A Preference-Based Approach to Inheritance

Lynn Andrea Stein

Brown University
Dept. of Computer Science
Providence, Rhode Island 02912

Technical Report No. CS-90-08
April 1990

A Preference-Based Approach to Inheritance

Lynn Andrea Stein*
Department of Computer Science
Brown University

Abstract

Existing theories of nonmonotonic inheritance are generally either computationally feasible or semantically well-founded, but not both. We present four approaches to inheritance reasoning:

- a Touretzky-style path-based theory
- a model-theoretic semantics
- a polynomial-time inference algorithm
- a reason-maintenance labeling scheme

and demonstrate that these four approaches yield equivalent results. The underlying framework further provides a unified theory of inheritance suitable for principled analysis of existing inheritance systems (rather than *ad hoc* comparison of specific examples).

1 Introduction

Inheritance reasoning is ubiquitous. Hierarchies provide a concise encoding of much of our commonsense knowledge. They appear in various guises in the literature of artificial intelligence, but also throughout the various arts and sciences. Open nearly any textbook, and some hierarchy is likely to appear.

Reasoning with inheritance hierarchies is of particular interest to the artificial intelligence community because hierarchies are not only useful and native to human reasoning, but also simple and often tractable. The simplicity of the "inheritance problem" makes it attractive as topic that can be considered in its entirety; its

*The author is supported by an IBM Graduate Fellowship.

tractability makes it practical for the implementation of knowledge representation systems.

In this paper, we present several equivalent approaches to inheritance reasoning. Each approach may be useful in a particular context; by demonstrating their equivalence, we show that the most appropriate definition may be chosen. Further, the general frameworks we describe apply not only to the inheritance scheme presented here, but to general inheritance reasoning. In this way, the current work provides not only new directions for inheritance research, but also commentary on and extensions to the existing inheritance literature.

We begin by providing an intuitive description of what we intend by inheritance reasoning. We view inheritance hierarchies as representing primitive assertions in the knowledge base of some reasoning agent, and the inheritance problem as that of determining the agent's derived beliefs. This provides an appealing base for a well-founded inheritance theory.

In section 3, we give a formal, path-based definition of inheritance. An inheritance theory combines a theory of the transitivity of primitive assertions with an ambiguity-resolution criterion to be invoked when two paths conflict. We describe the transitivity component by the notion of reachability; the resolution criterion is captured by specificity. This framework provides a rich description language for inheritance, making rigorous such central concepts of inheritance theory as admissibility and credulous extension.

In section 4, we describe a model-theoretic semantics for inheritance hierarchies. A hierarchy defines a space of possible world-states, or unambiguous interpretations. Specificity induces a preference relation over these world-states. In the resulting preferential semantics, each world-state has a classical model-theoretic interpretation, and the interpretation of an inheritance hierarchy is the set of models of preferred world-states. Previous approaches to semantics for inheritance hierarchies have been translational: a hierarchy is represented as axioms of some target logic. Transitivity and ambiguity resolution are generally fixed as a property of either the translation process or the target logic itself. In contrast, the approach that we present here uses the topology of the hierarchy, together with some specificity criterion, to resolve ambiguities and select the preferred extensions. By varying the definition of specificity, we select different sets of preferred extensions, providing model-theoretic semantics for diverse inheritance theories. We present an alternate specificity criterion in section 8.

Section 5 describes a tractable algorithm for computing the inferences supported by an inheritance hierarchy. This is the first tractable theory for credulous inheritance—reasoning about what *might* be. Previous tractable inheritance theories have been limited to skeptical reasoning—computing what *must* be.

In section 6, we present a reason maintenance update scheme for inheritance

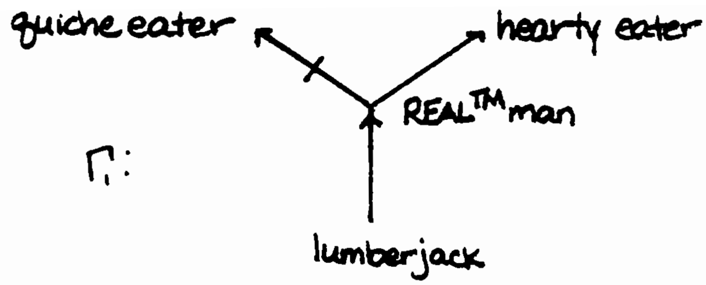

Figure 1: A simple inheritance hierarchy.

reasoners. This allows us to keep track of all of the possible (or preferred) interpretations of an inheritance hierarchy simultaneously. In section 7, we use this construct to demonstrate that several previous theories of "skeptical" inheritance do not compute precisely the intersection of interpretations—the certain conclusions of an inheritance hierarchy—and to prove that no purely path-based inheritance definition *can* compute this intersection.

2 Hierarchies as Belief Spaces

In this section, we present an intuitive interpretation of inheritance hierarchies. In later sections, we give formal definitions of these ideas; here, we hope to motivate that more formal work by answering the question of "what we mean" when we draw an inheritance hierarchy. Regrettably, most previous theories of inheritance omit such a statement of intents, and the lack of such an intuitive semantics has been one of the criticisms levelled against the entire inheritance endeavor.

Our interpretation of inheritance hierarchies is relatively simple. Each arc in an inheritance hierarchy represents an atomic assertion in the knowledge base of some reasoning agent—what this agent "believes," if you will.[1] Since this paper deals exclusively with defeasible inheritance, it is possible that an arc in the hierarchy—an atomic belief—is mistaken (e.g., *lumberjacks* might not, in fact, be *Real*$^{\mathrm{TM}}$ *men*); however, in this agent's world-model, each of these atomic assertions holds.[2]

[1]We put the word "believes" in quotation marks to emphasize that we are not proposing that edges of an inheritance hierarchy follow any realistic epistemic ontology; rather, we find the term belief, when removed from that ontologic context, to be suggestive of the kind of tentative assertion about the world we wish to describe.

[2]The exception to this is the case in which the knowledge base contains both the atomic assertion that a (defeasibly) is an x, and the atomic assertion that a (defeasibly) is not an x. Here, there

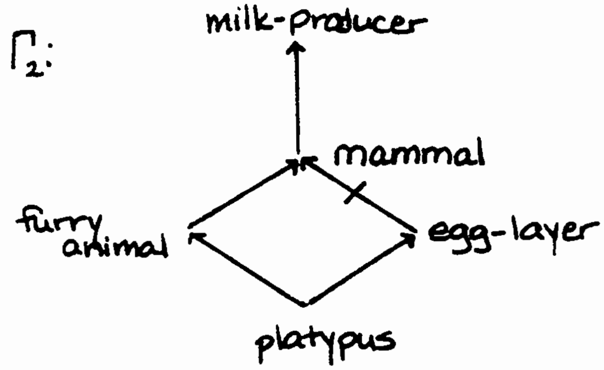

Figure 2: Is a *platypus* a *mammal*?

A second kind of defeasibility arises when we conjoin atomic beliefs. For example, the belief that *lumberjacks* are *Real*$^{\mathrm{TM}}$ *men*, and the belief that *Real*$^{\mathrm{TM}}$ *men* are *hearty eaters*, are more useful if we put them together to conclude that *lumberjacks* are *hearty eaters*. Here we have accepted both the defeasible atomic assertions, and the defeasible assumption that *lumberjacks* are *typical Real*$^{\mathrm{TM}}$ *men*, and therefore *hearty eaters*. A vital part of any inheritance theory is the notion of ***reachability***: under what circumstances are two assertions transitive?

Finally, we have the third kind of defeasibility in inheritance hierarchies: the defeasibility of ambiguous conclusions. Unlike the defeasibility of atomic assertions or that of derived (but uncontested) conclusions, ambiguity's defeasibility arises when there are two explicit and conflicting arguments. In figure 2, the derived conclusion that *platypi* are *mammals* is directly opposed by the (equally legitimate) conclusion that *platypi* are *not mammals*. In this case, even the reasoning agent is assumed to be aware of the defeasibility. Indeed, we interpret such a hierarchy as asserting that the world must be in such a state that *either platypi* are *mammals*, or they are not; but this reasoner does not know which. That is, the atomic assertions hold; and some (maximal) subset of their transitive closures hold; but there may be several such subsets, corresponding to several possible states of the world.

The situation here is not as hopeless as it may sound. First, the reasoner may prefer one of these possible world-states. For example, in figure 3, there are derived paths asserting that *blue whales* are *aquatic creatures* (by virtue of their being *whales*, which are explicitly *aquatic*) and that *blue whales* are not *aquatic* (by virtue of their being *mammals*). Because *blue whales* are *mammals only* by virtue of being *whales*, information about *whales* is more specific to *blue whales* than information about *mammals*, and the reasoner prefers to believe that the actual world-state is one in which *blue whales* are *aquatic*. We discuss the notion of specificity, and the

are (at least) two possibilities: the reasoner's beliefs may be ambiguous; or they may be (locally) inconsistent. We opt for the ambiguity interpretation (see below).

Γ_3:

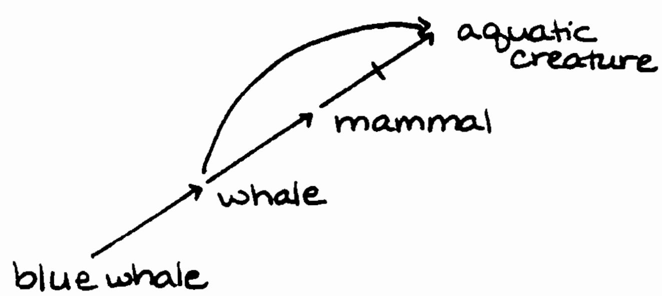

Figure 3: A *blue whale is* an *aquatic creature.*

selection of preferred world-state, in greater detail below.

Second, even in truly ambiguous cases—such as Γ_2—the inheritance hierarchy may contain contingent information. In this case, the reasoner cannot determine whether the actual world corresponds to the possible world-state in which *platypi* are *mammals*, or that in which they are not. However, the reasoner can assert that *if* the actual world-state is such that *platypi are mammals*, then it follows that *platypi* produce milk. In this fashion, the hierarchy supports conclusions relative to a particular possible world-state.

