point of view, this is obviously arbitrary. The only suitable candidate for the set of absolute default consequence of Σ is the intersection of all augmentations. Thus we have the following:

Definition 1

$\Sigma\vDash^{+}\alpha$ iff $\alpha\in\bigcap\{\Gamma;\ \Gamma$ is an augmentation of $\Sigma\}$

Σ^{+} is the natural candidate for the set of default consequences of Σ. Unfortunately it is no good, since it does no better than LP itself. In fact, Σ^{+} is exactly Σ^{LP}. The inclusion from right to left is obvious. The converse follows immediately from theorem 1. Thus $\vDash^{+}$ violates Maxim M.

3 *Consequence: LP*

If one considers the default rule (1) above, it might strike one simply to define the default consequences of Σ simply as the set of all those α such that for some β $\Sigma\vDash^{LP}\alpha\vee\beta!$, but $\Sigma\nvDash^{LP}\beta!$. This, however, is far too liberal, and will not do. For let $\Sigma=\{p!\}$. Then, as may easily be checked, $\Sigma\vDash^{LP}q\vee(\neg q\wedge p)!$, but $\Sigma\nvDash^{LP}(\neg q\wedge p)!$ Hence, an arbitrary q would be a consequence. This is not quite triviality (since it is only an arbitrary atomic formula that follows), but it is obviously near enough to render the definition of no value.

Is there, perhaps, some way of modifying this idea to make it work? There is. The matter is discussed in Priest [1987], and, to cut a long story short, the following account does the job. Let us define **consequence* as follows:

Definition 2

$\Sigma\vDash^{*}\alpha$ iff (i) $\Sigma\vDash^{LP}\alpha$ or
(ii) for some β $\Sigma\vDash^{LP}\alpha\vee\beta!$,
where $\Sigma\nvDash^{LP}\beta!$ and $\Sigma\nvDash^{LP}(\alpha\vee\beta!)(\alpha/\neg\alpha)$

where (γ/δ) denotes substitution of δ for γ. (The first disjunct is necessary since it does not imply the second.) What makes *consequence an appropriate notion of default consequence are the following facts.

Theorem 2

i) $\Sigma^{LP} \subseteq \Sigma^* \subseteq \Sigma^{CL}$.
ii) In general, these containments are proper. (Thus, for example, $\{p, \neg p \vee q\} \models^* q$ and $\{p!\} \not\models^* q!$)
iii) But if Σ is (classically) consistent $\Sigma^* = \Sigma^{CL}$.
iv) Moreover, $\alpha! \in \Sigma^*$ iff $\alpha! \in \Sigma^{LP}$.
v) Thus, Σ^* is trivial iff Σ^{LP} is trivial.

Proof

See Priest [1987], 8.6.

□

These facts show that $\models^*$ satisfies Maxim M. But despite these pleasing properties it has some awkward features. To start to see these, we need first a lemma and a theorem. Both of these concern only propositional logic.

Lemma

If there is a β which satisfies the second disjunct of the *definiens* of definition 2, it needs no more than 9 propositional parameters in addition to those occurring in $\Sigma \cup \{\alpha\}$. (These can, of course, be the first three on some arbitrary list.)

Proof

To see this, take any β which will do the job. We know that $\Sigma \models^{LP} \alpha \vee \beta!$, $\Sigma \not\models^{LP} \beta!$ and $\Sigma \not\models^{LP} \alpha \vee \beta!(\alpha/\neg\alpha)$. Let ν and μ be evaluations of the propositional parameters which establish the last two facts. Consider the propositional parameters that occur in β but not in $\Sigma \cup \{\alpha\}$. Define an equivalence relation, $\sim$, on these by: $p \sim q$ iff $\nu(p) = \nu(q)$ and $\mu(p) = \mu(q)$. Since an evaluation can take only 3 possible values, there are 9 possible equivalence classes. Choose one variable from each class, and substitute it for all the others in that class in β. Call this formula β'. Since β' is a substitution instance of β, which has not affected any of the variables in $\Sigma \cup \{\alpha\}$, $\Sigma \models^{LP} \alpha \vee \beta'!$. But, by construction, ν and μ refute $\Sigma \models^{LP} \beta!$ and $\Sigma \models^{LP} (\alpha \vee \beta!)(\alpha/\neg\alpha)$ respectively. Thus, β' will do the job.

