DIRECTLY SCEPTICAL INHERITANCE CANNOT CAPTURE THE INTERSECTION OF EXTENSIONS

Karl Schlechta
DFKI, Stuhlsatzenhausweg 3
D-6600 Saarbrücken
Germany

13.8.89, revised 6.12.92

Abstract

We show that, under some very weak assumptions about the definitions of sceptical and extension-based defeasible inheritance, directly sceptical inheritance cannot capture the intersection of extensions.

1 INTRODUCTION

The reader less familiar with defeasible inheritance is referred to the end of this introduction, where we repeat some basic definitions, intuitive interpretations and give a sample "directly sceptical" and extensions-based approach.

Throughout, all nets considered will be finite and acyclic. All diagrams are collected in an Appendix and referred to as Ex. ....

The main distinction of defeasible inheritance formalisms is between extensions-based and directly sceptical approaches. As a first approximation, an extension is a suitably chosen maximal consistent subset of the given information, whereas direct scepticism admits only (preferred) uncontested information. More precisely, defeasible inheritance systems may contain conflicting information. Some such conflicts can be resolved by a criterion in terms of specificity of preference of one information over the other (formally, by "preclusion", also called "preemption") while others are unresolvable by such preference. Both the directly sceptical and the extensions approach will treat the first type of conflict in the same way, choosing the preferred information. However, they differ on the second: a directly sceptical approach will come to no conclusion, and admits neither the positive nor the negative information as valid. The extensions approach will branch at this point, constructing two different extensions, one where the positive information is considered as valid, the other where the negative information is accepted. The extensions approach can, in a second step, be turned into an (indirectly) sceptical one by considering only the information contained in all extensions.

Thus, the problem we give an answer to lies at the very heart of inheritance theory: Can the two basic approaches be made equivalent - possibly through major modifications of existing formalisms, whilst, of course, preserving the basic spirit? Or is there an essential difference between cautiousness while performing an inductive construction and cautiousness in regarding the results of the completed construction? We show the latter.

It should be emphasized that our aim is not to show that some particular two definitions, or two narrow classes thereof, one via extensions, the other by direct scepticism, are not equivalent. For such results see e.g. [Ste91a], [Ste91b] and the author's [Sch89A]. Rather, we prove a generic result: wide classes of extension-based and directly sceptical approaches do not contain equivalent definitions. In particular, we show that an inflation of truth values, which can remedy some problems, among them the Double Diamond (Ex.4) - essentially by introducing paths which are not valid in the positive sense, but which still preserve destructive potential and therefore called "Zombie Paths" in [MS91] - cannot restore equivalence in general. Our counterexample is of a generic nature too: We give a construction schema which produces for every finite number of truth values a suitable counterexample. Such a schema is necessary because any single such example could still be overcome, essentially through encoding in a manner sketched at the end of this paper.

Historical Background On the Double Diamond (Ex.4), the sceptical definition of e.g. [HTT87] - which is essentially repeated below in the introduction - fails to give the same results as the intersection of extensions (this observation is due to [HTT87]). Recall that the path $a \rightarrow b \rightarrow d \rightarrow f$ is valid on the first, but not the second account. It is easy to find alternative sceptical definitions which solve the Double Diamond correctly in this sense (see e.g. [GV89], [Sch89A]); but they fail on other diagrams. So the question naturally arises whether it is possible to find a sceptical definition of inheritance which gives the same results as the intersection of extensions on all diagrams. This question has already been asked in [THT87], and also in [GV89]. The diagram of alternating preclusions (Ex.5a for off-path preclusion, and Ex.5b for on-path preclusion) led the author to the conjecture made precise and proved below as the Theorem, that no sceptical definition of inheritance, however modified, can ever give the same results as the intersection of extensions approach on all diagrams. (In our proof of the Theorem below, we will use another technique - essentially disjunction instead of negation - to keep the assumptions about the definition of sceptical inheritance weak.) The importance of these diagrams is that (e.g. in Ex.5a) we have a Nixon-Diamond in $a \rightarrow d \rightarrow g/a \rightarrow e \not\rightarrow g$: In an extensions approach, we can decide here for the positive or the negative possibility. If we take the positive choice, $a \rightarrow d \rightarrow g \rightarrow l \rightarrow k$ and $l \not\rightarrow m$ is a preclusion of $a \rightarrow i \rightarrow k \rightarrow m$, but the other possible preclusion $a \rightarrow f \rightarrow h \rightarrow i$ and $h \not\rightarrow k$ is itself precluded by $a \rightarrow d \rightarrow g \rightarrow f$ and $g \not\rightarrow h$. On the other hand, in the negative choice, $a \rightarrow d \rightarrow g$ does not exist, so neither does $a \rightarrow d \rightarrow g \rightarrow l \rightarrow k$, but $a \rightarrow f \rightarrow h \rightarrow i$ with $h \not\rightarrow k$ is now a preclusion. So, in either choice (=extension) the path $a \rightarrow i \rightarrow k \rightarrow m$ is destroyed by preclusion, but not by a common one. Consequently, in all extensions, we do not accept a positive path from a to m, in one extension, we have a valid negative path from a via l to m, in the other extension, we have no valid negative path either. In a directly sceptical approach, however, we will have no valid path from a to g, and have a situation similar to the negative choice above: $a \rightarrow d \rightarrow g \rightarrow l \rightarrow k$ does not exist, but $a \rightarrow f \rightarrow h \rightarrow i$ with $h \not\rightarrow k$ is a preclusion. So, again, the path $a \rightarrow i \rightarrow k \rightarrow m$ is destroyed by preclusion, and we have no valid paths, negative or positive, from a to m.

