THE IMPLICATIONS OF PARACONSISTENCY

John Slaney
Automated Reasoning Project
Australian National University
GPO Box 4, Canberra, ACT 2601
Australia

Abstract

This is a connected scries of arguments concerning paraconsistent logic. It is argued first that paraconsistency is an option worth pursuing in automated reasoning, then that the most popular paraconsistent logic, fde, is inadequate for the reconstruction of essential first order arguments. After a case is made for regarding quantifiers as dyadic rather than monadic operators, it is shown that the addition of such quantifiers to fde allows an implication connective to be defined yielding the known logic BN4. Refining the treatment of implication in a manner similar to that found in intuitionist logic leads to the more interesting system BN.

1 Paraconsistency

So many authors recently have speculated on the advantages of paraconsistent reasoning for the inference engines of intelligent systems managing large bodies of data that one hesitates to enter the lists again in support of the idea. Prominent among its champions are Belnap, Dunn, Sylvan, Da Costa, Arruda, Priest, Brady, Mortensen, Urbas and Patel-Schneider.[1]

Briefly, paraconsistent logic is motivated by distress at the classically valid form of inference

$$\frac{A \qquad \sim A}{B}$$

Given an inconsistent database and a classical logic for deducing information from it, inference might go *anywhere.* Classical logic is supposed to be the maximal system, closed as a logic should be under uniform substitution for its variables, in which inference always preserves truth. The claim that classically valid inference *does* preserve truth is already suspect on many grounds— for example such inferences as that from A to $B \rightarrow A$, that from $(A \rightarrow B) \rightarrow A$ to A and that from $A \rightarrow (A \rightarrow B)$ to $A \rightarrow B$ can be assessed for truth-preservation only on the basis of some *theory* as to what the arrow means, and the opponents of classical logic are not about to accept its theories of meaning without argument but in any case in the canvassed uses of logic truth preservation by itself is not really enough, because the presented information need not be truth-like. Especially, it need not be consistent. When we ask what pronouncements should be made by a mechanical reasoner, we typically have in mind a system on whom the World (Everything that is the Case) impinges not directly, for it has no senses or like means of immediate access, but indirectly through what it has been told. And of course what, it has been told is not constrained to be verisimilar, especially if it gets its tellings from many independent sources.

So even if it is necessary that the World be closed under classical logic, the best reasoning stratregy for beings like ourselves with access only to a corrupt and patchy version of it may well be to shy from the blithe thought that since inconsistency is impossible it does not matter what conclusions we draw from a contradiction. Moreover, even if we have the wisdom to be as gods, either never having inconsistent belief sets or at least never deducing too much silliness from them, the same is not true of our machines. Outside science fiction, automated reasoners hick the *nous* to recognise silliness when it happens. Now some fast and dirty checks for grossly contradictory data can usually be applied, and in current database technology, without really intelligent management systems, these suffice to keep things clean. Where the user interface at the input end is something as sophisticated and flexible as first order logical notation, however, they certainly will not. Even in fragments of the language for which consistency is decidable the complexity of known decision methods is exponential, so it is not generally feasible to perform a full consistency check for a large body of information. This is especially so where the information tends to change over a short span, so that the life of data in the record may be shorter than the duration of a consistency check. So where there is little control over the format in which data might be entered, and where the roll-over time for data is short, inconsistency must be treated just like any other falsehood: we do not want it, and if we discover it we shall isolate or eliminate it, but in the meantime we must continue to reason in the hope that it does not lead to too

[1]See the bibliography below. Speculation is all that any of the rioted authors has achieved, for despite the rash of papers extolling the virtues of such reasoning the Great Paraconsistent Database Manager has yet to be written.