□

We can now prove:

Theorem 3

$\models^*$ is decidable.

Proof

We show this by giving an algorithm to decide whether $\Sigma \models^* \alpha$ where Σ is finite. To check this, first check $\Sigma \models^{LP} \alpha$. (Since LP is decidable, this is effective.) If not, proceed. Next, we search for a β which witnesses the other disjunct of the *definiens* of $\Sigma\models^{*}\alpha$. If there is such a β it must have at most 9 parameters other than those occurring in $\Sigma\cup\{\alpha\}$, by the preceding lemma. There is, of course, an infinite number of such formulas, but they can all be put into logically equivalent disjunctive normal form. There are only a finite number of these (2^(2^(n+9))), where there are n parameters in $\Sigma\cup\{\alpha\}$. These can all be enumerated, and we can check to see if any of these is a suitable β. If not, the procedure fails.

□

This theorem illustrates the first awkward property of *consequence, which is that it is computationally complex to handle. The algorithm sketched in the proof of theorem 3 is of order of complexity 2^(2^n), which is horrendous. Doubtless it is not the most efficient algorithm, but because of the existential quantifier in the *definiens*, no algorithm is going to be very swift. The situation is worse if we move to first order logic. $\models^{LP}$ is Σ_1 in the arithmetic hierarchy; and as one can check, $\models^{*}$ is no worse than Σ_2. I suspect that in this case there is no way at all of putting a bound on the search for β (though I have no proof of this at the moment). If this suspicion is right then $\models^{*}$ is Σ_2. There is therefore no hope of even a semi-decision procedure.

A second, and more major, problem with *consequence is the following:

Theorem 4

Σ^{*} is not closed under $\models^{LP}$.

Proof

Let $\Sigma=\{p, \neg p\vee q, r, \neg r\vee\neg q\}$. Now, $\Sigma\models^{LP} q\vee p!$, but $\Sigma\not\models^{LP} p!$ and $\Sigma\not\models^{LP}\neg q\vee p!$ (as simple models show). Hence $\Sigma\models^{*}q$. Similarly, using r, $\Sigma\models^{*}\neg q$. But $\Sigma\not\models^{LP}q\wedge\neg q$. Hence $\Sigma\not\models^{*} q\wedge\neg q$. (Theorem 2 iv).) Thus, this instance of adjunction fails.

□

The non-closure of Σ^{*} under $\models^{LP}$ is problematic, since it means that we cannot establish certain *consequences of Σ and then carry on LP reasoning and be sure that we are safe. Thus, in a sense, *consequence undercuts the safeness of LP, upon which is piggy-backs.

4 Minimal Inconsistency: LP^{m}

Is there any way of finding a default reasoning for consistency that does not have these draw-backs? There is. To motivate it, consider the following. We are trying to make precise the idea that there is no more inconsistency around than we are forced to assume. A natural way of rendering this idea precise is by restricting ourselves to those models of the premises that are as consistent as possible, and then looking to see whether the conclusion is true in all of these. (This idea is due to Dirk Batens [1984].) Thus, suppose that we have defined an ordering of inconsistency, $<$, on interpretations, so that $\nu<\mu$ iff μ is more inconsistent than ν. Then we can define *minimal inconsistency (mi) consequence*, $\vDash^m$, as follows:

Definition 3

i) ν is a model of Σ iff every member of Σ is true under ν (i.e., $1\varepsilon\nu(\beta)$ for all $\beta\varepsilon\Sigma$).

ii) ν is a mi model of Σ iff ν is a model of Σ and if $\mu<\nu$, μ is not a model of Σ.

iii) $\Sigma\vDash^m\alpha$ iff every mi model of Σ is a model of $\{\alpha\}$.

The next question is how to define the ordering, $<$. For the rest of this section we will consider propositional logic only. One obvious possibility is to define $\nu<\mu$ to mean that ν makes fewer formulas (in the sense of cardinality) both true and false than μ does. However, it is easy to see that this will not work, since any inconsistent interpretation makes an infinite number of formulas both true and false; so any interpretation of an inconsistent theory would be minimal. $\vDash^m$ would therefore collapse back into $\vDash^{LP}$ for inconsistent sets of premises, violating Maxim M. We might try, instead, to define $\nu<\mu$ to mean that ν makes fewer propositional parameters both true and false than μ does. This is much better, but gives the wrong results, too, when Σ requires an infinite number of parameters to be inconsistent, for example, if $\Sigma=\{p_i!;\ i>10\}$ (where , $\{p_i;\ i\geq0\}$ is an enumeration of the propositional parameters). For then any LP interpretation will have the same number of inconsistent parameters, viz, $\aleph_0$. Thus, again, $\vDash^m$ would collapse into $\vDash^{LP}$. This is undesirable since, by the gloss on Maxim M, we would still like $\Sigma\cup\{p_0,\neg p_0\vee p_1\}$ to give us p_1 in the example.