Example Ex.5a involves "floating conclusions". They were first noted by L.Stein ([Ste89], [Ste91a], [Ste91b]), see also [MS91] for a discussion of their significance.

Relevance of the Question for Inheritance Theory and Beyond In an attempt to find suitable standards of adequacy of sceptical inheritance reasoning, we can, in a first try, present examples (labels) for our diagrams and show that the formalism solves them correctly or not. But, everyone can produce examples for and against such formalisms *ad libitum*, by the mere fact that any real-life example carries with it a load of unexpressed information, which makes the solution plausible or not. But also all semantics of inheritance systems known to the author (see [KK89], [KKW89a], [KKW89b], [Sch89A]) make themselves use of the underlying inheritance principles, so they cannot decide between the different approaches. Thus, given a set of principles like contradiction, preclusion (specificity), it seems natural to let the extensions defined using those principles decide the adequacy of a directly sceptical approach: In the Double Diamond example (Ex.4) e.g., the "directly sceptical" definition as in [HTT87] gives a positive path $a \to b \to d \to f$ as valid, whereas we have a contradictory *possibility* $a \to b \to e \not\to f$, i.e. an *extension*, where the opposite is true, thus the result of the "sceptical" approach is doubtful, and not sceptical in the intuitive sense of the word. Thus, our result tells us that we can never expect to fully meet the intuition - if defined via the result of the extensions approach - with sceptical inheritance definitions.

We would also like to emphasize that our result works in the very limited language of inheritance systems, and is thus different from related specific ones for e.g. logic programming. The question discussed here also has relevance beyond inheritance theory. It arises in the more general context of defeasible argumentations. As an example, we can refer the reader to the discussion by Nelson Goodman of the "grue emerald paradox" and his attempted solution, see [Goo55].

Basic definitions, intuitive interpretations, and approaches - a brief review A non-monotonic or defeasible inheritance net is a finite, directed, acyclic graph, whose nodes stand for elements or sets of elements, with two types of arrows, negative and positive ones, which are interpreted roughly as: $a \to A$: a is an element of A $a \not\to A$: a is not an element of A $A \to B$: almost all elements of A are in B $A \not\to B$: almost all elements of A are not in B All important questions can be discussed in diagrams consisting only of set type nodes, we will, however, give here the standard examples which also have element type nodes (and strict links).

The basic situations of conflicting information are exemplified by diagrams Ex.1 and Ex.2.

A common labelling of Ex.1 is: A=Nixon, B=Quaker, C=Republican, D=Pacifist, Nixon is a republican and Quaker, Quakers are usually pacifists, republicans usually not. This diagram is therefore called the Nixon Diamond. A common labelling of Ex.2 is: A=Tweety, B=Bird, C=Penguin, D=Flying thing, Tweety is a penguin and a bird, normally, penguins can't fly, normally, birds can fly, penguins are birds. One might call it the "Tweety Preclusion Diagram".

In the first case, using only the information contained in the diagram, we do not know how to decide whether Nixon is a pacifist or not, as the contradicting arguments have the same strength. In the second case, we have the additional information that penguins are birds, and the basic assumption is that the more specific information, i.e. that concerning the subclass penguins, is the more reliable one. This is also, roughly, intuitively correct. We will thus conclude here that Tweety can't fly.

More formally, a defeasible inheritance net permits a very limited - in contrast to standard logics, where there are usually infinitely many ways to prove a result - number of potential lines of reasoning or arguments by concatenating arrows to paths and drawing conclusions, like $A \to B$ and $B \to D$ to $A \to B \to D$, which thus stands for the argument: "Most elements of A are in B, most elements of B are in D, thus - probably, or defeasibly - most elements of A are in D." In Ex.1, however, we have the contradicting argument $A \to C \not\to D$ of equal strength, so the arguments $A \to B \to D$ and $A \to C \not\to D$ and thus the conclusions "Most A are in D."/"Most A are not in D." are in unresolvable conflict. There are two basic ways to proceed: The first, the "directly sceptical" one, will not accept any argument relating A to D, the second, the "credulous" one, will split into two "extensions" or "belief sets", one where $A \to B \to D$ is accepted, the other where $A \to C \not\to D$ is believed. The "indirectly sceptical" approach will, in a second step, either accept the *arguments* which are in all such extensions, or the resulting *conclusions*, which are derived by some argument in all extensions (see [MS91] for a discussion). In our example Ex.1, all sceptical approaches accept no path from A to D, and no conclusion linking A and D. In Ex.2, all three approaches will also agree in accepting as valid arguments $A \to C \not\to D$ (and $A \to C \to B$), as the contradicting one $A \to B \to D$ will be considered weaker, there will be only one extension, and the conclusion "Most A are not in D." will be considered valid in all approaches.