many wrong predictions. And there's the rub. What we want overwhelmingly is that our automatic reasoner shall give us predictions, strong and useful ones at that, and also that the claims it makes be *true.* These desiderata are in evident conflict. The sufficient strength re-(luirernent cannot be compromised very radically, or our system ceases to be useful, so we have no option but to weaken the second desdideratum to an attainable goal such as that *nearly* all of the predictions be true. Since inconsistency is widespread, undetectable and dangerous i.o this weaker goal, it seems prudent to base inference on a fault-tolerant logic in which the odd contradiction is allowed to occur without rendering the data *totally corrupt.*

2 First Degree Entailments

There are many pararonsisfent logical systems on the menu, but for the remainder of this paper we shall concentrate on the suggestion, emanating from work of Smiby's dated around I960, first given its paraconsistent interpretation by Dunn and taken seriously by Belnap, Sylvan, Priest, Levesque and Patel Schneider among others, that we liberate truth table reasoning minimally by allowing statements to take not just truth values but *sets* of truth values. There are evidently four such sets:

$$\{\top\} \qquad \{\bot\} \qquad \{\top, \bot\} \qquad \emptyset$$

That is, a statement can be true, false, both or neither. To assist readability, these four cases will be abbreviated to the single letters T, F, B and N respectively. Now conjunction, disjunction and negation, are easily seen to correspond to functions on these four cases:

~	
T	F
B	B
N	N
F	F

&	T	B	N	F
T	T	B	N	F
B	B	B	F	F
N	N	F	N	F
F	F	F	F	F

v	T	B	N	F
T	T	T	T	T
B	T	B	T	B
N	T	T	N	N
F	T	B	N	F

The relation of material implication also extends in the obvious way from truth values to sets of truth values; formula *A* materially implies formula. *B* iff in passing from *A* to *B* we neither lose truth nor gain falsehood, These two conditions are of course independent,. They induce an implication order *on* the sets of values thus:

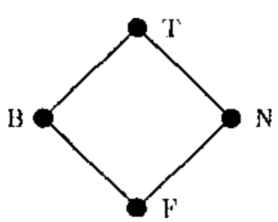

That is, F implies everything; everything implies T; B and N imply themselves but not each other.

The logic determined by the four-point structure is the system fde of relevant first degree entailments, sometimes called "tautological entailments". As is well known,[2] the alarming inference from an arbitrary contradiction to All Points West is avoided by fde, but as is again well known this avoidance comes at a price. The price is the invalidity of *resolution* in the form

$$\frac{A \vee B \qquad {\sim}B \vee C}{A \vee C}$$

and of its special case in which *C* is null, the *disjunctive syllogism*

$$\frac{A \vee B \qquad {\sim}B}{A}$$

Gven the disjunctive syllogism, we could argue thus:

$$\frac{\dfrac{B}{A \vee B} \qquad {\sim}B}{A}$$

After the. shock and horror at the thought of rejecting those principles has begun to subside it usually crystallises into two lines of opposition to this kind of paraconsistent logic. The first is that the disjunctive syllogism at least, and maybe resolution also, is *just obviously good reasoning.* This is usually said in a special voice reserved for that emphasis needed in philosophy when labouring the obvious in the face of some paradoxer inclined to point out that arguments in favour of the obvious are a little hard to come by. The second is that without resolution and material detachment reasoning would be hamstrung. This second claim deserves to be taken seriously, and is undecided in the present state of the research program. The present paper is motivated by the need to take it seriously, and will address parts of it below, but first it is necessary to say something in response to the former objection.

The simplest response is to challenge the claim of obviousness head-on. Given the need to reason under inconsistent assumptions, it just is *not* obvious that the disjunctive syllogism, for example, is applicable. What underlies the apparent naturalness of that argument form is perhaps some such thought as this: we are given that one of the disjuncts, either A or *B,* is the case, and we are given that it is not *B;* so what, can it be but *A?* That is, a disjunction cannot, be *barely* true, for its truth must arise out of that, of one of its disjuncts. So where we have assumed enough to make the disjunction true we must have assumed enough to make it appropriate to go on with one disjunct or the other. Now if ${\sim}B$ also holds under our assumptions, it can't be *B* that underpins the disjunction, on pain of inconsistency, so it must be *A.* Now the difficulty I find with this thought is not that the implied argument is circular, appealing precisely to the disjunctive syllogism in the metalogic, but that the claim that "it can't be *B* that underpins $A \vee B$" is in the disputed case plain false. Look at the displayed derivation: the disjunction *does* come from B, in the grossest and most blatant way, and if this fact causes us to incur the pain of inconsistency then such a pain we must bear.