3 A Path-Based Definition

In this section, we present a formal, path-based definition of inheritance hierarchies. The approach described here is similar to that of Touretzky [29, 30], save that it is upwards reasoning. This type of inheritance reasons about the properties of some particular object, rather than about the set of objects possessing some particular property. Some ramifications of upwards *vs.* downwards inheritance are discussed by Touretzky, *et al.* [31]; the computational advantages of upwards inheritance are described by Levesque and Selman [25].

The approach that we describe in this section is also credulous: it allows a conclusion whenever that conclusion is consistent with *some* (preferred) interpretation of the hierarchy; or, whenever it is a plausible conclusion of the reasoner's explicit beliefs. Thus, in a hierarchy like Γ_2, we expect to draw both the conclusion that a *platypus* is-a *mammal*, and that a *platypus* is-not-a *mammal*. In section 6, we discuss the problem of *contingent reasoning*—determining the assumptions that underlie credulous conclusions. For example, we conclude that a *platypus* is-a *milk-producer* whenever we assume that it is-a *mammal*. In section 7, we discuss skeptical

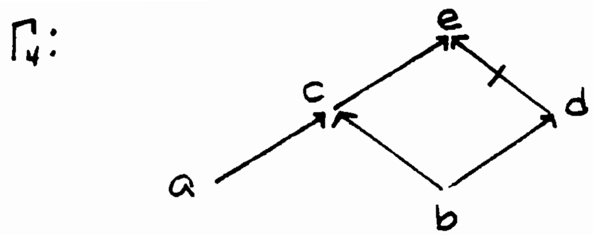

Figure 4: A hierarchy is ambiguous w.r.t. a focus node.

reasoning, in which only the certain—uncontested—conclusions of the hierarchy are considered.

3.1 The Framework

inheritance hierarchies

An *inheritance hierarchy* $\Gamma = \langle V_\Gamma, E_\Gamma \rangle$ is a directed acyclic graph with positive and negative edges, intended to denote "is-a" and "is-not-a" respectively. We write a positive edge from a to x as $a \cdot x$, and a negative edge $a \cdot \neg x$. We call a sequence of positive edges $a \cdot s_1 \cdots s_n \cdot x$ (n ≥ 0),[3] a *positive path*, and a sequence of positive edges followed by a single negative edge $a \cdot s_1 \cdots s_n \cdot \neg x$ (n ≥ 0) a *negative path*. *paths*

A path, or *argument*, $a \cdot s_1 \cdots s_n \cdot (\neg)x$ supports the *inference* "a is (not) an x." *inferences* We use the notation $a \rightarrow x$ (resp., $a \not\rightarrow x$) to stand for this inference, or conclusion, independently of the path through which it is derived. One inference—e.g., $a \rightarrow x$—may have many supporting arguments—$a \cdot s_1 \cdots s_n \cdot x$, $a \cdot \tau \cdot x$, *etc.*

Given an inheritance hierarchy $\Gamma = \langle V_\Gamma, E_\Gamma \rangle$ with nodes $a, x \in V_\Gamma$, we say that x is *reachable from* a (alternately, *a-reachable*) if there is some path $a \cdot s_1 \cdots s_n \cdot (\neg)x$ *reachability* in E_Γ. If the final edge is positive—$s_n \cdot x$—we say that x is *positively reachable from* a; similarly, $s_n \cdot \neg x$ and *negatively a-reachable*. By extension, we say that an *edge* $s \cdot (\neg)x$ is reachable from a if s is positively a-reachable, and a *path* $s_1 \cdots s_n \cdot (\neg)x$ is reachable from a if every edge on that path is a-reachable. We say that a hierarchy Γ is *a-connected* if every node in V_Γ and every edge in E_Γ is reachable from a. When *a-connected* reasoning about an inheritance hierarchy w.r.t. a particular node, we call that node the *focus node*. *focus node*

Ambiguity arises when two paths conflict. Formally, an inheritance hierarchy Γ is *ambiguous* w.r.t. a node a if there is some node $x \in V_\Gamma$ such that both $a \cdot s_1 \cdots s_n \cdot x$ *ambiguity*

[3]The notation $a \cdot s_1 \cdot s_2$ abbreviates the sequence of edges $\{a \cdot s_1, s_1 \cdot s_2\}$; $s_1 \cdots s_n$ abbreviates $\{s_1 \cdot s_2, s_2 \cdot s_3, \cdots s_{n-1} \cdot s_n\}$.

and $a \cdot t_1 \cdots t_m \cdot \neg x$ are in E_Γ. In this case, we say that the ambiguity is *at* x. Ambiguity is always relative to a focus node: for example, Γ_4 is unambiguous w.r.t. a, but ambiguous w.r.t. b (at e).

support

An inheritance hierarchy Γ *supports a path* $a \cdot s_1 \cdots s_n \cdot (\neg)x$, written $\Gamma \rhd a \cdot s_1 \cdots s_n \cdot (\neg)x$, if the corresponding sequence of edges $a \cdot s_1 \cdots s_n \cdot (\neg)x$ is in E_Γ *and it is admissible* according to specificity. Γ *supports an inference* $a \rightarrow x$ (resp, $a \not\rightarrow x$) if it supports some corresponding path. For simplicity, we also allow the degenerate path a, with the corresponding inference $a \rightarrow a$.

3.2 Specificity

specificity

A *specificity criterion*, or preemption strategy, makes choices among certain competing paths. The idea of a specificity criterion dates from Touretzky [29, 30]. Since then, many definitions of specificity have appeared in the literature, but all operate on the same underlying principle: more specific information is more directly relevant. For example, in figure 3, information about *whales* is more specific to *blue whales* than information about *mammals*, so we can infer that *blue whales are* aquatic.

The framework that we have presented so far is compatible with many of the existing definitions of specificity. In this section, we present an original definition that closely resembles Touretzky's original notion of specificity [29, 30], but which is computable in polynomial time (see section 5). In section 8, we describe an alternate definition of specificity—one resembling that of Sandewall [24] and Horty, *et al.* [14]—and show how to integrate that definition into this framework.

admissibility

An edge $v \cdot (\neg)x$ is *admissible* in Γ w.r.t. a if there is some path $a \cdot s_1 \cdots s_n \cdot v$, $(n \geq 0)$, in E_Γ, and

1. None of the edges of $a \cdot s_1 \cdots s_n \cdot v$ is redundant in Γ w.r.t. a,
2. Each of the edges of $a \cdot s_1 \cdots s_n \cdot v$ is admissible in Γ w.r.t. a, and
3. No intermediate node $a, s_1, \ldots s_n$ is a preemptor of $v \cdot (\neg)x$.

A *path is admissible* in Γ w.r.t. a if every edge in that path is admissible, and no edge except possibly the last is redundant in Γ w.r.t. a.

preemptor

A node s is a *preemptor* of $v \cdot x$ (resp., $v \cdot \neg x$) if $s \cdot \neg x$ (resp., $s \cdot x$) is admissible in Γ w.r.t. a.

The difficulties caused by redundant links were noted by Touretzky [29, 30]: Consider Γ_2 augmented by an additional edge from *blue whale* to *mammal*. This edge would be redundant: *blue whales* are typically *mammals* even without the explicit assertion. However, if the edge from *blue whale* to *mammal* were not excluded, there would be an admissible path *blue whale* $\cdot$ *mammal* $\cdot \neg$*aquatic creature*—no intermediate node is a preemptor of *mammal* $\cdot$ *aquaticcreature*. Clearly, this is not the intended meaning here (or, indeed, in any network of this form, since the "whale" node is always more specific than the "mammal" node (w.r.t. "blue whales")).

An edge $b \cdot w$ is *redundant* in Γ w.r.t. focus node a if there is some positive path $b \cdot t_1 \cdots t_m \cdot w \in E_\Gamma$, $n \gtrsim 0$, for which *redundancy*

1. Each of the edges of $b \cdot t_1 \cdots t_m$ is admissible in Γ w.r.t. a,
2. There are no c and i such that $c \cdot \neg t_i$ is admissible in Γ w.r.t. a; there is no c such that $c \cdot \neg w$ is admissible in Γ w.r.t. a.[4]

For example, the path-based definition of supports yields the following conclusions on the hierarchy Γ_2 in figure 2:

Γ_2	$\mid\!\!\triangleright$	*platypus*$\rightarrow$*platypus*	Γ_2	$\mid\!\!\triangleright$	*furry animal*$\rightarrow$*furry animal*
Γ_2	$\mid\!\!\triangleright$	*platypus*$\rightarrow$*furry animal*	Γ_2	$\mid\!\!\triangleright$	*furry animal*$\rightarrow$*mammal*
Γ_2	$\mid\!\!\triangleright$	*platypus*$\rightarrow$*egg-layer*	Γ_2	$\mid\!\!\triangleright$	*furry animal*$\rightarrow$*milk-producer*
$\star\,\Gamma_2$	$\mid\!\!\triangleright$	*platypus*$\rightarrow$*mammal*			
$\star\,\Gamma_2$	$\mid\!\!\triangleright$	*platypus*$\not\rightarrow$*mammal*	Γ_2	$\mid\!\!\triangleright$	*mammal*$\rightarrow$*mammal*
Γ_2	$\mid\!\!\triangleright$	*platypus*$\rightarrow$*milk-producer*	Γ_2	$\mid\!\!\triangleright$	*mammal*$\rightarrow$*milk-producer*
Γ_2	$\mid\!\!\triangleright$	*egg-layer*$\rightarrow$*egg-layer*	Γ_2	$\mid\!\!\triangleright$	*milk-producer*$\rightarrow$*milk-producer*
Γ_2	$\mid\!\!\triangleright$	*egg-layer*$\not\rightarrow$*mammal*			

In this case, specificity cannot resolve the ambiguity w.r.t. *platypus* at *mammal*, and Γ_2 supports both of the $\star$'d assertions: that *platypi* are *mammals*, and that they are not *mammals*. In such a case—when a hierarchy *supports* conflicting paths—we say that the hierarchy is *truly ambiguous*. *truly ambiguous*

The conclusion-set of Γ_3 is

Γ_3	$\mid\!\!\triangleright$	*blue whale*$\rightarrow$*blue whale*	Γ_3	$\mid\!\!\triangleright$	*whale*$\rightarrow$*whale*
Γ_3	$\mid\!\!\triangleright$	*blue whale*$\rightarrow$*whale*	Γ_3	$\mid\!\!\triangleright$	*whale*$\rightarrow$*mammal*
Γ_3	$\mid\!\!\triangleright$	*blue whale*$\rightarrow$*mammal*	Γ_3	$\mid\!\!\triangleright$	*whale*$\rightarrow$*aquatic creature*
Γ_3	$\mid\!\!\triangleright$	*blue whale*$\rightarrow$*aquatic creature*			
			Γ_3	$\mid\!\!\triangleright$	*mammal*$\rightarrow$*mammal*
Γ_3	$\mid\!\!\triangleright$	*aquatic creature*$\rightarrow$*aquatic creature*	Γ_3	$\mid\!\!\triangleright$	*mammal*$\not\rightarrow$*aquatic creature*

[4]Although the definitions of admissibility and redundancy are mutually dependent, they are not circular. Because the hierarchy is acyclic, it can be ordered topologically, and the definition of an admissible edge $v \cdot (\neg)x$ depends only on the redundancy of edges with endpoints strictly topologically earlier than x.

