CAISOR

Cognitive Autonomous and Information Systems Open Reference

Archive
1988-003

Erik Sandewall

1 Background and topic

The semantic entailment relation $\models$ of ordinary, monotonic logic is monotonic in the sense that if

$$\Gamma \models \alpha$$

and

$$\Gamma \subseteq \Gamma'$$

it follows

$$\Gamma' \models \alpha$$

where of course Γ and Γ' are sets of logical formulae, and α is one such formula. We wish to find a modified definition of $\models$ which is not monotonic in this sense, and which has plausible properties, particularly with respect to the uses of non-monotonic reasoning which are of current great interest in A.I. The definition should allow us to express and perform defeasible reasoning, where default "rules" (expressed as logical formulae) characterize what conclusions one can draw from lack of information about the truth of certain other formulae.

Shoham [**?**, **?**] and others have proposed definitions of non-monotonic entailment based on a preference ordering of models, where the ordering is a parameter for the entailment relation itself, and therefore is not specified or constrained by the set Γ of axioms. We are here addressing the more general case, where Γ specifies both the models and the preference ordering(s).

Other previous definitions of non-monotonic entailment use auxiliary mechanisms. Circumscription requires the introduction of abnormality predicates, and autoepistemic logic relies on "autoepistemic extensions" which are solutions to a particular equation. We believe that simpler and more direct definitions can be obtained by basing the semantics on *partial models.*

Thomason and Horty [**?**] define non-monotonic entailment using minimal partial models in a manner similar to our approach here. Compared to their work, the present paper adds the inclusion of a preference ordering into the models, so that it can be controlled by the axioms.

Turner [**?**] and Ginsberg [**?**] have also studied non-monotonic logic with unknown as a truth-value, but without defining entailment on that basis.

The present paper defines the logic for the propositional case, and illustrates it with a number of simple examples. Since the criterium for the usefulness of a proposed non-monotonic logic is how well it represents commen-sense reasoning, it is important to have a number of concrete examples.

A previous paper [**?**] is an earlier and less completed variant of the work, but represents the same basic ideas.

2 Definitions

We first re-write the definitions of conventional semantic entailment in a way which is only trivially different from the usual, but which facilitates the transition to the new definitions. If Γ is a set of logical formulae (= wff), and Δ is a set of interpretations, we define

$$\Delta \models \Gamma$$

to hold iff every formula $\alpha \in \Gamma$ has the value T in every interpretation $J \in \Delta$. We define the *model set* of a set of propositions as a set of interpretations in the following slightly roundabout way:

$$Mod(\Gamma) = Maxm_{\subset}(\{\Delta \mid \Delta \models \Gamma\})$$

where $Maxm_{\subset}$ takes a set of sets as argument, and has as value its $\subset$-maximal member. Semantic entailment is then defined as follows, where Π is also a set of logical formulae:

$$\Gamma \models \Pi \;\Leftrightarrow\; Mod(\Gamma) \models \Pi$$

Preferential entailment [**?**] is defined similarly:

$$\Delta \models_{<} \Gamma \;\Leftrightarrow\; Min_{<}(\Delta) \models \Gamma$$

$$\Gamma \models_{<} \Pi \;\Leftrightarrow\; Mod(\Gamma) \models_{<} \Pi$$

where the operation $Min_{<}$ takes a set and returns the set of its $<$-minimal members. Thus unlike $Maxm$ it returns a subset and not a member of the argument set. The relation $<$ is turned in such direction that "smaller" interpretations are preferred.

In the present paper we take two additional steps. First we generalize the domain of Δ to include also *partial interpretations,* which must be domains where one has defined a three-valued evaluation function, so that a logical formula has either of the values T, F, or U in the partial interpretation. The truth-value for a composite formula shall depend on the components according to Kleene's strong definitions for three-valued logic [**?**], i.e. $\wedge$ and $\vee$ are the *min* and *max* operations in an order where F is less than U which is less than T, and $\neg$ reverses T and F.