We now repeat the exemplary formal definition of validity of paths of [HTT87], (in slight refinement: our conditions II.1(c) and II.2(c) are new) and show how to transform it into an (intuitively) analogous extensions approach.

Fix now an inheritance diagram Γ, $x \to p \in \Gamma$ shall express that $x \to p$ is an arrow in Γ etc. We assume for simplicity that Γ contains no direct contradictions in the form of $u \to v \in \Gamma$ and $u \not\to v \in \Gamma$, and all nodes are of set type.

For two paths $\sigma : x \cdots > y$, $\tau : y \cdots > z$, $\sigma \circ \tau : x \cdots > y \cdots > z$ shall denote the obvious concatenation.

Definition 1.1 *Generalized and Potential Paths:*

a) Generalized paths: If $x \to p \in \Gamma$, then $x \to p$ is a generalized path. If $x \not\to p \in \Gamma$, then $x \not\to p$ is a generalized path. If $x \cdots\cdots > p$ is a generalized path, and $p \to q \in \Gamma$, then $x \cdots\cdots > p \to q$ is a generalized path. If $x \cdots\cdots > p$ is a generalized path, and $p \not\to q \in \Gamma$, then $x \cdots\cdots > p \not\to q$ is a generalized path.

b) Potential paths: If $x \to p \in \Gamma$, then $x \to p$ is a positive potential path (pp.). If $x \not\to p \in \Gamma$, then $x \not\to p$ is a negative pp. If $x \cdots \to p$ is a positive pp., and $p \to q \in \Gamma$, then $x \cdots \to p \to q$ is a positive pp. If $x \cdots \to p$ is a positive pp., and $p \not\to q \in \Gamma$, then $x \cdots \to p \not\to q$ is a negative pp. $\Box$

Thus a generalized path is a potential one iff all but possibly the last arrow are positive. It is easy to see by our intuitive reading that only potential paths can correspond to reasonable arguments.

Up to the end of this Section, Definition 1.5, we fix some node x in Γ and define the valid paths and preclusions starting in x for the directly sceptical and the extensions approach. We start by the auxiliary notion of degree (relative to x), a mapping f from the vertices to natural numbers s.th. $p \to q$ or $p \not\to q \in \Gamma$ implies $f(p) < f(q)$.

Definition 1.2 *Degree:*

Let σ be a generalized path beginning in x, ending in y, then $deg_\Gamma(\sigma) := deg_\Gamma(y)$:= the maximal length of any generalized path beginning in x and ending in y. $\Box$

Let now some enumeration ρ of all nodes y of Γ for which there exists a generalized path $\sigma : x \cdots > y$ be given, s.th. ρ respects deg_Γ, i.e. $deg_\Gamma(y) < deg_\Gamma(y') \rightarrow \rho(y) < \rho(y')$. (If $deg_\Gamma(y) = deg_\Gamma(y')$, then the order of $\rho(y)$ and $\rho(y')$ does not matter.) All inductions are now performed wrt. this enumeration ρ.

Definition 1.3 *Directly Skeptical Approach, Preclusion and Validity:*

We define by simultaneous induction validity of paths and results in Γ, denoted $\Gamma \models \sigma$ ($\Gamma \models xy / \Gamma \models x\overline{y}$), and preclusions in Γ. Assume that y is the first node in the enumeration ρ which has not yet been treated.

1) Preclusion in Γ

A pp. $\sigma : x \cdots > y$ is said to be precluded in Γ iff

(a) $length(\sigma) > 1$ and

(b) Case 1, σ is positive:

either $x \not\rightarrow y \in \Gamma$ or

there is u s.th. $\sigma : x \cdots \rightarrow u \rightarrow y$ and a positive pp. $\tau : x \cdots \rightarrow z \cdots \rightarrow u$ for some z s.th.

(1) $\Gamma \models \tau$ (thus, we accept that z is more specific than u, as seen from x - note that $\Gamma \models \tau$ has already been decided by induction hypothesis)

(2) $z \not\rightarrow y \in \Gamma$ (we have a contradiction, which, by (1) is stronger than $u \rightarrow y$).