In this case, specificity resolves the ambiguities w.r.t. *blue whale* and *whale* at *aquatic creature*. As a result,

$$\Gamma_3 \quad \not\approx \quad \textit{blue whale} \not\rightarrow \textit{aquatic creature} \qquad \text{and} \qquad \Gamma_3 \quad \not\approx \quad \textit{whale} \not\rightarrow \textit{aquatic creature}$$

The same set of conclusions would result if we added a redundant edge from *blue whale* to *mammal*.

4 Model-theoretic Semantics

Previous theories of inheritance reasoning have fallen into two major approaches: *translational* and ***path-based***. Translational theories include those of McCarthy [21]; Etherington and Reiter [7, 8, 9]; Haugh [12]; Przymusinska and Gelfond [11, 23]; Bacchus [3]; Boutilier [4]; and Krishnaprasad, Kifer, and Warren [17, 18]. These approaches provide traditional model-theoretic semantics for a hierarchy by encoding the hierarchy in some existing logic. While this results in a well-behaved semantics, the translation process itself loses the path-based nature and topological relations of the original hierarchy. It is our belief that the particular topology of an inheritance hierarchy is crucial to its interpretation,[5] and translational approaches which confound explicit edges of the hierarchy with derived conclusions cannot preserve this topology.

In contrast, path-based approaches—such as those of Touretzky, Horty, and Thomason [14, 13, 28, 29, 30, 31]; Sandewall [24]; and Geffner and Verma [10]—preserve precisely the fundamental nature of topology in inheritance; unfortunately, no previous path-based approach has led to a model-theoretic characterization of the conclusions of an inheritance hierarchy. Instead, path-based approaches provide "semantics" for a hierarchy only in terms of the supported paths or inferences—a sort-of "proof-theoretic" approach to semantics.

Further, virtually every previous approach to inheritance semantics contains a fixed ambiguity-resolving (preemption) strategy.[6] Although the selection of an appropriate preemption strategy is still a subject for debate in the inheritance literature (see, e.g., Touretzky, *et al.*'s discussion of the "Clash of Intuitions" [31]), existing systems of inheritance assume some single, fixed strategy. This makes principled analyses of preemption strategies at a theoretical level impossible. As a result, most so-called "comparisons of inheritance theories" are in reality ***ad hoc*** comparisons of system performances on selected examples.

[5] For example, most definitions of preemption make the subtle distinction between redundant and preempting paths based solely on fine topological distinctions.

[6] The exception to this is Haugh [12]; see the discussion of his work in section 9, below.

This section describes an approach to inheritance which reduces hierarchies to their credulous extensions—unambiguous possible interpretations. Since extensions are also hierarchies, and since the reduction to a set of preferred extensions is purely path-based, this approach takes seriously the primacy of topology in inheritance reasoning. However, these credulous extensions are unambiguous, and it is straightforward to provide a model-theoretic semantics for any single credulous extension. The resulting interpretation of an inheritance hierarchy is in terms of the models of its preferred credulous extensions. Unlike translational approaches, we do not rely on a logic as an intermediary; unlike past path-based approaches, we *do* obtain a model-theoretic interpretation as an end result.

An additional advantage of treating ambiguity-resolution as the separate problem of selecting preferred extensions is that the semantics for individual extensions described here can be used with a number of different selection strategies. This means we can compare the ambiguity-resolving heuristics of various inheritance theories directly and theoretically, rather than by resorting to *ad hoc* analysis of specific examples. We do this in section 8, below.

4.1 Semantics for Credulous Extensions

We focus first on the problem of providing a model-theoretic semantics for a single, unambiguous "credulous extension." In general, we believe that a translational approach to inheritance semantics is undesirable. Such an approach trades the topological information inherent in a hierarchy for the semantics of an existing logic. In particular, translational approaches provide less-than-satisfactory means of dealing with ambiguities. Often, they merely adopt the ambiguity-resolving strategy of the target logic, which may not be appropriate for inheritance hierarchies. Where they do provide additional, explicit ambiguity resolution, it is fixed as a part of the translation procedure.

The approach to semantics for credulous extensions we present here *is* translational. However, it is not subject to these criticisms of translational approaches precisely because it provides translations *only* for credulous extensions—*unambiguous* subhierarchies—and not for a hierarchy as a whole. Thus, it does not rely on the translation procedure or the underlying logic for ambiguity-resolution strategies. All ambiguity-resolution is done in *selecting* some preferred subset of the credulous extensions; once the set of preferred credulous extensions—or possible interpretations—for a hierarchy has been established, the semantics for any single extension are straightforward. In the next section, we discuss the problem of deriving the appropriate (preferred) set of credulous extensions for a hierarchy.

A credulous extension corresponds to a possible 'world-state'—one in the space of 'world-states' defined by taxonomic ambiguity. Formally, a *credulous extension* *credulous extension*

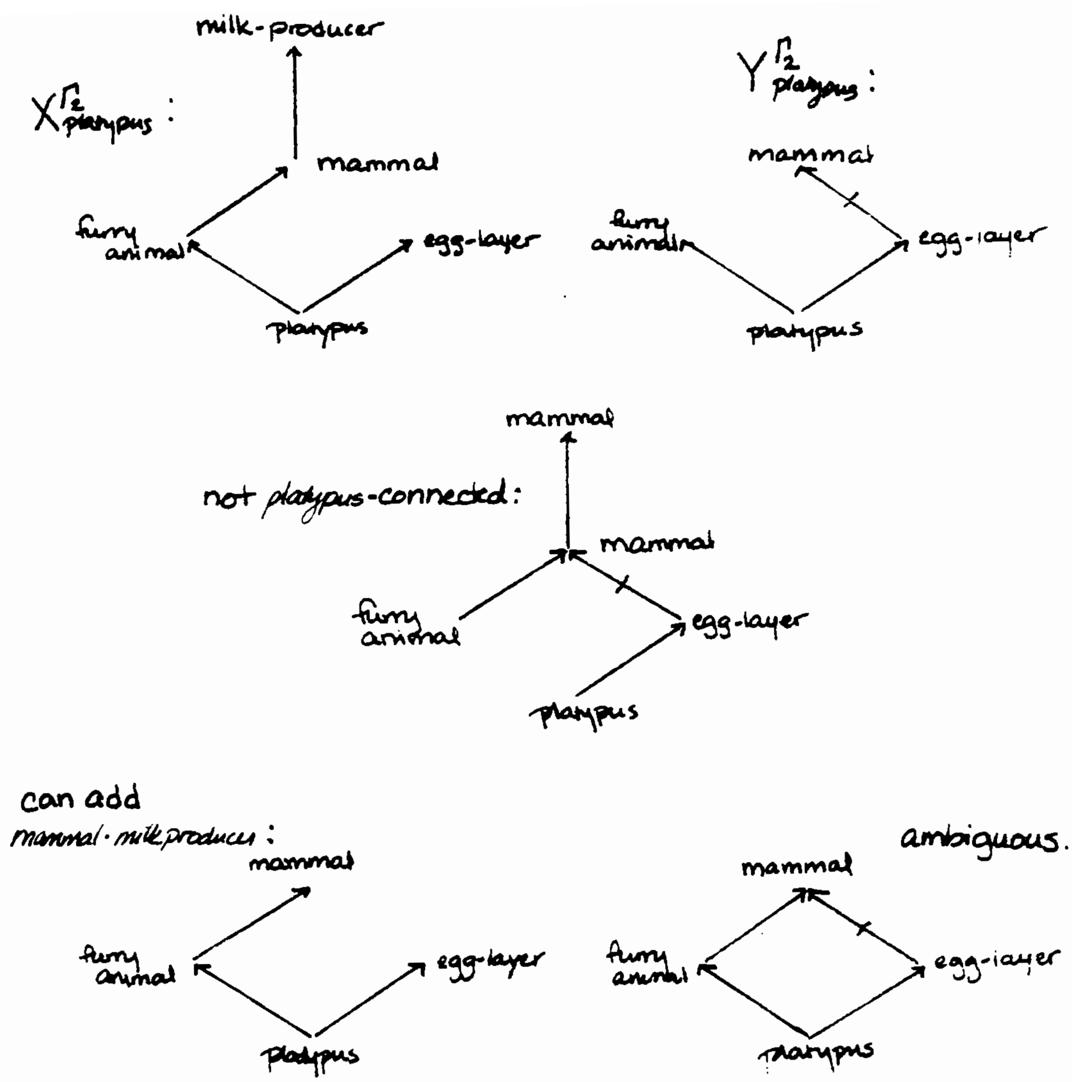

Figure 5: Some extensions and non-extensions of Γ_2.

of an inheritance hierarchy Γ with respect to a node a is a maximal unambiguous a-connected subhierarchy of Γ with respect to a: if X_a^Γ is a credulous extension of Γ w.r.t. a, then for every edge $v \cdot (\neg) x \in E_\Gamma - E_{\mathrm{X}_a^\Gamma}$, adding $v \cdot (\neg) x$ to X_a^Γ would make X_a^Γ unambiguous or not a-connected. An example of several credulous extensions—and some non-extensions—for the hierarchy of figure 2 is given in figure 5.