Thus, ordering by cardinality will not work. Another possibility is ordering by set theoretic containment. Thus, let $\nu!_P = \{q\ ;\ 1\varepsilon\nu(q!)\}$ and $\nu!_F = \{\alpha;\ 1\varepsilon\nu(\alpha!)\}$. Then we might define $\nu<\mu$ either as $\nu!_P\subset\mu!_P$, or as $\nu!_F\subset\mu!_F$ (where $\subset$ is *proper* subsethood). Note that in either case the ordering is only a quasi-ordering. This, however, is not important.

Which definition of $<$ should we use? Consider the second definition first (that employing formulas, not propositional parameters). This is, in fact, unsatisfactory; for if we did define it in this way we would have the following:

(Counterfactual) Theorem 5

If Σ is any inconsistent, finite, set of sentences, then $\Sigma^m=\Sigma^{LP}$.

Proof

Note first that $\Sigma^{LP}\subseteq\Sigma^m$. (If truth is preserved in all models it is preserved in all mi models.) For the converse, suppose that $\Sigma\nvDash^{LP}\alpha$. Clearly, we need concern ourselves only with interpretations of the parameters in $\Sigma\cup\{\alpha\}$. V, the set of such interpretations, is clearly finite. Let V' be the subset of V of evaluations which model Σ but not $\{\alpha\}$ (i.e. $\nu(\alpha)=\{0\}$). This is finite and is non-empty. Choose a member of V', ν, such that there is no member of V', μ, such that $\mu!_P \subset \nu!_P$. This is possible since V' is finite. We will show that ν is a mi model of Σ (NB, with respect to $!_F$) and hence $\Sigma \nvDash^m \alpha$. Suppose for *reductio* that μ is a model of Σ and $\mu!_F \subset \nu!_F$. From the second of these it follows that $\mu!_P \subset \nu!_P$. Since μ is a model of Σ is follows that $1 \in \nu(\alpha)$. Moreover, since Σ is inconsistent, it follows that for some propositional parameter, q, $\mu(q)=\{1,0\}$ (or μ would make every formula consistent). But then $\mu(\alpha\wedge q)=\{1,0\}$, and $\nu(\alpha\wedge q)=\{0\}$. Contradiction.

□

Counterfactual theorem 5 shows that if we were to define $<$ using $!_F$ then for (finite) inconsistent sets, $\vDash^m$ would not be an extension of $\vDash^{LP}$, violating Maxim M, as the gloss on it showed. Thus $\nu!_F$ is an inadequate measure of the consistency of ν. We therefore choose $\nu!_P$, and drop the '$_P$':

Definition 4

If μ and ν are interpretations:

i) $\nu! = \{p;\ 1 \in \nu(p!)\}$
ii) $\mu<\nu$ iff $\mu! \subset \nu!$

What, now, are the properties of $\vDash^m$? Let us start with a few examples.

a) First, $\Sigma=\{q!,\ p,\ \neg p \vee r\} \vDash^m r$. To see this, note that for any (mi) model, ν, of Σ $q \in \nu!$, and $1 \in \nu(p)$. But $1 \in \nu(\neg p \vee r)$. Hence $1 \in \nu(\neg p)$ or $1 \in \nu(r)$. But the first of these is not possible in an mi model since then $p \in \nu!$ and there are models of Σ, μ, for which $\mu!$ is just $\{q\}$. Hence, in an mi model, ν, $1 \in \nu(r)$. This shows that counterfactual theorem 5 is indeed counterfactual.

b) Next $\Sigma=\{p! \vee r,\ p! \vee q!\} \nvDash^m r$. For if $\nu(p)=\{0,1\}$ and $\nu(q)=\nu(r)=\{0\}$, ν is an mi model of Σ but not $\{r\}$. (ν is mi, since if we made p consistent we would have to make q inconsistent.) Note that $\Sigma \vDash^{LP} r \vee p!$, $\Sigma \nvDash^{LP} p!$ and $\Sigma \nvDash^{LP} \neg r \vee p!$, as simple models show. Hence $\Sigma \vDash^* r$. Thus *deducibility does not imply mi-deducibility.