Case 2, σ is negative: analogous, i.e.

either $x \rightarrow y \in \Gamma$ or

there is u s.th. $\sigma : x \cdots \rightarrow u \not\rightarrow y$ and a positive pp. $\tau : x \cdots \rightarrow z \cdots \rightarrow u$ for some z s.th.

(1) $\Gamma \models \tau$

(2) $z \rightarrow y \in \Gamma$.

2) Validity of Paths:

Let $\sigma : x \cdots > y$ *be a potential path. (Valid paths have to be potential paths.)*

Case I:

σ *is a direct link in* Γ. *Then* $\Gamma \models \sigma$.

Case II:

σ *is a potential path of length* > 1. *(Note that by induction hypothesis,* $\Gamma \models \tau$ *is decided for all* τ *with degree less than* $deg_{\Gamma}(\sigma)$.*)*

Case II.1:

σ *is a positive pp.* $x \cdots \to u \to y$, *let* $\tau : x \cdots \to u$ *be its initial segment.*

Then $\Gamma \models \sigma$ *iff*

(a) $\Gamma \models \tau$ *(The initial segment must be a valid path - decided already, by induction hypothesis.)*

(b) $u \to y \in \Gamma$

(c) σ *is not precluded in* Γ

(d) if there is a positive pp. $\tau' := x \cdots \to v$ *s.th.* $\Gamma \models \tau'$ *and* $v \not\to y \in \Gamma$, *then* $\sigma' := \tau' \circ v \not\to y$ *is precluded in* Γ *(Preclusion of conflicting paths).*

Diagram Ex.3 illustrates preclusion of a conflicting path.

Case II.2: The negative case is entirely symmetrical.

σ *is a negative pp.* $x \cdots \to u \not\to y$, *let again* $\tau : x \cdots \to u$ *be its initial segment. Then* $\Gamma \models \sigma$ *iff*

(a) $\Gamma \models \tau$

(b) $u \not\to y \in \Gamma$

(c) σ *is not precluded in* Γ

(d) if there is a positive pp. $\tau' := x \cdots \to v$ *s.th.* $\Gamma \models \tau'$ *and* $v \to y \in \Gamma$, *then* $\sigma' := \tau' \circ v \to y$ *is precluded in* Γ

3) Validity of Results:

Finally, we define $\Gamma \models xy$ *iff there is a positive pp.* $\sigma : x \cdots \to y$ *s.th.*

$\Gamma \models \sigma$, *likewise* $\Gamma \models x\overline{y}$, *iff there is a negative pp.* $\sigma : x \cdots \not\to y$ *s.th.* $\Gamma \models \sigma$.
□

Note that by II.1(a) and II.2(a), this formalism works by upward chaining of arguments, there are other formalisms, which work downwards. For more information, the reader is referred e.g. to [Sch92A].

We now transform this approach into an analogous one via extensions (this is not included in [HTT87], but our ad hoc definition). We inductively construct the extensions of Γ, which are sets of paths. For uniformity of notation, we shall continue to write $E \models \sigma$ for $\sigma \in E$, if σ is a path in the extension E.

Definition 1.4 *We start with one empty extension,* $\mathcal{E} := \{\emptyset\}$. *The construction proceeds in two directions: First, each extension constructed so far grows by adding new paths as valid, second, the number of extension may grow as the construction splits one extension into two differently developing cases.*

We assume to have (partially) constructed $\mathcal{E} := \{E_i : i \in I\}$ *so far, by considering all nodes prior to* y *in the enumeration* ρ. *The following operation for* y *will be performed simultaneously (or successively) on all* E_i, *so we fix one such partially constructed* $E := E_i$.

Again, we define by simultaneous induction preclusion and validity of paths. Many comments carry over from the directly sceptical case, we do not repeat them here.

1) Preclusion in Extension (analogous to preclusion in Γ*):*

A pp. $\sigma : x \cdots > y$ *is said to be precluded in E iff*

(a) length$(\sigma) > 1$ *and*

(b) Case 1, σ *is positive:*

either $x \not\to y \in \Gamma$ *or*

there is u s.th. $\sigma : x \cdots \to u \to y$ *and a positive pp.* $\tau : x \cdots \to z \cdots \to u$ *for some z s.th.*

(1) $E \models \tau$

(2) $z \not\to y \in \Gamma$

Case 2, σ *is negative: analogous, i.e.*

either $x \to y \in \Gamma$ *or*

there is u *s.th.* $\sigma : x \cdots \to u \not\to y$ *and a positive pp.* $\tau : x \cdots \to z \cdots \to u$ *for some* z *s.th.*

(1) $E \models \tau$

(2) $z \to y \in \Gamma$

2) Validity of Paths in Extensions

Case I (Direct Links):

If σ *is a direct link in* Γ, *then* $E \models \sigma$

Case II:

Potential paths of length > 1.