Because a credulous extension is unambiguous, every path is admissible. Thus, instead of $\mathrm{X}_a^\Gamma \rhd a \rightarrow x$—there is an admissible positive path from a to x—we need merely check that there is *some* positive path from a to x (resp., $a \not\rightarrow x$ and negative path). We make use of this to provide a straightforward model-theoretic semantics for a single extension:

For every vertex $x \in V_{\mathrm{X}_a^\Gamma}$, we create a unique propositional variable $\widehat{x}$. $\widehat{\mathrm{X}_a^\Gamma}$ is the theory (in the propositional calculus) given by $\widehat{\mathrm{X}_a^\Gamma}$

$$\widehat{a} \quad \wedge \bigwedge_{x \cdot y \in E_{\mathrm{X}_a^\Gamma}} (\widehat{x} \supset \widehat{y}) \quad \wedge \bigwedge_{x \cdot \neg y \in E_{\mathrm{X}_a^\Gamma}} (\widehat{x} \supset {\sim}\widehat{y})^7$$

Since X_a^Γ is acyclic and unambiguous, $\widehat{\mathrm{X}_a^\Gamma}$ is consistent and has a model.

In other words, we translate the edges of X_a^Γ into material implications, allowing inference chains exactly when there are paths in the credulous extension. It would be nice if all of inheritance semantics were this easy. However, Thomason and Horty demonstrate that the translation into propositional logic does not work even in the simpler case of strict (non-defeasible) inheritance [28]. In the general defeasible case, the propositional theory corresponding to an ambiguous inheritance hierarchy would be inconsistent. The only reason that the propositional theory works here is that it is the translation of a credulous extension: an *unambiguous* subhierarchy. Ambiguity resolution must therefore be applied in selecting the credulous extensions that are the preferred interpretations of the original hierarchy.

Once we have chosen the preferred extensions—as we discuss in the next section—this translational semantics provides the desired result. The following theorem states that support—the existence of an admissible path—is equivalent to derivability in the propositional theory. Thus the path-based definition of support, for credulous extensions, is sound and complete w.r.t. the model-theoretic semantics defined here.

Theorem 1 *Let Γ be an inheritance hierarchy, with $a, x \in V_\Gamma$. Let X_a^Γ be a credulous extension of Γ w.r.t. a, and let $\widehat{\mathrm{X}_a^\Gamma}$ be the propositional theory corresponding to X_a^Γ. The following are equivalent:*

1. $\mathrm{X}_a^\Gamma \rhd a \rightarrow x$ (resp., $a \not\rightarrow x$).

[7] $\wedge$, $\supset$, and $\sim$ should be read as propositional conjunction, material implication, and negation, respectively).

2. *x is positively a-reachable in* X_a^Γ *(resp., negatively a-reachable).*

3. $\widehat{X_a^\Gamma} \vdash \hat{x}$ *(resp.,* $\sim\hat{x}$*).*

4. $\widehat{X_a^\Gamma} \models \hat{x}$ *(resp.,* $\sim\hat{x}$*).*

The proofs of all theorems may be found in appendix A.

4.2 Selecting Preferred Extensions

Given a single credulous extension, we have seen how to derive a model-theoretic interpretation of that extension. In fact, the semantics of the previous section would work any time we began with an unambiguous inheritance hierarchy (which would have only a single credulous extension w.r.t. any given focus node). But inheritance hierarchies are generally ambiguous. As a result, Γ may have several extensions w.r.t. a given focus node, a. Some of these extensions may be more intuitive than others. In this section, we describe a means of selecting the *preferred* (more intuitive) extensions of Γ w.r.t. a. The semantics of Γ are then simply the sets of interpretations for the preferred extensions of Γ w.r.t. the nodes of V_Γ.

In inheritance hierarchies, *specificity* gives us a means of ruling out unintuitive interpretations. Thus, we use the definition of *admissibility* according to specificity, from section 3.2, above, to define the preference relation over credulous extensions. We say that one extension is *preferred* to another if it is more consistent with the constraints of specificity:

Let X_a^Γ and Y_a^Γ be two credulous extensions of an inheritance hierarchy Γ w.r.t. focus node a. Then *specificity prefers* X_a^Γ *to* Y_a^Γ ($X_a^\Gamma \preceq Y_a^\Gamma$) if there are some nodes v and x such that *preferred extension*

1. X_a^Γ and Y_a^Γ agree on all edges whose endpoints topologically precede x in any topological sort,

2. The edge $v \cdot (\neg)x$ is is inadmissible in Γ w.r.t. a.

3. $\exists s_1, \ldots s_n$, $Y_a^\Gamma \rhd a \cdot s_1 \cdots s_n \cdot v \cdot (\neg)x$, and

4. $X_a^\Gamma \not\rhd a \cdot s_1 \cdots s_n \cdot v \cdot (\neg)x$

If a credulous extension is minimal under this preorder—i.e. no other extension is preferred to it—we call it a *preferred extension* of the hierarchy:

$$\mathcal{P}ref(\Gamma, a) = \{X_a^\Gamma \mid \forall Y_a^\Gamma, Y_a^\Gamma \not\preceq X_a^\Gamma\}$$

Figure 6 shows the two credulous extensions of the hierarchy in figure 3. Specificity prefers $X_{\textit{blue whale}}^{\Gamma_3}$ to $Y_{\textit{blue whale}}^{\Gamma_3}$:

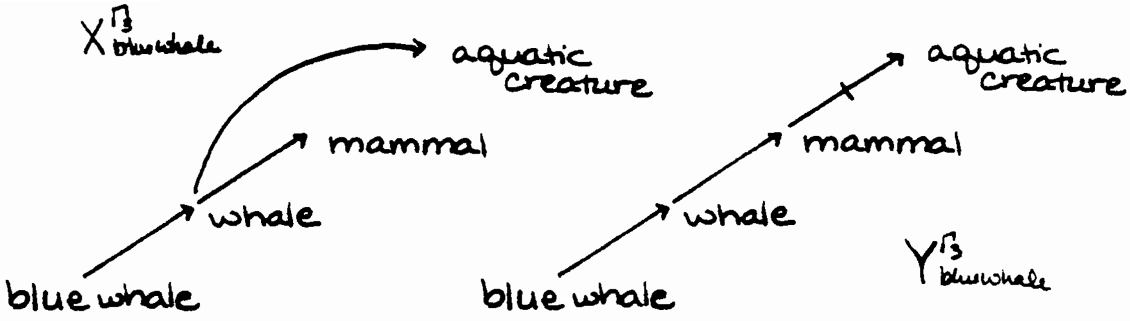

Figure 6: Specificity prefers $X^{\Gamma_3}_{\textit{blue whale}}$ to $Y^{\Gamma_3}_{\textit{blue whale}}$.

1. $X^{\Gamma_3}_{\textit{blue whale}}$ and $Y^{\Gamma_3}_{\textit{blue whale}}$ agree up to *aquatic creature*
2. *mammal* $\cdot$ *aquatic creature* is inadmissible in Γ_3 w.r.t. *blue whale*
3. $Y^{\Gamma_3}_{\textit{blue whale}} \rhd$ *blue whale* $\cdot$ *whale* $\cdot$ *mammal* $\cdot \neg$*aquatic creature*, and
4. $X^{\Gamma_3}_{\textit{blue whale}} \not\rhd$ *blue whale* $\cdot$ *whale* $\cdot$ *mammal* $\cdot \neg$*aquatic creature*.

Theorem 2 is a soundness and completeness theorem for the path-based definition of section 3 w.r.t. the complete model-theoretic semantics defined here: selecting preferred extensions and providing translational semantics for those preferred extensions. In section 8, we demonstrate that other inheritance theories, with other definitions of specificity, can be described in these terms. This enables us to use the definitions of this section and section 3 to give similar soundness and completeness results for several existing inheritance theories.

Theorem 2 *Let Γ be an inheritance hierarchy, with $a, x \in V_\Gamma$, and let $\mathcal{P}ref(\Gamma, a)$be the set of preferred extensions of Γ w.r.t. a. Then $\Gamma \rhd a \rightarrow x$ iff there is some preferred extension $X^\Gamma_a \in \mathcal{P}ref(\Gamma, a)$ such that $\widehat{X^\Gamma_a} \models \widehat{x}$ (resp., $a \not\rightarrow x$ and $\sim\widehat{x}$).*

The complete model-theoretic interpretation of Γ_2 in figure 2, is therefore the set of models of preferred extensions. In this case, specificity can't disambiguate the hierarchy, and both extensions are preferred. The propositional theories w.r.t. *platypus* are

$$\begin{array}{lclcl} \widehat{\textit{platypus}} & \wedge & (\widehat{\textit{platypus}} \supset \widehat{\textit{furry animal}}) & \wedge & (\widehat{\textit{furry animal}} \supset \widehat{\textit{mammal}}) \\ & \wedge & (\widehat{\textit{mammal}} \supset \widehat{\textit{milk-producer}}) & \wedge & (\widehat{\textit{platypus}} \supset \widehat{\textit{egg-layer}}) \end{array}$$

entailing the conclusions

$$\widehat{\textit{platypus}} \quad \widehat{\textit{furry animal}} \quad \widehat{\textit{mammal}} \quad \widehat{\textit{milk-producer}} \quad \widehat{\textit{egg-layer}}$$

and

$$\widehat{\mathit{platypus}} \;\wedge\; (\widehat{\mathit{platypus}} \supset \widehat{\mathit{furry\ animal}}) \;\wedge\; (\widehat{\mathit{platypus}} \supset \widehat{\mathit{egg\text{-}layer}})$$
$$\wedge\; (\widehat{\mathit{egg\text{-}layer}} \supset {\sim}\widehat{\mathit{mammal}})$$

which entails

$$\widehat{\mathit{platypus}} \quad \widehat{\mathit{furry\ animal}} \quad {\sim}\widehat{\mathit{mammal}} \quad \widehat{\mathit{egg\text{-}layer}}$$

The preferential semantics allows the conclusion that *platypus* is-a x if x is entailed by the models of either of these theories, so *playtpus* is-a

platypus   *furry animal*   *mammal*   *milk-producer*   *egg-layer*

and *platypus* is-not-a *mammal*. The hierarchy is unambiguous w.r.t. its other nodes. The conclusions w.r.t. *furry animal* are $\widehat{\mathit{furry\ animal}}$, $\widehat{\mathit{mammal}}$, and $\widehat{\mathit{milk\text{-}producer}}$, since $X^{\Gamma_2}_{\widehat{\mathit{furry\ animal}}}$ is

$$\widehat{\mathit{furry\ animal}} \;\wedge\; (\widehat{\mathit{furry\ animal}} \supset \widehat{\mathit{mammal}}) \;\wedge\; (\widehat{\mathit{mammal}} \supset \widehat{\mathit{milk\text{-}producer}})$$

$X^{\Gamma_2}_{\widehat{\mathit{mammal}}}$ is

$$\widehat{\mathit{mammal}} \;\wedge\; (\widehat{\mathit{mammal}} \supset \widehat{\mathit{milk\text{-}producer}})$$

which entails $\widehat{\mathit{mammal}}$ and $\widehat{\mathit{milk\text{-}producer}}$, and $X^{\Gamma_2}_{\widehat{\mathit{milk\text{-}producer}}}$ is $\widehat{\mathit{milk\text{-}producer}}$. $X^{\Gamma_2}_{\widehat{\mathit{egg\text{-}layer}}}$ is

$$\widehat{\mathit{egg\text{-}layer}} \;\wedge\; (\widehat{\mathit{egg\text{-}layer}} \supset {\sim}\widehat{\mathit{mammal}})$$

entailing $\widehat{\mathit{furry\ animal}}$ and ${\sim}\widehat{\mathit{mammal}}$.