The partial order $\subset$ is defined for the truthvalues so that $U \subset F$, $U \subset T$, and is extended to interpretations so that $J \subseteq J'$ iff

$$\forall\alpha[J(\alpha) \subseteq J'(\alpha)]$$

It is seen that evaluation of a formula is $\subset$-monotone. We will use preferential entailment with respect to $\subset$, i.e. consider entailments of the form

$$\Gamma \models_{\subset} \Pi$$

which will be called *weakest model entailment.*

As three-valued evaluation is introduced, we also introduce the monary propositional operators L, M, and N defined by the following truth-tables:

α	$\neg\alpha$	$L\alpha$	$M\alpha$	$N\alpha$
T	F	T	T	F
U	U	F	T	T
F	T	F	F	F

For the purpose of the examples in the present paper we will simply use partial interpretations which are mappings from proposition symbols to the three truth-values. It would presumably also be possible to use Kripke style interpretations, and identify the truthvalue T for α with $\neg\diamond\neg\alpha$, F with $\neg\diamond\neg\neg\alpha$, and U with $\neg\neg\diamond\neg\alpha \wedge \diamond\alpha$.

Let Δ be a set of partial interpretations and $J \in \Delta$. J is said to be *maximal in* Δ iff no partial interpretation in Δ is $\supset J$. Also, the *completion of* J *in* Δ is the set of those J' which are maximal in Δ and $\supseteq J$.

Our other step is to combine a partial order $\ll$ with the set of interpretations which occurs in the above definitions. An *aggregate* is defined to be a pair $\langle\Delta, \ll\rangle$ where $\ll$ is a partial order on the maximal members in Δ. For those logical formulas introduced above, i.e. formulas formed using $\wedge$, $\vee$, $\neg$, L, M, N, a formula evaluates in $\langle\Delta, \ll\rangle$ in the same way as in Δ.

However at the same time we introduce an additional propositional operator D, and say that the *value of* $D\alpha$ *in an interpretation* J *in* $\langle\Delta, \ll\rangle$, also written

$$val(D\alpha, J, \langle\Delta, \ll\rangle)$$

where $J \in \Delta$, is T iff α is T in all the $\ll$-minimal completions of J in Δ, and F otherwise. Consequently, with a circle on the entailment symbol which flags the new convention,

$$\langle\Delta, \ll\rangle \models^{\circ} \{\alpha\}$$

iff α is T in every element[1] J of $\langle\Delta, \ll\rangle$, or formally

$$\forall J(J \in \Delta \rightarrow val(\alpha, J, \langle\Delta, \ll\rangle) = T)$$

In accordance with the previous definitions, only provided that we have a good definition for the ordering used in the $Maxm$ operation for aggregates, we immediately have the definition for

$$\Gamma \models^{\circ} \Pi$$

which we call *preferential model entailment* . The definition for $\models^{\circ}_{\subset}$, *preferential weakest model (PWM)* entailment then follows directly although with one important detail: we define $\langle\Delta, \ll\rangle \models^{\circ}_{\subset} \{\alpha\}$ iff

$$\forall J(J \in Min_{\subset}(\Delta) \rightarrow val(\alpha, J, \langle\Delta, \ll\rangle) = T)$$

where the third argument of val shall contain Δ, not $Min_{\subset}(\Delta)$.

The purpose of this paper is to show that common-sense default reasoning is well expressed as entailment under $\models^{\circ}_{\subset}$.

[1]We say that J is an *element* of $\langle\Delta, \ll\rangle$ iff it is a member of Δ.

3 Examples of $\models_{\subset}$ entailment

Before we proceed to default rules and preferential weakest model ($\models^{\circ}_{\subset}$) entailment, let us show a few simple examples of ordinary weakest model entailment (according to $\models_{\subset}$) in three-valued logic. A similar approach is taken by Thomason and Horty in their paper [**?**] in the present volume.

Example 1. $\{a\} \models_{\subset} \{Nb\}$

For the purpose of examples, we write interpretations as strings of truth-values, so for example TF stands for the interpretation where a is T and b is F. The model set for a is $\{TT, TF, TU\}$ and there is only one $\subset$-minimal model namely TU. The value of b in TU is U so the value of Nb is T there. Informally this can be understood as saying "if a is all that is known, then it is not known whether b".