[2]Sec for example [Belnap, 1977].

There is a slightly deeper point to be made here, about the nature of logical theory. The starting point of all logic is the question of which are the valid (perfect, reliable, necessarily rational) forms of argument. What we do in answer to this question is to think up some argument forms which seem good to us, isolate what we take to be the logical constants involved, formulate rules of inference to govern the behaviour of these and thus arrive at a formal calculus. It is usually after this that we devise a semantic story suited to the inferences we find acceptable. What, we thus produce is a theory sufficing to catalogue the allowable argument, forms across the chosen range. We somehow have the impression that our logic is inexorable, so that to question it is not even intelligible. ligible. clearly this inexorability is an illusion. The formed theory goes a long way beyond the intuitive reflections that gave rise to it, so that it. applies to many arguments of sorts not considered at all when we so readily assented to the rules. When we accepted *modus ponens* as part of the meaning of 'if we may have been unable to imagine what a counter-example could possibly be like; but later we come across the sorites paradox or the semantic antinomies and perhaps our imagination is expanded to encompass such a thing at least we might well come to the view that the orthodox formulations of *modus ponens* miss some of the subtleties of the rules which really are or ought to be parts of the meaning of 'if. Similarly, when we considered resolution or the disjunctive syllogism we may have thought: yes, 1 reason like that; 1 would regard it, as quite irrational not to. But of course we were not. then thinking of reasoning situations that involve taking inconsistent assumptions seriously. When we do confront the problem of the logic of inconsistent theories, such as large unstructured databases, we find our intuitions less than ideally firm. It is disquieting to have our logic challenged, but the mature philosophy of logic should *start* from tin.' reflection that logical theories, like any theories, are perpetually open to challenge and that logic can no more be shielded from the hard questions in its domain of application than can any other study.

So much for the assertion of obviousness. The other major objection to paraconsistent logic in the presently canvassed style is that. it. \s too weak to allow reconstruction of standard reasonings whether their premises be consistent or inconsistent. Weak it is, and its weakness is of two sorts. Firstly it omits many of the classically valid inferences, such as the noted resolution and disjunctive syllogism, which turn just on the properties of the extensional connectives. It remains to be seen whether the omitted principles really pay their way in classical logic or whether anything they are good for can be achieved paraconsistenfly by some other means. This paper is not really about weakness in that sense. The second sense in which fde is weak is slightly less obvious but none the less important. It, is weak in that it is fragmentary. Classically, all the other connectives needed for truth-table reasoning can be defined in terms of negation and either conjunction or disjunction, so once the logic of these connectives is settled there is no more to be said, but this is not the case once we move to fde. Crucially, the operation of material implication is inexpressible in terms of the given extensional connectives. True, we can define a connective in the classical way

$$A \supset B =df \sim A \vee B$$

but this is no sort of implication in the present context as it fails to satisfy any suitable principle of detachment. So the language stands in need of enrichment in places where that of classical logic does not.