c) Finally, $\Sigma=\{r!,\ s \vee p!\} \vDash^m (r\wedge s)!$. For in any model of Σ, r! is true. Moreover, in any mi model s is true. Hence $(r\wedge s)!$ is true. Note that $\Sigma \nvDash^* (r\wedge s)!$ since $\Sigma \nvDash^{LP} (r\wedge s)!$ (by theorem 2 iv). Hence mi-deducibility does not imply *deducibility.

Some general properties of mi-consequence may be summarised in the following theorem.

Theorem 6

i) $\Sigma^{LP} \subseteq \Sigma^m$.
ii) In general, this containment is proper.
iii) $\Sigma^m \subseteq \Sigma^{CL}$.
iv) In general, this containment is proper.
v) If Σ is (classically) consistent then $\Sigma^{CL}=\Sigma^m$.
vi) Σ^m is closed under $\vDash^{LP}$.

Proof

i) If truth is preserved in all models it is certainly preserved in all mi models.
ii) See example a) above. $\Sigma \nvDash^{LP} r$.
iii) Define an interpretation to be *classical* iff no propositional parameter takes the value {1,0}. The class of classical interpretations characterises classical logic. Obviously if a model is classical it is mi. Hence if truth is preserved in all mi models it is preserved in all classical models.
iv) Clearly, $\{p!\} \nvDash^{m} q!$. But $\{p!\} \vDash^{CL} q!$.
v) If Σ is consistent then the mi models are exactly the classical models.
vi) Suppose $\Sigma^{m} \vDash^{LP} \alpha$. Then α is true in all models of Σ^{m}. *A fortiori*, α is true in all models of Σ, and hence all mi models of Σ.

□

These facts show that mi-consequence has the right kind of properties for a notion of default inference, and, in particular, unlike Σ^{*}, Σ^{m} is closed under $\vDash^{LP}$. $\vDash^{m}$ is also easier to handle computationally than $\vDash^{*}$, as the proof of the following theorem shows:

Theorem 7

$\vDash^{m}$ is decidable.

Proof

The following algorithm is obviously a decision procedure for $\vDash^{m}$ (for finite sets of premises). To determine whether $\Sigma \vDash^{m} \alpha$, enumerate the set of all the parameters in $\Sigma \cup \{\alpha\}$. This is finite. Suppose it has cardinality n. Enumerate all the interpretations for these parameters. (This set has cardinality 3^n.) Check to see which of these is a model of Σ and discard those that are not. Compute $\nu!$ for each remaining model, ν. Discard all those for which there is some μ such that $\mu! < \nu!$, This is essentially a sorting operation, and so can be done in time 3^n.log(3^n). What remains are the mi models of Σ. If α is not true in any one of these, consequence fails; otherwise it passes.

□

This algorithm is not the most efficient. It is of order n.3^n, which, though computationally horrible, is still much better than the 2^(2^n) of $\vDash^{*}$. Thus, we see that, at least from propositional inferences, $\vDash^{m}$ betters $\vDash^{*}$ on the two problems I pointed out in section 3.

A further fact of interest about mi-consequence is the following:

Theorem 8

If Σ is finite then Σ^{m} is trivial iff Σ^{LP} is trivial.

Proof

The proof from right to left is trivial in virtue of i) of theorem 6. In the other direction: Suppose that Σ^{LP} is non-trivial, i.e. for some formula, α, $\Sigma\nvDash^{LP}\alpha$. It follows that for some propositional parameter, p, $\Sigma\nvDash^{LP}p!$. (Since if you can prove q! for all q then you can prove all formulas, as a simple induction shows.) Then there is an LP interpretation, ν, which models Σ but not {p!}. Let V be the set of all models of Σ ordered by <. We need consider only evaluations of the propositional parameters of $\Sigma\cup\{p!\}$, and since $\Sigma\cup\{p!\}$ is finite, V is finite. It follows that either ν is mi, or there is a mi model < ν. In either case, there is a mi model of Σ which is not a model of {p!}. Hence $\Sigma\nvDash^{m}p!$, i.e., Σ is not mi-trivial.