If we now add also b as an axiom, then the model set is reduced to merely $\{TT\}$, and Nb is no longer entailed. Thus $\models_{\subset}$ is non-monotonic with respect to its left argument.

Example 2. $\{a \vee b\} \models_{\subset} \{Ma, Mb\}$

In this case it is not known whether a is T, and it is not known whether b is T; it is just known that $a \vee b$ is T. The first two conclusions can be informally read as "maybe a" and "maybe b". Contrary to what one would expect, maybe, Na and Nb are not entailed. This is because the set of models is

$$\{TT, TF, FT, TU, UT\}$$

and does not include UU. The set of minimal models is $\{TU, UT\}$ and in order to be entailed, a formula must be T in both of those. Therefore Na is not entailed, but notice, neither is $\neg Na$.

This example also shows that the formula

$$\alpha \vee \neg\alpha \vee N\alpha$$

is <u>not</u> T in every interpretation.

4 Entailment using $\models^{\circ}_{\subset}$

Let us now proceed to preferential weakest model entailment using $\models^{\circ}_{\subset}$ and the full range of definitions in section 2. We shall show how to represent the classical examples of default rules in that logic.

Single default. Let us first consider the simplest possible example, namely the default rule "unless a is known to be false, assume that it is true".

The correct way to express that rule for PWM is as a formula

$$Ma \rightarrow Da$$

which can mnemonically be read as "if maybe a, then default a". This formula is T in all elements of the following aggregates

$$\langle\{T, F, U\}, [T \ll F]\rangle$$

where $[T \ll F]$ stands for the (weakest) ordering $\ll$ which satisfies the conditions in the brackets, and in

$$\langle\{T, F\}, [\,]\rangle$$

$$\langle\{T, U\}, [\,]\rangle$$

plus of course all aggregates obtained from either of these three by reducing the Δ component to a subset. The ordering which is used in the $Maxm$ operation, in the definition of the model set $Mod(\Delta)$ in section 2, has not yet been defined, but we see now that it must be such that the first one of these aggregates is preferred. We can see it as a meta level preference ordering. It maximizes primarily the set Δ, and secondarily minimizes the strength of the $\ll$ ordering. (More about this in section 5).

With these criteria, the aggregate

$$\langle\{T, F, U\}, [\,]\rangle$$

would be even more preferable, but its element U does not satisfy the axiom, since Ma is true[2] in the interpretation [where the value of a is] U; the set of $\ll$-minimal completions of U is $\{T, F\}$, a does not have the value true in the interpretation F, and therefore Da is not true in the element U. Therefore the aggregate with a stronger $\ll$ ordering where $T \ll F$, is the best one possible.

The following entailment is therefore obtained:

$$\{Ma \to Da\} \models^{\circ}_{\subset} \{Na, Da\}$$

which can be read to say that we do not know whether a, but we have it by default. If we add the additional fact $\neg a$ we have

$$\{Ma \to Da, \neg a\} \models^{\circ}_{\subset} \{\neg Na, \neg Da\}$$

since the model aggregate is

$$\langle\{F\}, [\,]\rangle$$

One may wonder why the rule has to be written with the Ma antecedent. The reason why we can not write the rule simply as Da, is that Da is F in the interpretation F. Therefore if Da is one member of an axiom set, then the interpretation F will never be an element in a model aggregate. The default rule together with information that overrides it, for example $\{Da, \neg a\}$, does not have any model then.

Implied default. Next we consider rules of the form "unless a, which is unusual, conclude b as default". Such a rule should be written as two axioms

$$M\neg a \to D\neg a$$

$$L\neg a \wedge Mb \to Db$$

and obtains the following model aggregate

$$\langle \mathbf{B} \times \mathbf{B}, [FT \ll TT, FT \ll FF \ll TF]\rangle$$

[2]In this particular paragraph we write the truth-values as "true" and "false" rather than T and F, in order to avoid confusion with the single-component interpretations being used here.

where $\mathbf{B}$ is the set $\{T, F, U\}$ of all three truth-values. We should expect that no interpretation is ruled out from the Δ component of the aggregate, since we have only default rules so nothing has been definitely excluded.