This need for enrichment is even more obvious at the quantiiicational level. Clearly we can add quantifiers of the usual classical sort, but the effects of so doing are not the usual classical effects. With a model we may associate a nonempty domain of individuals, letting an interpretation assign to each predicate symbol two subsets of the relevant power of the domain: positively, those n-tuples that satisfy the predicate, and negatively those that dissatisfy it. Then an atomic formula

$$Ft_1 \ldots t_n$$

is true for the interpretation iff the tuple consisting of the individuals assigned to terms $t_1 \ldots t_n$ is in the positive set assigned to F, and false iff that tuple is in the negative set assigned to *F.* Truth and falsehood for com pounds built up with connectives is as before, and the clauses for quantifiers are that $\forall v\, A$ is true iff *A* is true for all V-variants of the given interpretation and false iff *A* is false for some invariant, while the existential $\exists v\, A$ is true iff *A* is true for some invariant and false iff *A* is false for all of them. This is fine as far as it goes but it does not go far enough. Most of the important forms of quantificafional argument are not captured by this logic because they are not even expressible. In order to validate

All footballers are bipeds
Socrates is a footballer
So Socrates is a biped

—an argument commonly supposed to be of an acceptable form we need to be able to say not only that everything is bipedal but that all *footballers* are. That is, we cannot get by for perfectly ordinary reasoning purposes without restricted quantifiers. Now classically there is a familiar trick available to get us out of this problem: we can re-express 'Every footballer is a biped' as 'Everything is if~a-footballer-then-a-biped'.The latter we can express using the conditional connective and an unrestricted quantifier. But where no genuine conditional is available there is no such trick available either.

3 Quantifiers

The simple quantiiicational devices we were all taught in our first, year of learning logic are quite inadequate to express most of the quantifiers used in reasoning. Even in classical logic, where many conceptually distinct notions are collapsed together in order to allow a poor base vocabulary to give rough and ready expression to as wide a range of constructions as possible, we have no means of proving logically valid such simple arguments as

Most logicians are animals
Most logicians are underpaid
Therefore some animals are underpaid

and no convenient way of bringing out the ambiguity of a sentence like

Most students are smarter than most wombats.

Nor are we given a logical form, identical up to variations in the quantity indicator, for examples like

All Fs are Gs
Several Fs are Gs
At least six Fs are Gs
Hardly any Fs are Gs
Finitely many Fs are Gs

We are unlikely to want to investigate the logic of many such quantifiers: in the case of vague quantifiers like 'several', 'few' and 'hardly any* it is insufficiently interesting to repay the effort, while for such examples as 'finitely many* and 'most' it is inconvenient because not compact. Nonetheless, we should give thought to their logical form, if only because the general form of the quantifiers we *do* need to investigate also underlies the others.[3]

So the least we need in the way of notation for quantifiers *Q* is a tripartite structure

$$\mathcal{Q} = \langle Q,\ V,\ R\rangle$$

where *Q* is a quantity indicator ('all*, 'some', 'most', 'thirteen*, etc.), *V* is the variable, or in the more general case the several variables, bound by the quantifier, and *R* is a formula specifying the subset of the domain over which the quantifier ranges. Thus we might represent Tor every student, x' as $(\forall x|Sx)$, and write 'All students are more intelligent than some wombats'

$$(\forall x|Sx)\,(\exists y|Wy)\,Ixy$$

That is, for every thing, x, such that x is a student, for some thing, y, such that y is a wombat, x is more intelligent than y. This connective-free notation for restricted quantifiers can easily accomodate the "non-standard" examples above: given a notation *M* for 'most' we can write

$$(\mathcal{M}x|Sx)\,(\mathcal{M}y|Wy)\,Ixy$$

as in 'Most students, even some Logic 1 students, are smarter than most wombats', and distinguish it, from

$$(\mathcal{M}y|Wy)\,(\mathcal{M}x|Sx)\,Ixy$$

as in 'Most students are smarter than *most* wombats; that's even true of Willie the Wonder Wombat'.