□

The above theorem is reassuring since it shows that in propositional logic Σ^{m} does not blow up unless Σ^{LP} does (for finite Σ). Hence there is no more danger of collapse into triviality with $\vDash^{m}$ than with $\vDash^{LP}$. Unfortunately, the proof of the above theorem depends essentially on the fact that Σ is finite. Whether or not the theorem is true if Σ is infinite I don't know.

Note also that there are some extensions of LP where mi-consequence may well blow up even where ordinary consequence doesn't. Consider the extension of propositional LP to include i) an implication connective, →, satisfying *modus ponens* ii) propositional quantifiers (for example, a suitable relevance logic). Now consider the set of formulas $\Sigma=\{\exists p(p!)\}\cup\{p_i!\rightarrow p_{i+1}!;\ i\geq 0\}$, where the indexation enumerates propositional parameters. Clearly, Σ may have a non-trivial model (for example, one in which p_0 is not true). But it has no minimal model: Suppose it did. Call it ν. Then there must be a least i such that $p_i\in\nu!$. Now consider the interpretation, μ, that is the same as ν except that the least i such that $p_i\in\mu!$ is one greater. This is a model, and $\mu<\nu$. Contradiction. Thus, Σ^{m} is trivial since any formula is true in all mi models of Σ (since there aren't any).

Theorems such as theorem 8 are not, therefore, to be taken for granted. (Compare this situation with that for $\vDash^{*}$, where, as we noted (theorem 2v)), quite generally, Σ^{*} is trivial iff Σ^{LP} is. Σ^{*} does, therefore, have some advantages over Σ^{m}.

5 First Order $\vDash^{m}$

Let us now turn to the first-order case. To apply the machinery of minimal inconsistency to this we need to define the notion $\nu!$, where ν is an interpretation of the language, L. How is this to be done? What corresponds to the set of propositional parameters in the propositional case is the set of atomic facts in ν. To make this easy to talk about, let L_ν be the language, L, with all constants removed, but extended by a canonical name for each member of the domain of ν. (We might as well take the domain of ν, itself, as the set of names, each naming itself.) Now define $\nu!$ as follows:

Definition 5

$\nu! = \{p \,;\ p$ is a(n atomic) formula of L_ν and $1 \in \nu(p!)\}$

A little care needs to be taken at this point. For if we now defined $<$ as in the propositional case, the wrong results would ensue. Let L be the language with just two monadic predicates, P and Q. Let $\Sigma = \{\forall x(Px!)\}$. For every interpretation of Σ, ν, with domain D, $\nu! \supseteq \{Pd;\ d \in D\}$. This is minimised by setting the cardinality of D to 1 (since an interpretation must have at least one member). Thus, every ν with minimal $\nu!$ has one member. Hence, e.g., $\forall xQx \vee \forall x \neg Qx$ would be an mi consequence of Σ, which it obviously ought not to be.

Clearly, what has gone wrong is that we ought not to be comparing interpretations with different domains. Thus, we should define:

Definition 6

$\mu < \nu$ iff μ and ν have the same domain and $\mu! \subset \nu!$

Since at least one of every atomic formula and its negation must occur in the diagram of an interpretation, this is equivalent to saying that μ and ν have the same domain and the diagram of μ is a proper subset of the diagram of ν.

Having got the appropriate definition straight, let us now consider what happens to the theorems of the previous section in first order logic. Theorem 6 continues to hold, as may easily be checked. Theorem 7 we would hardly expect to hold: since first order $\models^{LP}$ is Σ_1 in the arithmetic hierarchy, $\models^m$ is not likely to be decidable. (In fact, it must be at least Σ_1 since its set of logical truths coincides with that of LP and CL.) What recursive complexity it does have I do not know. The proof of theorem 8 appeals to the fact that the interpretations of a fixed finite vocabulary are finite in number. This is no longer true in the first order case, of course; so the proof fails. Whether there is any analogue of theorem 8 for first order logic, I do not know.

6 Conclusion

It is clear from the above that $\models^m$ seems to have the right kinds of property for a default logic for consistency. I therefore propose it as such. It will be clear from the above, however, that some of the open questions concerning it may have an important bearing on the issue. I therefore conclude with a summary of these.

i) Does theorem 8 hold for infinite Σ?
ii) Do any analogues of theorem 8 hold for first order logic?
iii) Where is first order $\models^m$ in the arithmetic (or analytic!) hierarchy?
iv) How can first order $\models^m$ be handled computationally? That is, what algorithms capture (a significant part of) it.