In contrast, specificity does provide a preference over the credulous extensions of the hierarchy in figure 3 w.r.t. *blue whale* and *whale*. The interpretation of Γ_3 w.r.t. *blue whale* is therefore the interpretation of the single preferred extension, $X^{\Gamma_3}_{\widehat{\mathit{blue\ whale}}}$, from figure 6:

$$\widehat{\mathit{blue\ whale}} \;\wedge\; (\widehat{\mathit{blue\ whale}} \supset \widehat{\mathit{whale}}) \;\wedge\; (\widehat{\mathit{whale}} \supset \widehat{\mathit{mammal}})$$
$$\wedge\; (\widehat{\mathit{whale}} \supset \widehat{\mathit{aquatic\ creature}})$$

This entails the conclusions (w.r.t. *blue whale*)

$$\widehat{\mathit{blue\ whale}} \quad \widehat{\mathit{whale}} \quad \widehat{\mathit{mammal}} \quad \widehat{\mathit{aquatic\ creature}}$$

Similarly, specificity prefers the extension w.r.t. *whale* in which *whales* are *aquatic*, with the corresponding propositional theory

$$\widehat{\mathit{whale}} \;\wedge\; (\widehat{\mathit{whale}} \supset \widehat{\mathit{mammal}}) \;\wedge\; (\widehat{\mathit{whale}} \supset \widehat{\mathit{aquatic\ creature}})$$

and the hierarchy entails the inferences that a *whale* is-a *whale*, *mammal*, and *aquatic creature*. The hierarchy is unambiguous w.r.t. *mammals* and *aquatic creatures*, so there is only one credulous extension w.r.t. each of these:

$$\widehat{\mathit{mammal}} \;\wedge\; (\widehat{\mathit{mammal}} \supset {\sim}\widehat{\mathit{aquatic\ creature}})$$

and a *mammal* is-a *mammal*, but is-not-a *aquatic creature*; and

$$\widehat{\mathit{aquatic\ creature}}$$

so an *aquatic creature* is-a *aquatic creature*.

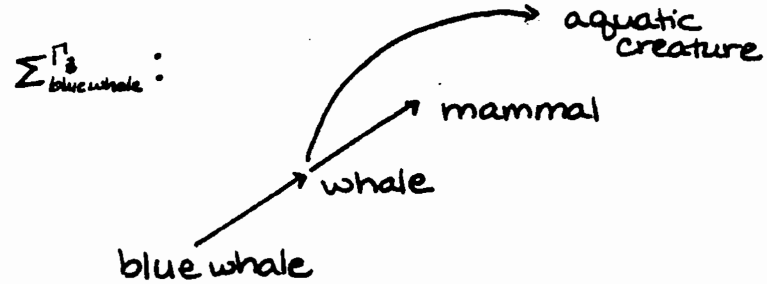

Figure 7: $\Sigma^{\Gamma_3}_{blue\ whale}$.

5 Computing Specificity

The preference criterion approach to ambiguity resolution is satisfying from a semantic perspective, but it does not tell us much about how to *derive* the preferred extensions of a hierarchy. In this section, we give a polynomial-time algorithm for computing Σ^{Γ}_{a}, the *specificity extension* of Γ w.r.t. a. Σ^{Γ}_{a} is a subhierarchy of Γ with the inadmissible edges (w.r.t. a) removed. For example, the specificity extension of the hierarchy in figure 3 w.r.t. *blue whale* is shown in figure 7. The following algorithm always yields a unique specificity extension for a hierarchy w.r.t. a focus node.

specificity extension

Compute–Specificity–Extension

Let Γ be an inheritance hierarchy.
Sort the nodes of Γ topologically
For each focus node a ; Σ^{Γ}_{a} will be the *specificity*
  $\Sigma^{\Gamma}_{a} := \Gamma$; *extension* of Γ w.r.t. a.
  For each node x reachable from a in Σ^{Γ}_{a}, in topological order
    ; from a to ...
    For each edge $v \cdot (\neg)\, x \in E_{\Sigma^{\Gamma}_{a}}$, in reverse topological order
      ; check if $v \cdot (\neg)\, x$ is admissible in Γ
      Let $\Gamma^{\star} = \Sigma^{\Gamma}_{a} - \{q \mid q$ is a preemptor of $v \cdot (\neg)\, x$ in $\Sigma^{\Gamma}_{a}\}$
      If $\Gamma^{\star}$ no longer contains a positive path from a to v,
        then remove the edge $v \cdot (\neg)\, x$ from Σ^{Γ}_{a}
    For each remaining positive edge $p \cdot x$
      ; check if $p \cdot x$ is redundant in Γ
      If Σ^{Γ}_{a} contains a path $p \cdot q_1 \cdots q_n \cdot x$ ($n \geq 1$)
        such that there is no negative path through

admissible edges from a to any of q_i, x, in Σ_a^Γ
then mark the edge $p \cdot x$ as redundant

Not all of the ambiguities in an inheritance hierarchy are susceptible to specificity. For example, "diamond" ambiguities such as the "platypus diamond" of figure 2 cannot be resolved using any preemption technique. For this reason, Σ_a^Γ may still contain ambiguities, and may yield several credulous extensions w.r.t. a. However, the ambiguities that *can* be resolved by specificity have been, and the credulous extensions of Σ_a^Γ w.r.t. a are the interpretations of Γ (w.r.t. a) that are consistent with the specificity criterion. Theorem 3 says that any inference in Σ_a^Γ or its extensions is admissible. This means that computing admissible inferences in Σ_a^Γ simply requires verifying that there is *some* supporting path.

Theorem 3 *Let Γ be an inheritance hierarchy, with $a, x \in V_\Gamma$, and let Σ_a^Γ be the specificity extension of Γ w.r.t. a. Then the following are equivalent:*

1. *$\Sigma_a^\Gamma \rhd a \rightarrow x$ (resp., $a \not\rightarrow x$).*
2. *x is positively a-reachable in Σ_a^Γ (resp., negatively a-reachable).*
3. *There is an extension $X_a^{\Sigma_a^\Gamma}$ of Σ_a^Γ w.r.t. a, such that $X_a^{\Sigma_a^\Gamma} \rhd a \rightarrow x$ (resp., $a \not\rightarrow x$).*

The next theorem is a "proof of correctness" for the Compute-Specificity-Extension algorithm. It says that the inferences supported by the specificity extension (w.r.t. a) are precisely the admissible inferences in Γ w.r.t. a. Since both Compute-Specificity-Extension and path-finding are polynomial, we can use them to determine whether an inference is supported in an inheritance hierarchy in polynomial time.

Theorem 4 *Let Γ be an inheritance hierarchy, with $a, x \in V_\Gamma$, and let Σ_a^Γ be the specificity extension of Γ w.r.t. a. Then $\Gamma \rhd a \rightarrow x$ iff x is positively a-reachable in Σ_a^Γ (resp., $a \not\rightarrow x$ and negatively a-reachable).*

For example, $\Sigma_{\mathit{blue\ whale}}^{\Gamma_3}$ has paths from *blue whale* to *whale*, *mammal*, and *aquatic creature*. These are precisely the desired inferences w.r.t. *blue whale*.

6 Reason Maintenance and Inheritance

In this section, we define a reason maintenance system for inheritance hierarchies. A reason maintenance system is a construct that keeps track of the assumptions or justifications supporting a particular conclusion; several systems of this type appear in the literature [5, 6, 19, 20, 22]. In the case of inheritance hierarchies, the reason maintenance system keeps track of the multiple credulous extensions of a hierarchy (w.r.t. each node), using a set of propositional labels for the nodes of that hierarchy. This function mimics the behavior of de Kleer's assumption-based truth maintenance system (*ATMS*) [5]; when we introduced the labeling scheme in [26], we called it the *ATMS-labeling* of the hierarchy. Like de Kleer's *ATMS*, our reason maintenance labeling keeps track of contingent conclusions—what follows if we make a certain assumption.

6.1 The Labels

For each pair of nodes, a and x, in V_Γ we define two labels: $[\![x]\!]^\Gamma_a$, the conditions under which a is an x, and $[\![\overline{x}]\!]^\Gamma_a$, the conditions under which a is not an x. The rules for labeling a hierarchy mimic the construction of paths by concatenating the edges of the hierarchy. $[\![x]\!]^\Gamma_a$, $[\![\overline{x}]\!]^\Gamma_a$

Consider an inheritance hierarchy Γ w.r.t. focus node a. Reasoning about whether a is an x, for some particular $x \in V_\Gamma$, proceeds as follows:

Let $\mathcal{C}pos_\Gamma(x)$ be the "positive children" of x in Γ—the nodes $p_i \in V_\Gamma$ with positive edges $p \cdot x \in E_\Gamma$—and let $\mathcal{C}neg_\Gamma(x)$ be the "negative children" of x in Γ—$n \in V_\Gamma \mid n \cdot \neg x \in E_\Gamma$. Suppose that, for all $p \in \mathcal{C}pos_\Gamma(x)$, a is not a p; i.e., $\Gamma \not\mathrel{|\!\!\!>} a \rightarrow p$. Then $\Gamma \not\mathrel{|\!\!\!>} a \rightarrow x$, because there is no support for x. So

$$\sim[\![\mathcal{C}pos_\Gamma(x)]\!]^\Gamma_a \quad \supset \quad \sim\widehat{x}$$

Here, we've introduced a notational shorthand: $[\![\mathcal{C}pos_\Gamma(x)]\!]^\Gamma_a$ really means $(\bigvee_{p \in \mathcal{C}pos_\Gamma(x)} [\![p]\!]^\Gamma_a)$, where an empty disjunction is to be read as $\perp$.