So we have some notation for quantifiers which does not depend on the connectives in our logic. Now in the cases of the usual quantifiers to which logical theory applies, the universal and existential ones, we need to ask what inferences a decent logic should validate. In the classical case the answer is easy. For the universal quantifier we want premises *V* to entail $(\forall v|A)\,B$ if *B* can be deduced from *T* and *A* taken together without cheating. Cheating here is appealing to free occurrences of *v* in the premises P. Conversely, we want the above argument form about Socrates the footballer: for any term t,[4] from $(\forall v|A)\,B$ and A^t_v we can deduce B^t_v. The existential quantifier is governed by the dual principles. For any t, $(\exists v|A)\,B$ follows from A^t_v and B^t_v, and conversely we want the rule of choice: from $(\exists v|A)\,B$ we can get from whatever follows from *A* and *Li,* again without cheating (in this case, no free *v* in the conclusion or in any side premises appealed to). Semantically, the universal $(\forall v|A)\,B$ is true iff whatever individual *i* from the domain is assigned to v, *A* materially implies *B* while the existential $(\exists v|A)\,B$ is true iff the assignment of some individual to *v* satisfies both *A* and *B.*

For the four-valued paraconsistent logic we can keep the classical truth conditions, making $(\forall v|A)\,B$ true for V iff for every *v*-variant *V'* of V

$$\mathcal{V}'(A) \leq \mathcal{V}'(B)$$

where the order is the relation of material implication as specified earlier. For the existential we may take $(\exists v|A)\,B$ to be true iff for some *v* variant V of V

$$\top \in \mathcal{V}'(A) \cap \mathcal{V}'(B)$$

However, we also need falsehood conditions, and these are less clearly dictated by the need to capture simple syllogistic inferences. The most, natural condition under which a universal claim is to be denied is the existence of a counter-example to it. That is $\bot$ is in $\mathcal{V}((\forall v|A)\,B)$ ifr there is some v-variant V of V, such that

$$\top \in \mathcal{V}'(A) \text{ and } \bot \in \mathcal{V}'(B)$$

For the nonexistence of an *A* that is *B* the expected condition would surely be that it be false of everything in the domain that *it* is an *A* which is *B* 'Thai is, I_ is in $\mathcal{V}((\exists v|A)\,B)$ iff for every *v*- variant *V*

$$\bot \in \mathcal{V}'(A) \cup \mathcal{V}'(B)$$

The upshot of all this is that the existential quantifier, just as classically, is equivalent to the unrestricted form with a connective (conjunction).

$$\mathcal{V}((\exists v|A)\,B) \;=\; \mathcal{V}(\exists v\,(A \,\&\, B))$$

The universal quantifier, however, is something new.

4 Implication: BN4 and beyond

One oddity of classical logic, retained by some authors in order to smooth the definition of 'formula[1] and the statements of axioms or rules, removed by others in the interests if intelligibility is the vacuous quantifier. A formula such as

$$\forall y\,\exists x\,Fx$$

is either nonsense or merely a long-winded way of saving

$$\exists x\,Fx$$

[3]See Belnap's paper [Belnap, 1973] on restricted quantification for a start on some of the issues. The line taken here is rather different though not completely incompatible.

[4]Strictly, t must he free for *v* in *A.* The same applies to the introduction rule for the existential quantifier. Technical detail in these matters is not appropriate to this sketch, however.

Now when quantifiers are correctly construed as being of the type to make *two* predicates into a sentence, vacuous quantification is no longer just circumlocutionary. Consider the existential quantifier first. What does $(\exists x|A)\,B$ mean where x does not occur free in A or B'i Well it is true iff some x-variant makes both A and B true, and false iff they all make A or B false; but since x does not occur in either formula what holds for one x variant holds for all of them. In other words, $(\exists x|A)\,B$ is exactly equivalent in the vacuous case to A & B. Since no connectives are needed to define the dyadic existential quantifier, any logic with such a quantifier does not need a primitive conjunction: the connective & is definable in terms of the quantifier. This is so classically and remains so in our present paraconsistent logic— as it does, in fact, over a very wide range of first order systems.