UU is here the $\subset$-minimal interpretation, entailing $D\neg a \wedge Db$. As we should expect, if we add a as a second axiom we have the model set $\{TT, TF, TU\}$ and no preference, so the conclusion Db is lost. Also if we instead add the axiom $\neg b$ the conclusion $D\neg a$ is retained as it should. Neither addition causes any contradiction, of course.

The royal elephant. This is one of the classical toy examples. Consider the statements "royal elephants are typically albinos" and "if an elephant is not an albino, then it is typically gray". We represent the proposition "the elephant at hand is a royal one" as the propositions symbol r, and similarly use a for albino and g for gray.

The two rules now have the general form

$$r \to a$$

$$\neg a \to g$$

where the exact interpretation of the $\to$ symbol remains to be specified. In spite of the apparent similarity, there is a significant difference between these two rules, with respect to how the antecedents are defaulted. For the first rule, the default is that the antecedent does not hold, since most elephants are not royal ones, whereas for the second rule the antecedent holds by default. We may of course also encounter rules where the antecedent doesn't default either way.

In accordance with the solution for the previous example, we write the rules for this example as follows:

$$M\neg r \to D\neg r$$

$$Lr \wedge Ma \to Da$$

$$\neg Lr \wedge M\neg a \to D\neg a$$

$$L\neg a \wedge Mg \to Dg$$

The condition $\neg Lr$ in the third axiom is the specificity constraint; without it we would have two conflicting defaults if r is T. Let us now consider what model aggregates we obtain for these axioms, using the entailment defined above. If we take only the first three axioms, and consider interpretations over the proposition symbols r and a (in that order), we obtain the following model aggregate

$$\langle \mathbf{B} \times \mathbf{B},\ [FF \ll FT \ll TT \ll TF] \rangle$$

The ordering can not be weaker; we require its elements for the following reasons:

$FF \ll FT$, so that FU satisfies axiom 3

$FT \ll TT$ so that UT satisfies axiom 1

$TT \ll TF$, so that TU satisfies axiom 2.

It is also readily verified that all J in the model set satsify all axioms 1-3 in the model set.

Since UU is the only $\ll$-minimal element of the model aggregate, the set Γ of the first three axioms satisfies

$$\Gamma \models^{\circ}_{\subset} \{Nr, Na, D\neg r, D\neg a\}$$

as they should do: the elephant at hand is by default neither royal nor an albino.

If we add the axiom r saying that the elephant is royal, then the extended set of axioms has the model set

$$\langle \{TT, TF, TU\}, [TT \ll TF] \rangle$$

and $\models^{\circ}_{\subset}$-entails $\{Na, Da\}$ i.e. it is not known, but assumed as default that the elephant is an albino.

When the fourth axiom is also taken into account, we have one more proposition symbol, so each of the maximal interpretations splits into two "descendants". Clearly the $\ll$ relations that we had before are "inherited" pairwise by the descendants, so the first three axioms are satisfied iff we have the ordering

$$[[FFT, FFF] \ll [FTT, FTF] \ll [TTT, TTF] \ll [TFT, TFF]]$$

where $[J_1, J_2] \ll [J_3, J_4]$ means that

$$J_1 \ll J_3 \wedge J_2 \ll J_4$$

The first three axioms do not require any extra cross-relationships such as $FFT \ll FTF$.

It remains to check out whether the fourth axiom requires any strengthenings of the order within any of these sets, and whether it contradicts the ordering required by the first three axioms. Interpretations where the third element g is either T or F satisfy axiom 4 regardless of $\ll$. Interpretations where the second element is T or U also satisfy the axiom trivially. The remaining interpretations, i.e. TFU, FFU, and UFU require the ordering to be such that

$$[TFT \ll TFF, FFT \ll FFF]$$

Thus the model aggregate for all four axioms is

$$\langle \mathbf{B} \times \mathbf{B} \times \mathbf{B},\ [FFT \ll FTT \ll TTT \ll TFT \ll TFF, \\ FFT \ll FFF \ll FTF \ll TTF \ll TFF] \rangle$$

where the $\subset$-minimal element is UUU. The model aggregate therefore $\models^{\circ}_{\subset}$-entails that the elephant at hand is not royal, not an albino, and gray. It is easy to verify how additional axioms that override the default will remove some of the elements from the model aggregate, without affecting the ordering of the remaining elements.