On the other hand, if a might be a p, for some $p \in \mathcal{C}pos_\Gamma(x)$, but $\Gamma \not\mathrel{|\!\!\!>} a \rightarrow n$, for all of x's negative children n, then a might also be an x: x's positive children provide it with a supporting argument, and there are no explicit counterarguments. That yields

$$[\![\mathcal{C}pos_\Gamma(x)]\!]^\Gamma_a \wedge (\sim[\![\mathcal{C}neg_\Gamma(x)]\!]^\Gamma_a) \quad \supset \quad [\![x]\!]^\Gamma_a$$

Finally, it could be that at least one of $\mathcal{C}pos_\Gamma(x)$ is true, but at least one of $\mathcal{C}neg_\Gamma(x)$ is also true. In this case, x is genuinely ambiguous: it has some support, but there is also a counterargument. There is no constraint corresponding to this third condition: if x is ambiguous, its value is unconstrained. When this situation arises, x becomes a choice-point, and we introduce $\widehat{x}$—the free variable corresponding to the node x—into the labeling scheme.

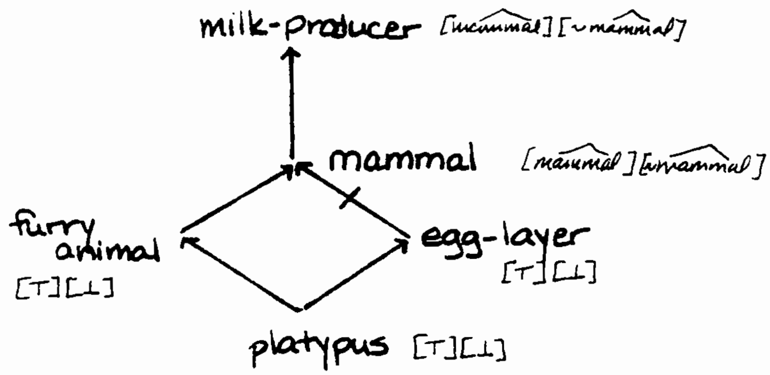

Figure 8: The labeling of Γ_2 w.r.t. *platypus*.

For example, the labeling of the "*platypus* diamond" of figure 2 w.r.t. *platypus* is shown in figure 8. In this hierarchy, $[\![mammal]\!]^{\Gamma_2}_{platypus} = \widehat{mammal}$. Although this label says very little—*platypi* are *mammals* in those extensions in which they are *mammals*—the propagation of labels yields $[\![milk\text{-}producer]\!]^{\Gamma_2}_{platypus} = \widehat{mammal}$—*platypi* are also *milk-producers* in precisely those extensions in which they are *mammals*. This is exactly the sort of contingent reasoning that we would expect our interpretation of inheritance hierarchies to capture.

The conjunction of these constraints gives us the following biconditional:

$$[\![x]\!]^{\Gamma}_{a} \equiv ([\![\mathcal{C}pos_{\Gamma}(x)]\!]^{\Gamma}_{a} \wedge (\sim [\![\mathcal{C}neg_{\Gamma}(x)]\!]^{\Gamma}_{a})) \vee ([\![\mathcal{C}pos_{\Gamma}(x)]\!]^{\Gamma}_{a} \wedge [\![\mathcal{C}neg_{\Gamma}(x)]\!]^{\Gamma}_{a} \wedge \widehat{x})$$

In addition, we have $[\![a]\!]^{\Gamma}_{a} = [\top]$—$a$ is always an a—and $\forall x$, if $x \neq a$ is a leaf of the hierarchy, $[\![x]\!]^{\Gamma}_{a} = [\perp]$—trivially, $\Gamma \not\triangleright a \rightarrow x$.

There is a similar labeling scheme for $[\![\overline{x}]\!]^{\Gamma}_{a}$, the conditions under which $\Gamma \triangleright a \not\rightarrow x$:

1. $[\![\overline{a}]\!]^{\Gamma}_{a} = [\perp]$: a is never a non-a.

2. If $x \neq a$ is a leaf of Γ, $[\![\overline{x}]\!]^{\Gamma}_{a} = [\perp]$: since x cannot be a-reachable, $\Gamma \not\triangleright a \not\rightarrow x$.

3. For a generic node, x,

$$[\![\overline{x}]\!]^{\Gamma}_{a} \equiv ((\sim [\![\mathcal{C}pos_{\Gamma}(x)]\!]^{\Gamma}_{a}) \wedge [\![\mathcal{C}neg_{\Gamma}(x)]\!]^{\Gamma}_{a}) \vee ([\![\mathcal{C}pos_{\Gamma}(x)]\!]^{\Gamma}_{a} \wedge [\![\mathcal{C}neg_{\Gamma}(x)]\!]^{\Gamma}_{a} \wedge \sim\widehat{x})$$

Theorem 5 demonstrates that these labels do keep track of all of the credulous extensions of a hierarchy at once. In the following sections, we present several useful corollaries of this theorem.

Theorem 5 *Let Γ be an inheritance hierarchy, with $a, x \in V_{\Gamma}$. Then there are mappings between the set of credulous extensions of Γ w.r.t. a, and the truth-assignments to the labels of Γ w.r.t. a, such that the truth-assignment assigns $[\![x]\!]^{\Gamma}_{a} = [\top]$ iff the corresponding extension supports $a \rightarrow x$ (resp., $[\![\overline{x}]\!]^{\Gamma}_{a}$ and $a \not\rightarrow x$).*

The mapping from extensions of Γ w.r.t. a to truth-assignments assigns a variable $\widehat{x}$ true iff a is-a x in the extension; the inverse mapping creates an extension containing those edges $s{\cdot}t \in E_\Gamma$ for which $[\![s]\!]_a^\Gamma = [\![t]\!]_a^\Gamma = [\top]$, and $s{\cdot}\neg t \in E_\Gamma$ whenever $[\![s]\!]_a^\Gamma = [\top]$ and $[\![\overline{t}]\!]_a^\Gamma = [\bot]$. A proof of the legitimacy of these mappings, and of theorem 5, may be found in appendix A.

6.2 Labels and Specificity

Theorem 5 demonstrates that the valuations of x's labels w.r.t. a correspond to the extensions in which a is-(not-)a x. It follows immediately that the labels are computing reachability:

Corollary 5.1 *Let Γ be an inheritance hierarchy, with $a, x \in V_\Gamma$. Then $[\![x]\!]_a^\Gamma$ is satisfiable iff x is positively a-reachable in Γ. (resp., $[\![\overline{x}]\!]_a^\Gamma$ and negatively a-reachable).*

Of course, reachability is not really what we are interested in. Γ supports an inference if its conclusion is reachable by an *admissible* path, not just any path. That is, the labels are not taking into account specificity—the preemption of an argument by a more specific counterargument. In this section we discuss the integration of specificity with this labeling scheme.

There are two ways that we can take specificity into account. Perhaps the most obvious is to integrate preemption into the labeling scheme itself, and to condition the acceptance of a node on its *non-preempted* children. Unfortunately, even for a relatively simple preemption scheme such as the one described in section 3.2, these labels are quite complicated. For example, the "specificity label" for $a \rightarrow x$ is

$$\left[\begin{array}{l} \left(\begin{array}{l} \displaystyle\bigvee_{p \in Cpos_\Gamma(x)} \left([\![p]\!]_a^\Gamma \wedge \sim\left(\bigwedge_{\{k|a\cdot s_1^k \cdots s_n^k \cdot p \text{ is admissible}\}} \bigvee_{\{i|s_i^k\cdot\neg x \in E_\Gamma\}} [\![s_i^k]\!]_a^\Gamma\right)\right) \\ \wedge \sim\left(\displaystyle\bigvee_{n \in Cneg_\Gamma(x)} [\![n]\!]_a^\Gamma \wedge \sim\left(\bigwedge_{\{k|a\cdot s_1^k \cdots s_n^k \cdot n \text{ is admissible}\}} \bigvee_{\{i|s_i^k\cdot x \in E_\Gamma\}} [\![s_i^k]\!]_a^\Gamma\right)\right) \end{array}\right) \\ \vee \left(\begin{array}{l} \displaystyle\bigvee_{p \in Cpos_\Gamma(x)} \left([\![p]\!]_a^\Gamma \wedge \sim\left(\bigwedge_{\{k|a\cdot s_1^k \cdots s_n^k \cdot p \text{ is admissible}\}} \bigvee_{\{i|s_i^k\cdot\neg x \in E_\Gamma\}} [\![s_i^k]\!]_a^\Gamma\right)\right) \\ \wedge \displaystyle\bigvee_{n \in Cneg_\Gamma(x)} \left([\![n]\!]_a^\Gamma \wedge \sim\left(\bigwedge_{\{k|a\cdot s_1^k \cdots s_n^k \cdot n \text{ is admissible}\}} \bigvee_{\{i|s_i^k\cdot x \in E_\Gamma\}} [\![s_i^k]\!]_a^\Gamma\right)\right) \wedge (\sim\widehat{x}) \end{array}\right) \end{array}\right]$$

For more complex preemption schemes, such as that described in section 8, the labels become still more complicated.

But we do not have to resort to embedding specificity in the labeling scheme. The specificity extension of a hierarchy provides us with precisely the information we need—a subhierarchy with the inadmissible edges removed—and simple reachability in the specificity extension is equivalent to reachability by an admissible path in the complete hierarchy:

Corollary 5.2 *Let Γ be an inheritance hierarchy, with $a, x \in V_\Gamma$, and let Σ_a^Γ be the specificity extension of Γ w.r.t. a. Then $[\![x]\!]_a^{\Sigma_a^\Gamma}$ is satisfiable iff $\Gamma \triangleright a \rightarrow x$ (resp., $[\![\overline{x}]\!]_a^{\Sigma_a^\Gamma}$ and $a \not\rightarrow x$).*

Thus, instead of using the more complex "specificity labels" on the full hierarchy, we can compute the specificity extension of the hierarchy and use the simpler labels defined in section 6.1, above.