The universal quantifier is even more interesting in this regard. Like the existential one it defines a connective which we can easily calculate to correspond to the following table of values:

	T	B	N	F
T	T	F	N	F
B	T	B	N	F
N	T	N	T	F
F	T	T	T	T

Unlike conjunction, this connective is not definable in terms of the usual extensional ones. It is however familiar, at least to the vast and happy readership of [Meyer *et al*., 1984], as the implication matrix of the logic I3N4. Let us fix it as a properly defined connective, then, by the usual trick of letting w be the first variable in the standard enumeration of variables which does not occur in A or B. Then

$$A \rightarrow B \;=df\; (\forall w|A)\,B$$

and it will readily be seen that $(\forall v|A)\,B$ is equivalent to $\forall v\,(A \rightarrow B)$, so the classical move of confining attention to unrestricted quantifiers is still available except in some vacuous cases. Once again, the connective lias been defined using quantification, so not only conjunction but also implication is redundant in a logic with dyadic quantifiers. Clearly, if we were to go on with the other Aristotelian primitive 'No *S* is *P*' we should be able to define negation as well, and thus arrive at a logic in which the only undefined constants are quantifiers. It is amusing to note that Aristotle was so close to an adequate logical vocabulary and theory: only the fact that he tried to manage without variables prevented him from completely anticipating Frege. It is also worth noting that conjunction, implication and negation continue to be definable in the present way over a wide range of logics, from all of the Anderson-Belnap relevant systems to Lukasiewicz many-valued ones. The motivational significance of this fact is considerable, since the intuitions regarding universality are much more firmly rooted than those regarding implication, to the extent that while one still hears complaints that the logic of implication is esoteric and contrived, superfluous to requirements or even downright incomprehensible, *no-one* dare say such things about the logic of 'all'.

BN4 may be a known logic, but it is not the end of the journey. While its underlying four-point lattice is as a De Morgan lattice polynomially free, so that it tells the whole story about first degree entailments, the same is not obviously true of its theory of implication. Such strange theorems as

$$(A \rightarrow B) \vee (B \rightarrow C) \vee (B \leftrightarrow \sim B)$$

are brought about not by any offered insight into the nature of reasoning but simply by the lack of different places to put B. These reflect the fact that the impli cation of BN4 is a *material* implication. While this is not necessarily a bad thing, and indeed BN4 is an interesting logic, there are related systems which are prooftheoretically more elegant and philosophically more satisfactory. Hence we move on towards a more refined account of implication.

The simplest way to loosen up BN4, giving it a minimally intensional implication operator, is by basing models not on a single assignment of values but on a partially ordered set of them. This idea is familiar from intuition ist logic: the values being epistemic rather than ontic, the "truth" at any stage is partial. The reasoner should therefore take account not only of the values of propositions as they are given but. also of the values as they might be on receipt of more information. For this purpose *deletions* from the data are not considered, though if we wished to investigate a non-nionotonic construal of logic they could be. Possible extensions to the data, being sets, are partially ordered by inclusion, inducing a natural partial order on the associated valuations. There should be a matching inclusion requirement on the domains of individuals as in the intuitionist case. Since there is addition only -never subtraction we should expect that for any atomic formula p, if *p* is true at a valuation point it remains true at all greater points, and if it is false it remains false. After all, one can hardly *lose* a warrant, for assertion or denial as m*ore* information comes in. The clauses for evaluating compounds should be so formulated that this heredity fact extends from the atomic case to all formulae.

The extensional connectives, defined if we like in terms of the existential quantifier and negation, behave in just the BN4 way at each individual evaluation point. The conditional, however, like the universal quantifier in terms of which it may be presented, is an inference ticket. As such it requires us to look beyond the presently given state of information to its possible extensions. It assures us of the availability of an inference from its antecedent to its consequent, or from the denial of its consequent, to that of its antecedent. Hence we must demand appropriate closure not only at the valuation point in question but at all its extensions:

$$\top \in \mathcal{V}_\alpha(A \rightarrow B) \iff (\forall \beta|\alpha \subseteq \beta)\Big(\mathcal{V}_\beta(A) \leq \mathcal{V}_\beta(B)\Big)$$

Recall that the order $\leq$ is the partial order of material implication on the four sets of values as for BN4. On the other side of the street, we may take an implication to be deniable just when there is a concrete counter-example to it. That is

$$\bot \in \mathcal{V}_\alpha(A \rightarrow B) \iff \top \in \mathcal{V}_\alpha(A) \wedge \bot \in \mathcal{V}_\alpha(B)$$

In this way we arrive at another known logic, the system BN (without the "4" since it is no longer 4-valued).