Interlocked defaults. This was the original problematic example of non-monotonic logic, where it was first observed that default reasoning may give multiple extensions (Sandewall, [**?**]). It has the general form

$$Unless\ a \rightarrow b$$

$$Unless\ b \rightarrow a$$

and is expected to give two separate "extensions" on the procedural grounds that a reasoning system could either use the first rule first, and conclude b since a is not (yet) known, but then after that it could not infer a by the second rule since the condition there is violated; or else it takes the rules in the other order and obtains a but not b as a "theorem".

For PWM entailment we would like to write the rules as follows:

$$M\neg a \rightarrow D\neg a$$

$$La \wedge Mb \rightarrow Db$$

$$M\neg b \rightarrow D\neg b$$

$$Lb \wedge Ma \rightarrow Da$$

However here again there is a need for specificity constraints, like in the previous example, in order to avoid conflicting defaults for one single variable. Inserting specificity constraints we obtain

$$\neg Lb \wedge M\neg a \rightarrow D\neg a$$

$$La \wedge Mb \rightarrow Db$$

$$\neg La \wedge M\neg b \rightarrow D\neg b$$

$$Lb \wedge Ma \rightarrow Da$$

We obtain the following model aggregate

$$\langle \mathbf{B} \times \mathbf{B}, [FF \ll TT \ll \{TF, FT\}] \rangle$$

with UU as the $\subset$-minimal model, and FF as its only $\ll$-minimal completion. Surprisingly, therefore, the axioms in this example entail $D(\neg a \wedge \neg b)$, and we have no trace of the two extensions. This could be taken as an indication that the lack of (perceived) persistency constraints was the only problem with the original formula.

Nixon diamond. Yet another classical example besides the royal elephant regards Nixon who is both a republican and a quaker. For the purpose of the example quakers are said to generally be pacifists, and republicans to generally not be pacifists. We represent "Nixon is a republican" by r, "Nixon is a quaker" by q, and "Nixon is a pacifist" with p, and write the four statements in the diamond as

$$q$$

$$r$$

$$Lq \wedge Mp \rightarrow Dp$$

$$Lr \wedge M\neg p \rightarrow D\neg p$$

Of course in a more general setting one would have something like $q(n)$, $r(n)$, and general rules of the form

$$\forall j[Lr(j) \wedge M\neg p(j) \rightarrow D\neg p(j)]$$

but we stay with the propositional case for the example. The first two axioms rule out all interpretations except those where both q and r are T, so we only have the consider the three interpretations for the proposition symbol p. The model aggregate is then

$$\langle \{T, F\}, [\,] \rangle$$

The interpretation U can not be included because the last two axioms would then require $T \ll F$ and $F \ll T$. Therefore no interesting conclusion is entailed by these axioms.

Competing defaults. Consider the rules “unless a conclude b” and “unless b conclude c”. The combination of these two rules is similar to the Nixon diamond in that it is somewhat debatable whether they make practical sense: one may interpret the first rule as implicitly saying “usually b”, and the other as implicitly saying “usually not b”. Still (or therefore) it is interesting to figure out which conclusions are formally entailed.

We read both rules as defaulting to false antecedent, so the first rule says “usually not a, and if not a then by default b”. Informally, we might expect to obtain two extensions, by the following argument: in one extension we assume $\neg a$ by default, and obtain b as a default consequence. In the other extension we assume $\neg b$, conclude a (by using the first rule backwards), and also conclude c (by the second rule). These two expected extensions can be informally characterized as FTU and TFT.