6.3 Using the Labels

We can use this labeling for incremental update of the hierarchy. Consider, for example, the hierarchy of figure 2. If we later discover that *platypi* are *mammals*, the labeling automatically tells us that they are milk-producers as well (since $[\![\textit{milk-producer}]\!]_{\textit{platypus}}^{\Gamma_2} = [\widehat{\textit{mammal}}]$, and now $\widehat{\textit{mammal}} = \top$. In fact, we can incorporate all sorts of ambiguity resolving information—from domain-specific knowledge to updated beliefs—into this labeling simply by adding further constraints.

The complexity of this extended labeling algorithm, including further constraints, is unknown. Since it is a special case of boolean satisfiability, the problem may be $\mathcal{NP}$-hard. However, for the limited case of determining that a label is falsifiable—i.e., that there is some credulous extension in which the corresponding inference does not hold—there is a polynomial algorithm. In the next section, we exploit corollary 5.3 to solve the problem of *skeptical* inheritance: computing what *must* be true.

Corollary 5.3 *Let Γ be an inheritance hierarchy, with $a, x \in V_\Gamma$, and let Σ_a^Γ be the specificity extension of Γ w.r.t. a. Then $[\![x]\!]_a^\Gamma$ is valid (tautological) iff $a \rightarrow x$ holds in every credulous extension of Γ, and $[\![x]\!]_a^{\Sigma_a^\Gamma}$ is valid iff $a \rightarrow x$ holds in every preferred credulous extension of Γ (resp., $[\![\overline{x}]\!]_a^\Gamma$ and negatively a-reachable).*

7 Skeptical Inheritance

Up to this point, we have been discussing *credulous* inheritance—reasoning in which a conclusion that holds in *some* plausible (preferred) extension is acceptable. In this

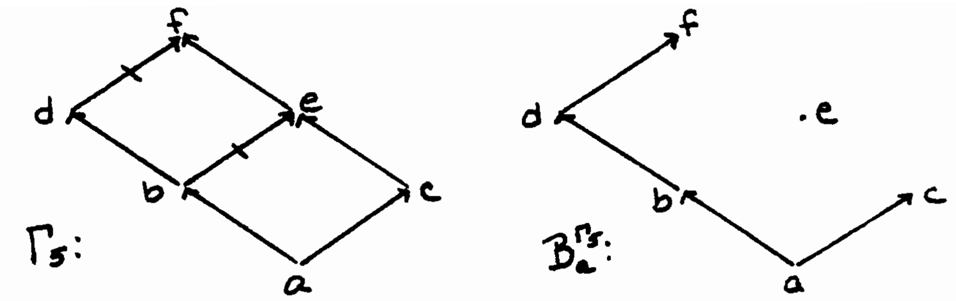

Figure 9: Applying ambiguity blocking inheritance to Γ_5 w.r.t. a yields $B_a^{\Gamma_5}$.

section, we discuss skeptical inheritance: computing those inferences that hold in *all* plausible interpretations. These conclusions are the necessary consequences of the hierarchy–what the reasoner *must* believe, no matter which possible world-state actually exists.

Formally, *ideally skeptical inheritance* supports exactly those conclusions true in every (preferred) credulous extension. The labeling scheme of section 6 provides an elegant language for expressing these two types of inheritance: credulous inheritance permits a conclusion $a \rightarrow x$ iff $[\![x]\!]_a^\Gamma$ is satisfiable; ideally skeptical inheritance permits $a \rightarrow x$ exactly when $[\![x]\!]_a^\Gamma$ is valid (resp., $a \not\rightarrow x$ and $[\![\overline{x}]\!]_a^\Gamma$). If we wish to consider only *preferred* extensions, we can simply substitute $[\![x]\!]_a^{\Sigma_a^\Gamma}$ (resp., $[\![\overline{x}]\!]_a^{\Sigma_a^\Gamma}$) here. Happily, Kautz and Selman [16] provide a polynomial-time algorithm for determining the validity of the label of a node.

In the next sections, we present two previous path-based approaches to skeptical inheritance, ambiguity blocking and ambiguity propagating inheritance. We demonstrate that these approaches are not ideally skeptical, but only approximate the always-true conclusions of a hierarchy. The failure of path-based approaches to be both sound and complete for ideally skeptical inheritance indicates the importance of reasoning about conclusions, or inferences, rather than about their supporting paths, or arguments.

7.1 Ambiguity Blocking Inheritance

The first attempt at skeptical inheritance is due to Horty, *et al.* [14];[8] Haugh [12] gives an equivalent circumscriptive definition. They argue that an ambiguous line

[8]According to Horty (personal communication), a "skeptical" approach to inheritance is one which offers a unique, unambiguous set of conclusions for any inheritance hierarchy. This differs with our intuition that "skeptical" means "unwilling to believe uncertain conclusions." In Horty's view, computing the intersection of the credulous extensions is only one way to reason "skeptically."

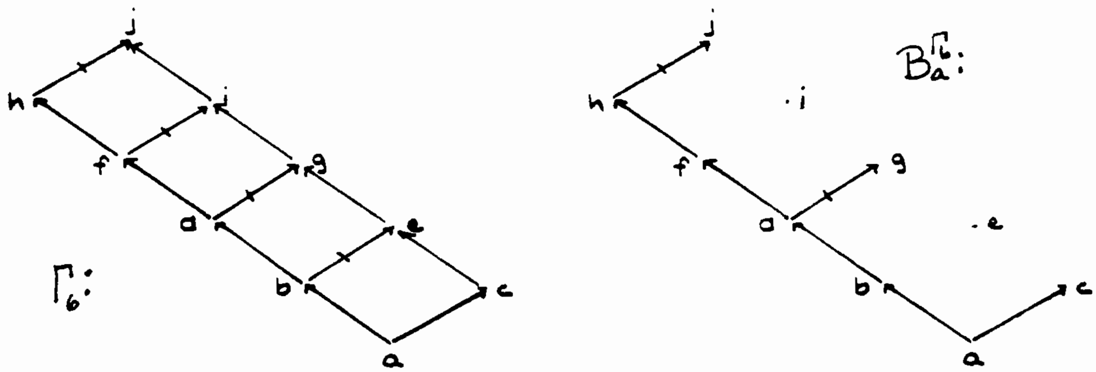

Figure 10: $B_a^{\Gamma_6}$ computes "parity."

of reasoning should not be allowed to interfere with other potential conclusions. Because this approach discontinues a line of reasoning as soon as an ambiguity has been reached, Haugh calls it *ambiguity blocking* inheritance. Although Horty, *et al.*, and Haugh describe a specific theory—including, e.g., a particular specificity criterion—ambiguity blocking inheritance defines a general approach:

> Starting from the focus node a, if a node x is truly ambiguous w.r.t. a in Γ,[9] eliminate all edges into and out of x. When the entire hierarchy has been scanned, the remaining edges form a new network, B_a^Γ, which is unambiguous w.r.t. a. This is the *ambiguity blocking* skeptical extension of Γ; ambiguity blocking inheritance concludes that a network Γ admits $a \rightarrow x$ exactly when $B_a^\Gamma \rhd a \cdot s_1 \cdots s_n \cdot x$ (resp., $a \not\rightarrow x$ and $a \cdot s_1 \cdots s_n \cdot \neg x$).

While ambiguity blocking inheritance seems reasonable, it results in some anomalous conclusions. Consider, for example, figure 9. Ambiguity blocking inheritance on Γ with focus node a determines that e is ambiguous w.r.t. a, so it eliminates all edges to and from e. In particular, it eliminates the edge $e \cdot f$, making f unambiguous w.r.t. a: $B_a^{\Gamma_5} \rhd a \rightarrow f$. This is certainly one possibility. But it is also possible that $a \rightarrow e$; and if $a \rightarrow e$, it is ambiguous whether $a \rightarrow f$—that is, a might not be an f. It is certainly not safe to assume from the ambiguity at e that the path $a \cdot b \cdot d \cdot f$ is always true. But this is precisely what ambiguity blocking inheritance does.[10]

A more severe anomaly follows from this first. Ambiguity blocking inheritance computes a kind of "parity" on the number of ambiguities in a path. According to ambiguity blocking inheritance, the network in figure 10 is skeptical as to whether a is-a e or an i, but allows the conclusions that a is-a g and a j. Similarly, this net is

[9]A node x is *truly ambiguous* in Γ w.r.t. a if $\Gamma \rhd a \rightarrow x$ and $\Gamma \rhd a \not\rightarrow x$.

[10]Horty, *et al.* originally pointed out this difference between ambiguity blocking inheritance and ideally skeptical inheritance in [14] and [31].

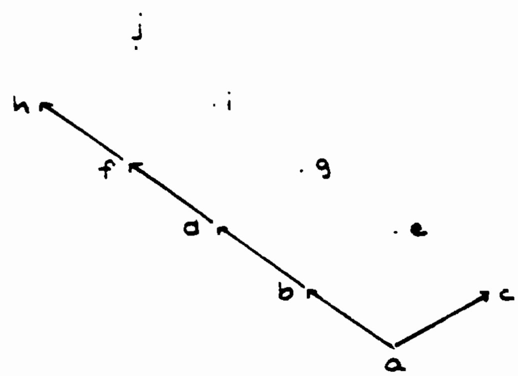

Figure 11: $\Pi_a^{\Gamma_6}$.

skeptical about whether b or f is-a j, but allows the paths from a and d to j. More than the first anomaly, this result calls into question the intuitiveness of ambiguity blocking inheritance. In any case, ambiguity blocking inheritance is *promiscuous*: there are inferences $a \rightarrow x$ such that $\mathrm{B}_a^\Gamma \rhd a \rightarrow x$, but $[\![x]\!]_a^\Gamma$ is falsifiable.

7.2 Ambiguity Propagating Inheritance

Unlike ambiguity blocking inheritance, *ambiguity propagating* inheritance allows ambiguous lines of reasoning to proceed. An argument thus cannot be certain unless there are *no* counterarguments; in contrast, ambiguity blocking inheritance considers only *unambiguous* counterarguments. Like ambiguity blocking inheritance, ambiguity propagation defines a family of algorithms. Haugh [12] gives a circumscriptive definition; in [27], we describe an $O(E_\Gamma)$ algorithm.