Let

$< x, y >^+ :=$ { $< \tau : x \cdots \to u, u \to y >$: *(a)* τ *is a positive pp., (b)* $E \models \tau$, *(c)* $u \to y \in \Gamma$, *(d)* $\tau \circ u \to y$ *is not precluded in* E }

and

$< x, y >^- :=$ { $< \tau : x \cdots \to u, u \not\to y >$: *(a)* τ *is a positive pp., (b)* $E \models \tau$, *(c)* $u \not\to y \in \Gamma$, *(d)* $\tau \circ u \not\to y$ *is not precluded in* E }

Case II.1: $< x, y >^+ = < x, y >^- = \emptyset$: *E is left unchanged, it contains no valid paths from* x *to* y

Case II.2: $< x, y >^+ \neq \emptyset$, $< x, y >^- = \emptyset$: *We add to E all paths* $\tau \circ u \to y$ *with* $< \tau, u \to y > \in < x, y >^+$ *as valid*

Case II.3: $< x, y >^+ = \emptyset$, $< x, y >^- \neq \emptyset$: *We add to E all paths* $\tau \circ u \not\to y$ *with* $< \tau, u \not\to y > \in < x, y >^-$ *as valid*

Case II.4: $< x, y >^+ \neq \emptyset$, $< x, y >^- \neq \emptyset$: *The construction of E will split into two branches:*

E' contains E and all paths $\tau \circ u \to y$ *with* $< \tau, u \to y > \in < x, y >^+$ *as valid*

E" contains E and all paths $\tau \circ u \not\to y$ *with* $< \tau, u \not\to y > \in < x, y >^-$ *as valid.*

(Thus, E is replaced in $\mathcal{E}$ by the new E' and E".)

3) Validity of Results:

Finally, we define $E \models xy$ iff there is a positive pp. $\sigma : x \cdots \rightarrow y$ s.th. $E \models \sigma$, likewise $E \models x\overline{y}$, iff there is a negative pp. $\sigma : x \cdots \not\rightarrow y$ s.th. $E \models \sigma$. □

The reader will note that this is again an upward chaining formalism, as the initial segment of a valid path has to be valid.

Definition 1.5 *Intersection of Extensions - Indirectly Skeptical Approach:*

We define

(a) $\bigcap Ext(\Gamma) \models \sigma$ iff σ is in all extensions of Γ, and

(b) $\bigcap Ext(\Gamma) \models xy$ iff in all extensions E of Γ there is some $\sigma : x \cdots \rightarrow y$; alternatively,

(c) $\bigcap Ext(\Gamma) \models xy$ iff there is some $\sigma : x \cdots \rightarrow y$ in all extensions E of Γ (and $\bigcap Ext(\Gamma) \models x\overline{y}$ analogously). (A discussion of these possibilities is found in [MS91].) □

Our Theorem will show that the directly sceptical approach and the definition of $\bigcap Ext(\Gamma) \models xy$ by (b) are fundamentally different.

The reader might check that the directly sceptical approach results in a valid path $a \rightarrow b \rightarrow d \rightarrow f$ in diagram Ex.4 below (as there is no path from a to e), but that there is an extension with $a \rightarrow b \rightarrow e \not\rightarrow f$, and no path or result linking a and f common to all extensions.

2 DEFINITIONS, STATEMENT AND PROOF OF THEOREM

Definition 2.1 *Fix some net Γ.*

For a definition of generalized and potential paths, the reader is referred to the last section of the introduction, Definition 1.1.

a) Let x, y be nodes in Γ*, then [x,y] will be the set of all generalized paths from x to y and their subpaths.*
b) Apart from the distinction on-path/off-path preclusion (not made above), there is another fundamental one to be found in the literature. We will call the alternatives split-validitiy and total-validity preclusion. Split-validity preclusion can e.g. be found in [GV89], total-validity preclusion e.g. in [HTT87]. Consider the diagram Ex.6. For split-validity preclusion, the pair $a \cdots \to c \cdots \to b$, $c \not\to d$ *is a preclusion of* $a \cdots \to b \to d$ *iff both* $a \cdots \to c$ *and* $c \cdots \to b$ *are valid paths, whereas total validity preclusion requires validity of the concatenated path* $a \cdots \to c \cdots \to b$*. The difference is discussed in detail in [Sch92A].* $\Box$