Formally we write these rules as

$$M\neg a \rightarrow D\neg a$$

$$L\neg a \wedge Mb \rightarrow Db$$

$$M\neg b \rightarrow D\neg b$$

$$L\neg b \wedge Mc \rightarrow Dc$$

Again the need for a specificity constraint comes up, since Mb and $M\neg b$ are both satisfied in partial interpretations where b is U, so we extend the third axioms to read

$$\neg L\neg a \wedge M\neg b \rightarrow D\neg b$$

The first three axioms can be analyzed in terms of the nine two-element interpretations, and have the model structure where

$$FT \ll FF \ll TF \ll TT$$

When also c and the fourth axiom are added, we obtain two parallel sequences like in the gray elephant example:

$$FTT \ll FFT \ll TFT \ll TTT$$

$$FTF \ll FFF \ll TFF \ll TTF$$

and in addition the requirements $FFT \ll FFF$, $TFT \ll TFF$. The interpretation UUU is $\subset$-minimal, and has two $\ll$-minimal completions FTT and FTF. Thus the given axioms entail

$$Na \wedge Nb \wedge Nc \wedge D\neg a \wedge Db$$

Of course without the specificity constraints we get another analysis.

5 Aggregate preference

The meta-preference relation $<$ is a necessary part of the definition of the model set, as we saw in section 2. In this section we discuss the exact definition of that ordering. We shall show that, unlike in the classical case, we can not have any reasonable criterium whereby there is a single, preferred model aggregate. Therefore the formal definitions of entailment in section 2 must be revised to handle the multiple model aggregates correctly.

First of all, we turn the meta-preference relation in the same direction as the other ones, so that "smaller" aggregates are preferred over "larger" ones. We shall write this relation using the symbol $<$. Clearly it must satisfy

$$\ll \subset \ll' \rightarrow \langle \Delta, \ll \rangle < \langle \Delta, \ll' \rangle$$

so that for equal Δ we take the weakest possible ordering, and

$$\Delta \supset \Delta' \rightarrow \langle \Delta, \ll \rangle < \langle \Delta', \ll' \rangle$$

so that larger sets Δ are always preferred, regardless of the $\ll$ ordering.

Unfortunately there does not seem to be any reasonable way of defining the ordering so that the minimum is always unique. Consider for example the following four axioms

$$Na \wedge Lb \rightarrow Da$$

$$Na \wedge L\neg b \rightarrow D\neg a$$

$$Nb \wedge La \rightarrow D\neg b$$

$$Nb \wedge L\neg a \rightarrow Db$$

In each of the axioms, the antecedent is only T in one single interpretation in $\mathbf{B} \times \mathbf{B}$, for example the antecedent of the first axiom is only satisfied in UT. The four maximal interpretations TT, TF, FT, FF satisfy all the axioms. On the next level down, each of the four interpretations with one U in it will imply a $\ll$ ordering between its two completions, and do so in a fashion which creates a circularity, i.e.

$$TT \ll FT \ll FF \ll TF \ll TT$$

It is therefore necessary to omit at least some of these eight interpretations from the model aggregate, and since the interpretation preferences are entirely symmetric there is no rule that could distinguish between them, as long as T and F are treated symmetrically.

It may be reasonable to have a "stratification" rule which maximizes $\subset$-stronger interpretations with higher priority, although the examples we have described here do not give any indication whether that is a good idea from a common sense point of view, or even whether it makes any difference. However such a stratification strategy does not solve the problem for the example just mentioned.

We therefore require from the meta-preference ordering $<$ on aggregates that it at least satisfies the two conditions stated at the beginning of this section. Also the definitions of entailment which were given in section 2 must now be technically revised as follows: we define $Mods(\Gamma)$ to be a set of sets of interpretations, as follows:

$$Mods(\Gamma) = Min_{<}(\{\Delta \mid \Delta \models \Gamma\})$$

For classical entailment, $<$ is the same as $\supset$, and $Mods(\Gamma)$ is the singleton set of the set of models for Γ. In the case of $\models^{\circ}_{\subset}$, $Mods(\Gamma)$ is the set of $<$-minimal aggregates, consisting usually of one, but as we have seen occasionally of a larger number of aggregates, each being a pair $\langle\Delta, \ll\rangle$.