For the formal investigation of BN and its many splendid properties I can do no better than to advertise [Slaney *et al,* 1989].[5] What is worth remarking here is that we have come to BN not in order to secure those formal properties but as the result of a chain of philosophical reflections in which they played no part. We were led here specifically by consideration of the need for paraconsistency, then of the expressive weakness of certain paraconsistent logics and finally of the desirability, where our notions of truth and falsehood are epistemieally based, of looking beyond the given state of information in order to determine implication. What has not been shown is that BN is the One True Logic The reader seeking demonstrations that this is the only way to go will seek in vain: as an honest logician 1 can merely indicate that this is a *good* way to go. Nor, having got this far, can we rest content. At least, given the motivation in terms of automated reasoning, there must be some attempt to apply the ideas empirically;[6] and to invite such empirical testing is to give hostages to fortune. So be it then: having long urged that logic become an empirical science, 1 am in no position to flinch from such a thought.

References

[Belnap, 1973] Nuel D. Belnap. Restricted Quantification and Conditional Assertion. In H. Leblanc, editor, *Truth, Syntax and Modality,* pages xx-xx. Amsterdam, 1973.

IBelnap, 1977] Nuel I). Belnap. A Useful Four-Valued Logic, in J.M. Dunn & G. Epstein, editors, *Modcm Uses of Multiple- Valued Logics,* pages 8 37. Dordrecht, 1977.

[Dunn, 1976a] J. Michael Dunn. *Intuitive Semantics for First-Degree Entailments and "Coupled Irecs". Philosophical Studies,* 29:149 168, 1976.

[Dunn, 1976b] J. Michael Dunn. *A Kripkt-Style Semantics for R-Mingle Using a Binary Accessibility Relation. Studia Logtca,* 35:163 172, 1976.

[Levesque, 1984] Hector J. Levesque. A Fundamental Tradeoff in Knowledge Representation and Reasoning. FLAIR Technical Report 35, Fairchild Laboratory for AI Research, Ontario, 1984.

[5]It will be evident that Dunn's semantics in [Dunn, 1976b] for the semi-relevant logic R.M provided input to the present reflections. Thanks should be recorded to Dunn and to the co-authors of [Slaney *ct al.,* 1989] for many illuminating conversations on these matters as well as for their published formal work.

[6]R.A. Girle has produced tableau-based theorem provers for BN and some related systems. These are a good start, but clearly must be taken much further and applied to some realistic examples. The Automated Reasoning group in Canberra intends to produce workable software embodying some of the above ideas.

[Meyer *ct aL,* 1984] Robert K. Meyer, S. Giambrone & Ross T. Brady. Where Gamma Fails. *Studia Logica,* 43:247-256,1984.

[Patel-Schncider, 199l] Peter F. Patel-Schneider. A Decidable First-Order Logic for Knowledge Representation, *journal of Automated Reasoning,* forthcoming.

[Priest, 1980] Graham G. Priest. Sense, Truth and Modus Ponens. *Journal of Philosophical Logic,* 9:415 435, 1980.

[Slaney *et al,* 1989] John Slaney, Timothy Surendonk & Roderic Girle. Time, Truth and Logic. Technical Report TR-ARP-11/89, Automated Reasoning Project, Australian National University, Canberra, 1989.

[Urbas, 1990] Igor F. Urbas. Paraconsistency. *Studies in Soviet Thought,* 39:343 354, 1990.