Definition 2.2 *An evaluation* $\mathcal{E}$ *assigns for each finite acyclic net* Γ*, every potential path* σ *in* Γ*, and every pair of nodes* $< x, y >$ *in* Γ *a truth value* $[\sigma]_{\mathcal{E},\Gamma}$, $[x,y]_{\mathcal{E},\Gamma}$, $[x,\overline{y}]_{\mathcal{E},\Gamma}$ *like valid, not valid, maybe valid etc. The latter two stand for the degree to which it is believed that all x are y (not y) etc.*
More precisely, if $T(\mathcal{E})$ *is the set of truth-values* $\mathcal{E}$ *can assign, then* $\mathcal{E}$ *consists of three functions, one assigning for each finite acyclic net* Γ*, and every potential path* σ *in* Γ *the value* $[\sigma]_{\mathcal{E},\Gamma} \in T(\mathcal{E})$*, the second and third assigning for each finite acyclic net* Γ*, and every pair of nodes* $< x, y >$ *in* Γ *the values* $[x,y]_{\mathcal{E},\Gamma}$, $[x,\overline{y}]_{\mathcal{E},\Gamma} \in T(\mathcal{E})$.
We shall leave the number of truth values deliberately free, but finite. The indices $\mathcal{E}$, Γ *will be omitted, being always obvious from the context. For completeness, let* $V(\mathcal{E}) \subseteq T(\mathcal{E})$ *be the set of truth values considered finally "valid", and* $\rho(\mathcal{E}) := card(T(\mathcal{E}))$ *be the "resolution" of* $\mathcal{E}$*.* $\Box$

Definition 2.3 *Let a net* Γ *be given. Let m be a natural number (we follow the usual convention of identifying natural numbers with the set of their predecessors). Let a,* $x_i : i \in m$*, b, b' be nodes in* Γ *such that*

1) There are positive potential paths σ_i from a to x_i (see Ex.7) unique in the following sense:
2) All other generalized paths τ from a to x_i (or some z on σ_i) end by simple negative preclusion before reaching σ_i, and below x_i (as illustrated in Ex.8). Thus, all generalized paths τ other than σ_i from a to x_i are definitely "dead", and not to be considered any more. In particular, they are not meant to carry any structural information.
3) Let $I \subseteq m$, $n, n' \in m - I$, $n \neq n'$, $J := I \cup \{n\}$, $J' := I \cup \{n'\}$. Let the only links going into b (b') or leaving b (b') be the positive links $\{x_i \to b: i \in J\}$ ($\{x_i \to b': i \in J'\}$) (Ex.9).
If, given these assumptions 1)-3), $[a,b] \neq [a,b']$ always implies $[a,x_n] \neq [a,x_{n'}]$, we say that the evaluation satisfies the $Principle of Disjunction$. □

In other words, in this situation, the truth value e.g. $[a,b]$ is determined by the truth values $[a,x_i]$, $i \in J$, given the additional assumption that the σ_i are the only possible candidates for paths from a to x_i. In such a situation, it seems very reasonable to demand that the truth value $[a,b]$ depends only on the truth values $[a,x_i]$ (or $[\sigma_i]$), since all other paths are dead for good. This assumption, or another like, seems to be essential for our theorem.

Theorem 2.1 *Assume that the definition of validity in extensions respects a very primitive from of preclusion, i.e. at least on-path preclusion, split-validity or total-validity (see Ex.10 (Basic Cell), for the exact picture). Then: Any evaluation $\mathcal{E}$ such that $\rho(\mathcal{E})$ is finite and $\mathcal{E}$ satisfies the Principle of Disjunction cannot capture the intersection of extensions in the sense of Definition 1.5, (b).*

Proof of the Theorem: Let $t := \rho(\mathcal{E})$. Assume that $\mathcal{E}$ gives the same results as the intersection of extensions, i.e. $[x,y]$ ($[x,\overline{y}]$) $\in V(\mathcal{E})$ iff xy ($x\overline{y}$) is true in all extensions. Using the Principle of Disjunction, we shall work for a contradiction and construct a diagram Γ necessitating more than t truth values. Our aim is to encode every extension into a path. The problem is to switch the paths simultaneously on or off (depending on the path) by the choice of extension and yet to keep the paths separate as demanded by the Principle of Disjunction. We proceed as follows: First, we describe a Basic Cell of the construction, which has a Nixon Diamond as core, and several "wings". Every path we will construct moves through one wing, and is switched on and off by the Diamond. The construction via wings serves to keep the paths isolated from each other. Every such cell corresponds to a propositional variable, and n cells will give us 2^n different paths (propositional assignments), kept apart as the Principle of Disjunction demands in its premise. Choosing n s.th. $2^n > t$ will show the result.

Basic Cell (Ex.10) The Basic Cell B_i has a core c_i, d_i, e_i, f_i, and each wing $W_{i,j}$ has three "pipes", the inner pipe with (upper) index 0, the positive pipe with index p, and the negative pipe with index n. The inner pipe serves to isolate paths (wings) from each other, the positive pipe transports positive information, and the negative one negative information. We consider the situations where we put a positive path to $p^p_{i,j}$ or $p^n_{i,j}$, and examine how far it can be extended and what the resulting polarity will be. The results are marked in the diagram, where e.g. n+:- at node y is to read: if we have a positive path going into $p^n_{i,j}$, and the Diamond is switched +, then we have a negative (-) path up to y. "/" indicates that no path is possible. In particular, all generalized paths that move into an other wing j' have to go via f_i, and end up by negative preclusion in the triangle shown in Ex.11, so they are killed in the inner pipe. The purpose of this construction is to satisfy the prerequisites of the Principle of Disjunction.