*This paper was written while I was a Project Visitor at the Automated Reasoning Project, Australian National University, during January 1988. I am grateful to them for providing such a stimulating environment.

References

Asher N.M. and Kamp J.A.W. [1986] 'The Knower's Paradox and Representational Theories of Attitudes', in Halpern J.Y. (ed), *Reasoning About Knowledge*, Morgan Kaufmann.

Batens D. [1984] 'Dialectical Dynamics within Formal Logics', *Proceedings of the 5th Joint International Conference on the History and Philosophy of Science*, Hungary, 1984.

Belnap N. [1977] 'A Useful Four-Valued Logic', in Dunn J.M. and Epstein G. (eds), *Modern Uses of Multiple Valued Logic*, Reidel.

McCarthy J. [1980] 'Circumscription - a form of non-monotonic reasoning', *Artificial Intelligence* 13, 27-39.

McDermott D. [1982] 'Non-Monotonic Logic II: non-monotonic modal theories', *Journal of the Association for Computing Machinery* 29, 33-57.

McDermott D. and Doyle J. [1980] 'Non-Monotonic Logic I', *Artificial Intelligence* 13, 41-72.

Priest G. [1980] 'Sense, Entailment and *Modus Ponens*', *Journal of Philosophical Logic* 9, 415-35.

Priest G. [1987] *In Contradiction: a study of the transconsistent*, Nijhoff.

Priest G. [198+] 'Reasoning about Truth', *Artificial Intelligence*, to appear.

Priest G. and Routley R. [1983] 'On Paraconsistency', Research Report #13, Department of Philosophy, Australian National University. Reprinted as the introductory chapters of Priest G. *et al* (eds), *Paraconsistent Logics*, Philosophia Verlag, 1988.

Reiter R. [1980] 'A Logic for Default Reasoning', *Artificial Intelligence* 13, 81-132.

Routley R. *et al* [1982] *Relevant Logics and their Rivals*, Ridgeview.

Automated Reasoning Project
Technical Report Series

Meyer, R.K. & Slaney, J; *Computing Demorgan Monoids*

Brink, C; *Some Background on Multisets*

Brink, C; $R^{\neg}$*-Algebras and* $R^{\neg}$*-model Structures as Power Constructs*

Brink, C. & Heidema, J; *A Verisimilar Ordering of Theories Phrased in a Propositional Language*

McRobbie, M.A., Meyer, R.K. & Thistlewaite, P.B.; *Towards Efficient Knowledge-Based Automated Theorem Proving for Non-Standard Logics*

Malkin, P.K. & Martin, E.P.; *Logical Matrix Generation and Testing*

Lavers, P.; *Relevance and Disjunctive Syllogism*

Martin, E.P.; *A Formalized Metalanguage for* P-W *and* S

Meyer, R.K. & Mortensen, C.; *Alien Intruders in Relevant Arithmetic*

Fuhrmann, A.; *Reflective Modalities and Theory Change*

Meyer, R.K., Giambrone, S., Urquhart, A., & Martin, E.P.; *Commutative Monoid and Semilattice Semantics for* P-W

McRobbie, M.A. & Siekmann, J.H.; *Artificial Intelligence: Perspesctives and Predictions*

Read, S.; *Proof Theory and Semantics for Relevant Logics*

Hazen, A.; *Predicative Arithmetic*

Bollen, A.W.; *A Relevant Extension to* PROLOG

Pelletier, J.; *Further Developments in* THINKER, *an Automated Theorem Prover: Proof Condensation, Adding Identity, 'Empirical' Issues in Computational Complexity, Pseudo-Parallel Subproof Development*

Siekmann, J.; *Unification Theory*

Siekmann, J. & Szabo, P.; *The Undecidability of the DA-Unification Problem*

Priest, G.; *Reasoning About Truth.*

Priest, G.; *Consistency by Default.*

Slaney, J.K.; *General Logic.*

Slaney, J.K.; *Solution to a Problem of Ono and Komori.*

Ulrich, D.; *Small Matrices Validating the Two-Variable Theses of the Classical Equivalential and Implicational Calculi.*

Slaney, J.K.; *The Ackermann Constant Theorem: A Computer-Assisted Investigation.*