Then entailment from sets of formulas to sets of formulas is defined straightforwardly as

$$\Gamma \models \Pi \;\Leftrightarrow\; [\Lambda \in Mods(\Gamma) \Rightarrow \Lambda \models \Pi]$$

and the definition of $\models^{\circ}_{\subset}$ is also updated accordingly.

6 Abstract properties of NME logic

The $\models^{\circ}_{\subset}$ entailment defined here does not have the same abstract properties as ordinary entailment. This is not to say that it is entirely misbehaved; it does entail its own axioms; conclusion sets may be joined and split; and modus ponens is sound:

$$\Gamma \cup \Pi \models^{\circ}_{\subset} \Gamma$$

$$\Gamma \models^{\circ}_{\subset} \Pi \cup \Sigma \Leftrightarrow [\Gamma \models^{\circ}_{\subset} \Pi \wedge \mathit{\Gamma} \models^{\circ}_{\subset} \mathit{\Sigma}]$$

$$\Gamma \models^{\circ}_{\subset} \{\alpha, \alpha \rightarrow \beta\} \Rightarrow \Gamma \models^{\circ}_{\subset} \{\beta\}$$

However it also has the following less standard properties, besides being non-monotonic.

Cumulative monotony [**?**, **?**] is not satisfied: if $\Gamma \models^{\circ}_{\subset} \Pi \cup \Sigma$ we can not in general conclude $\Gamma \cup \Pi \models^{\circ}_{\subset} \Sigma$

Counterexample:

$$\{Ma \rightarrow Da\} \models^{\circ}_{\subset} \{Na, Da\}$$

$$\{Ma \rightarrow Da, Na\} \not\models^{\circ}_{\subset} \{Da\}$$

Cumulative transitivity [**?**, **?**] is not satsified either: if $\Gamma \models^{\circ}_{\subset} \Pi$ and $\Gamma \cup \Pi \models^{\circ}_{\subset} \Sigma$, we can not in general conclude $\Gamma \models^{\circ}_{\subset} \Sigma$. Counterexample:

$$\{a\} \models^{\circ}_{\subset} \{Nb\}$$

$$\{a, Nb\} \models^{\circ}_{\subset} \{\neg D(b \vee \neg b)\}$$

$$\{a\} \not\models^{\circ}_{\subset} \{\neg D(b \vee \neg b\}$$

because $\{a\}$ has the model set $\{TT, TF, TU\}$ and $\{a, Nb\}$ has the model set $\{TU\}$.

Tautologies are not entailed: if α is a tautology in the sense of conventional propositional logic, we can not in general conclude $\Gamma \models^{\circ}_{\subset} \{\alpha\}$, nor even $\Gamma \models^{\circ}_{\subset} \{D\alpha\}$.

Counterexample: $\Gamma = \{Na\}$, $\alpha = a \vee \neg a$.

Conclusions not always consistent: if $\Gamma \models^{\circ}_{\subset} \Pi$ and Γ is consistent (i.e. has a non-empty model aggregate), then we can not in general conclude that $\Gamma \cup \Pi$ is consistent.

Counterexample:

$$\{Ma \rightarrow Da\} \models^{\circ}_{\subset} \{Na\}$$

but $\{Ma \rightarrow Da, Na\}$ has only the model aggregate $\langle\{\}, []\rangle$

All the counterexamples to the conventional properties however involve using formulas which are different in syntax and style from those used for expressing default rules. The language of this logic is therefore considerably more general than what default rules require, and it is when the expressiveness of the logic is exploited that the non-standard behavior is obtained.

These observations suggest that it could be worthwhile to look for syntactic or other constraints on the formulas which would recover at least some of the conventional properties, for example cumulative monotony and transitivity. Also one would hope to regain uniqueness of the $<$-best model aggregate as a result of these syntactic constraints.

7 Conclusion

We have described a logic for non-monotonic entailment based on partial models, and showed that it can express default rules and obtain the common-sense conclusions from them. We intend to analyze a number of additional examples.

A major intended application area for this logic is in the context of non-monotonic reasoning about time and action, with the approach which is presented in another current paper, [**?**].