Combining Basic Cells (Ex.12) The path $\sigma_0 : a \cdots \rightarrow x_0$ will be valid only in the extension where A_1 and A_2 are chosen positive, and $\sigma_1 : a \cdots \rightarrow x_1$ is valid iff A_1 is positive, and A_2 negative, thus σ_0 encodes the extension $A_1 \wedge A_2$, σ_1 the extension $A_1 \wedge \neg A_2$.

Final Construction of Γ Let $2^n > t$. Put n isomor phic Basic Cells $B_i : i < n$ with 2^{n+1} wings $W_{i,j} : j < 2^{n+1}$ each "one above the other". (It will become clear in a moment why we take 2^{n+1} wings instead of 2^n.) For each of the 2^n different extensions E we choose 2 wings in every cell to encode

E (i.e. twice in the same way) by joining the positive or negative pipes as shown above. Next, we join all these paths in a bottom element a, and put a positive link ending in some x on each path leaving the topmost cell. We thus construct 2^{n+1} different paths $\sigma^0_j : a \cdots \to x^0_j$, $\sigma^1_j : a \cdots \to x^1_j$, $j < 2^n$, where σ^0_j and σ^1_j are alike.

Formally, we first put the basic cells $B_i : i < n$ into Γ. For each $\phi_j : n \to \{0 = true, 1 = false\}$ $(j < 2^n)$ we construct two positive potential paths σ^0_j, σ^1_j. Let $\alpha := 2j$ and work with the wings $W_{i,\alpha}$, $W_{i,\alpha+1}$ in each cell B_i, making "vertical" connections.

For i=0,
if $\phi_j(0) = 0$, add $a \to p^p_{0,\alpha}$ and $a \to p^p_{0,\alpha+1}$,
if $\phi_j(0) = 1$, add $a \to p^n_{0,\alpha}$ and $a \to p^n_{0,\alpha+1}$
for i=n-1,
if $\phi_j(n-1) = 0$, add $r^p_{n-1,\alpha} \to x^0_j$ and $r^p_{n-1,\alpha+1} \to x^1_j$,
if $\phi_j(n-1) = 1$, add $r^n_{n-1,\alpha} \to x^0_j$ and $r^n_{n-1,\alpha+1} \to x^1_j$
for $i < n-1$,
if $\phi_j(i) = 0$ and $\phi_j(i+1) = 0$ add $r^p_{i,\alpha} \to p^p_{i+1,\alpha}$ and $r^p_{i,\alpha+1} \to p^p_{i+1,\alpha+1}$,
if $\phi_j(i) = 0$ and $\phi_j(i+1) = 1$ add $r^p_{i,\alpha} \to p^n_{i+1,\alpha}$ and $r^p_{i,\alpha+1} \to p^n_{i+1,\alpha+1}$,
if $\phi_j(i) = 1$ and $\phi_j(i+1) = 0$ add $r^n_{i,\alpha} \to p^p_{i+1,\alpha}$ and $r^n_{i,\alpha+1} \to p^p_{i+1,\alpha+1}$,
if $\phi_j(i) = 1$ and $\phi_j(i+1) = 1$ add $r^n_{i,\alpha} \to p^n_{i+1,\alpha}$ and $r^n_{i,\alpha+1} \to p^n_{i+1,\alpha+1}$

From the discussion of the Basic Cells B_i, it is clear that the following holds:

1) $\sigma^0_j : a \cdots \to x^0_j$ and $\sigma^1_j : a \cdots \to x^1_j$ are valid (in the usual sense) positive paths exactly in the extension which corresponds to ϕ_j, i.e. makes a positive choice in the core of B_i iff $\phi_j(i) = 0$.

2) all generalized paths from a to x^k_j other than σ^k_j end in simple negative preclusions.

We now show that all $[a, x^1_j]$, $[a, x^1_{j'}]$ for $j \neq j' < 2^n$ are different, and we are finished by $2^n > t = the$ number of truth values of $\mathcal{E}$.

For this purpose, we add for each pair $< j, j' >$, $j \neq j'$, $j, j' < 2^n$, new nodes $b_{j,j'}$, $b'_{,j'}$ and positive links $x^0_k \to b_{j,j'}$ and $x^0_k \to b'_{,j'}$ for all $k \neq j$, $k < 2^n$ and the links $x^1_j \to b_{j,j'}$, $x^1_{j'} \to b'_{,j'}$, see Ex.13. Thus, in all extensions, and for all j, j', there will be a path from a to $b_{j,j'}$, because the "missing path" σ^0_j is replaced by its twin σ^1_j, but not from a to $b'_{,j'}$, as $\sigma^1_{j'}$ and $\sigma^0_{j'}$ are true in the same extension, and the extension corresponding to σ^0_j has no valid path from a to $b'_{,j'}$. As $\mathcal{E}$ is supposed to capture the intersection of extensions, $[a, b_{j,j'}] \neq [a, b'_{,j'}]$. By 1) and 2) above, our construction satisfies the prerequisites of the Principle of Disjunction, so the pair $b_{j,j'}$, $b'_{,j'}$ shows that $[a, x^1_j] \neq [a, x^1_{j'}]$. Repeating the argument, we see that all $[a, x^1_j]$, $j < 2^n$ are different, and we are done by $2^n > t$. □

It is natural to ask whether the finiteness assumption was necessary. Let me give a short argument to show informally that it is. If we admit unboundedly many truth values, we can encode for each potential path how it behaves with respect to each decision (which corresponds essentially to a Nixon Diamond). Thus, roughly speaking, a path will be in every extension iff every outcome of every decision is in its set of truth values. And there will be a valid path between two nodes, iff for every outcome of every decision there is a path between the two, incorporating this outcome as a truth value. Thus, looking at the (arbitrarily large) set of its truth values, we can determine if a path is valid in every extension, and, looking at the sets of truth values of each potential path linking two nodes x, y, we can determine if there is a valid one in each extension.

3 NOTE

A preliminary version of this paper was presented at a German Workshop on Nonmonotonic Reasoning held at Gesellschaft für Mathematik und Datenverarbeitung, Bonn, West Germany, in December 1989. The proceedings have appeared as a technical report, see [Sch89B].

References

[Goo55] N.Goodman, "Fact, Fiction, and Forecast", Harvard Univ. Press, 1955

[GV89] H.Geffner, T.Verma, "Inheritance = Chaining+Defeat", UCLA Technical Report CSD-890039, R-129-L, June 1989, also in: Proceed. Fourth International Symposium on Methodologies for Intelligent Systems, North Holland, 1989, p.411-418

[HTT87] J.F.Horty, R.H.Thomason, D.S.Touretzky, "A Sceptical Theory of Inheritance in Nonmonotonic Semantic Networks", Dept. Comp. Sci., Carnegie Mellon Univ., CMU-CS-87-175, October 1987 also in: Proceedings AAAI-87 (1987), p.358-363, or: AIJ 42 (1990), p.311-348

[KK89] T.Krishnaprasad, M.Kifer, "An Evidence-based Framework for a Theory of Inheritance", Proceed. IJCAI 89, p.1093-1098

[KKW89a] T.Krishnaprasad, M.Kifer, D.Warren, "On the Circumscriptive Semantics of Inheritance Networks", Proceed. Fourth International Symposium on Methodologies for Intelligent Systems, North Holland, 1989, p.448-456

[KKW89b] T.Krishnaprasad, M.Kifer, D.Warren, "On the Declarative Semantics of Inheritance Networks", Proceed. IJCAI 89, p.1098-1103

[MS91] D.Makinson, K.Schlechta: "Floating Conclusions and Zombie Paths", AI Journal (Research Note) 48 (1991) p. 199-209

[Sch89A] K.Schlechta: "Defeasible Inheritance: Coherence Properties and Semantics", Proceedings of Tübingen Workshop on Semantic Networks and Non-Monotonic Reasoning, M.Morreau ed., SNS-Report 89-48, Seminar für natürlichsprachliche Systeme, Universität Tübingen, (1989)

[Sch89B] K.Schlechta: "Directly Sceptical Inheritance cannot Capture the Intersection of Extensions", Proceedings Workshop Non-Monotonic Reasoning 1989, G.Brewka, H.Freitag eds., GMD-Report, Arbeitspapiere der Gesellschaft für Mathematik und Datenverarbeitung, 443, (1989)

[Sch92A] K.Schlechta: "Results on Non-Monotonic Logics" (Habilitationsschrift), IBM Germany, IWBS Report 204, IWBS, POB 80 08 80, D-7000 Stuttgart 80, Germany

[Ste89] L.Stein, "Skeptical Inheritance: Computing the Intersection of Credulous Extensions", Proceedings IJCAI 1989, p.1153-1158

[Ste91a] L.Stein, "Resolving Ambiguity in Nonmonotonic Reasoning", PhD Thesis, Tech.Rept. CS-90-18, Brown University, Dept. Comp. Sci., Providence, Rhode Island 02912, USA, May 1991

[Ste91b] L.Stein, "Resolving Ambiguity in Nonmonotonic Inheritance Hierarchies", AI Memo No. 1316, MIT, Artificial Intelligence Laboratory, August 1991

[THT87] D.S.Touretzky, J.F.Horty, R.H.Thomason, "A Clash of Intuitions: The Current State of Nonmonotonic Multiple Inheritance Systems", Proceed. IJCAI 1987, p.476-482