> Starting from the focus node a, if a node x is truly ambiguous w.r.t. a in Γ, rather than eliminating all edges to and from x, retain x but mark it as ambiguous, and continue inheriting. Although paths to and from x will not be in Π_a^Γ, they can still act as counterarguments and prevent other nodes from being unambiguous. Π_a^Γ is the subgraph of Γ with those edges $x \cdot (\neg)y \in E_\Gamma$ such that x and y are both unambiguous w.r.t. a.

For example, the cascading ambiguities of figures 9 and 10, which gave ambiguity blocking inheritance difficulty, present no problem for ambiguity propagating inheritance. Figure 11 shows $\Pi_a^{\Gamma_6}$.

In fact, Π_a^Γ never supports extra paths. Π_a^Γ computes the subhierarchy of Γ containing exactly those edges of Γ that are in every (preferred) credulous extension of Γ w.r.t. a. However, there are some *inferences* $a \rightarrow x$ that are supported by every credulous extension of Γ w.r.t. a, but have different supporting arguments in different

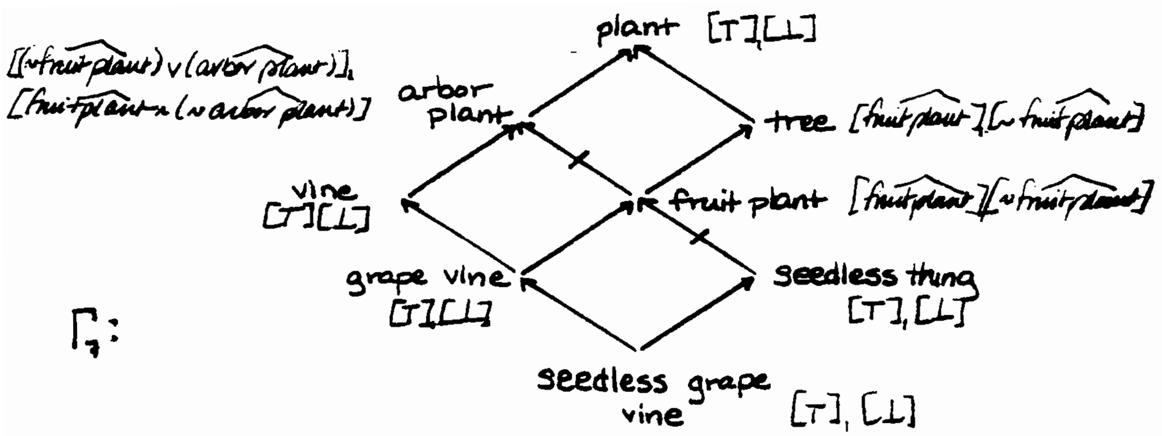

Figure 12: The intersection of credulous extensions is not a path-based notion.

extensions. In these circumstances, we need to reason about inferences rather than paths.

7.3 Ideally Skeptical Inheritance

Consider the hierarchy in figure 12.[11] Every credulous extension w.r.t. *seedless grape vine* supports the inference *seedless grape vine* $\rightarrow$*plant*, so ideally skeptical inheritance concludes that *seedless grape vines* are *plants*. Suppose, for example, that a *seedless grape vine* is a *fruit plant*; then it is a *plant*. Suppose that it is *not* a *fruit plant*; then it is unambiguously an *arbor plant*, and therefore a *plant*. In any state of the world, no matter how we resolve the ambiguities of the taxonomy, a *seedless grape vine* is a *plant*. This is reflected in the fact that $[\![plant]\!]^{\Gamma_7}_{seedless\ grape\ vine} = [\top]$.

If we wish to determine what is true in all possible worlds, we cannot avoid this kind of reasoning. There are facts which are true in all credulous extensions, but which have no justification in the intersection of those extensions. This is why we cannot generate an "ideallly skeptical extension"—no *particular* set of edges of Γ from *seedless grape vine* to *plant* is in every credulous extension, so no such path can be in the "ideally skeptical extension." Thus every path-based approach to skeptical inheritance will always be either unsound or incomplete with respect to ideally skeptical inheritance. We can only compute the always-true inferences by, in effect, reasoning about all of the credulous extensions. Fortunately, in acyclic hierarchies, such reasoning is tractable. By combining Kautz and Selman's algorithm for falsifying labels [16] with the polynomial-time computation of specificity in section 5, we *can* compute the intersection of credulous extensions—ideally skeptical inheritance—in polynomial time.

[11]Apparently, Matt Ginsberg independently proposed a hierarchy with similar properties, in which *Nixon* is always *politically motivated*.

Γ_8:

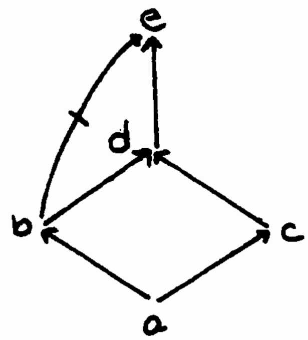

Figure 13: On- *vs.* off-path preemption.

8 Off-Path Inheritance

The inheritance theory as presented so far uses a specificity criterion called *on-path* preemption. In this type of ambiguity-resolution, an edge is inadmissible if every non-redundant path leading to that edge contains a preempting intermediary. In this section, we consider an alternate specificity criterion: that of *off-path* preemption.[12]

Off-path preemption is described in a credulous version by Sandewall [24],[13], in an ambiguity-blocking skeptical version by Horty, *et al.* [14], and in both ambiguity-blocking and ambiguity-propagating versions by Haugh [12]. Here, an edge is inadmissible if *any* edge contains a preempting intermediary. That is, off-path preemption replaces the universal quantifier—in the on-path definition of admissibility—with an existential. The difference is reflected in the hierarchy of figure 13 In Γ_8, the edge $b \cdot \neg e$ preempts the inference $a \rightarrow e$—barring both $a \cdot b \cdot d \cdot e$, and $a \cdot c \cdot d \cdot e$. On-path preemption allows $a \rightarrow e$, since $a \cdot c \cdot d \cdot e$.

Previous comparisons between on- and off-path preemption have resorted to arguments about the correct semantics of hierarchies with the topology of Γ_8. These arguments, based on the "intuitive" interpretation of Γ_8 with various names assigned to its nodes, are nothing more than *ad hoc* attempts at proof-by-example. In this section, we present the same approaches to off-path preemption as we have described for on-path, above. At the end of this section, we show that the principles brought out by couching the two types of inheritance in our framework can be constructively compared, and give some opinions as to the relative merits of on- *vs.* off-path preemption.

[12]The terms *on-* and *off-path* preemption were popularized by Touretzky, *et al.* [31].

[13]Sandewall's notion of credulous extension, however, is closer to our notion of *preferred* credulous extension.

Path-Based Definition The off-path definition of *admissibility* is:

> An edge $v \cdot (\neg)x$ is admissible in Γ w.r.t. a if *no* path $a \cdot s_1 \cdots s_n \cdot v$, contains a preempting intermediary.

In contrast, the on-path definition, in section 3.2, says that an edge is admissible if *some* path does not contain a preempting intermediary—on-path preemption uses an existential quantifier where off-path uses a universal. Because off-path preemption says that *no* path through v to x may contain a preemptor of $v \cdot (\neg)x$, this definition does not need to explicitly exclude redundant paths.

Model-Theoretic Semantics Here, we simply replace the definition of admissibility in the model-theoretic semantics of section 4.2 with that of the preceding paragraph. By replacing the definition of admissibility in both the path-based and model-theoretic approaches to inheritance, we obtain a model-theoretic semantics—complete with soundness and completeness proofs—for off-path inheritance. In general, it should be possible to obtain these results using a preemption strategy corresponding to *any* upwards theory of inheritance.[14]

Computing Specificity The algorithm for computing the specificity extension of a hierarchy w.r.t. a focus node can be adapted for off-path preemption:

Compute–Off–Path–Specificity–Extension

```
Let Γ be an inheritance hierarchy.
Sort the nodes of Γ topologically
For each focus node a                 ; Ω_a^Γ will be the specificity
   Ω_a^Γ := Γ                         ; extension of Γ w.r.t. a.
   For each node x reachable from a in Ω_a^Γ, in topological order
      ; from a to ...
      For each edge v · (¬) x in Ω_a^Γ, in reverse topological order
         ; check if v · (¬) x is admissible
         If there is some node w ∈ Cpos_{Ω_a^Γ}(x)
               such that w is positively a-reachable in Ω_a^Γ
               and v is positively w-reachable in Ω_a^Γ
      then remove the edge v · (¬) x from Ω_a^Γ
```

Like Σ_a^Γ, Ω_a^Γ has removed the inadmissible edges of Γ w.r.t. a. Computing support is reduced to verifying reachability, and the complete procedure runs in polynomial time.

---

[14]Because the definition of credulous extension is upwards, the approach presented here does not lend itself directly to the analysis of downwards inheritance theories.

Labeling Since Ω_a^Γ eliminates the inadmissible edges of Γ w.r.t. a, we can use the simple labels of section 6.1 to calculate contingent conclusions and ideally skeptical inheritance, in much the same way as we did for on-path preemption.

On- or Off-Path? By defining off-path preemption in this framework, we can compare its underlying assumptions to those of on-path preemption, presented above. By examining the differences in preference criteria, we see the principles behind these two types of reasoning: let $t_1, ..., t_n$ be many types of ts, each differing from the typical t in specific ways. For example, t might be *birds*, and the t_i might be *flightless birds*, *tiny birds*, *songbirds*, *etc.*. Each of these categories inherits most of the default properties of birds, but overrides some particular default. Now, if some particular bird inherits from several of these t_i, what is the expected outcome?

On-path preemption says that a particular bird is typical for a particular default if it is a member of *any* t_i that is normal for that default—e.g., it flies as long as it is tiny, or sings, or....—even if it is a member of the class that overrides that default—i.e. even if it is a *flightless bird*. That is, on-path preemption inherits default behavior whenever possible.

In contrast, off-path preemption *overrides* default behavior whenever possible. If a bird can *possibly* be atypical—say, flightless—then it *is* atypical for that property. Since the reason for defining $t_1, ..., t_n$ as subtypes of t is generally to enforce certain distinctions between the t_is and typical ts, anything which is a t_i should probably be assumed to be at least as unusual as t_is. That is, off-path preemption's strategy of assuming maximal atypicality is probably a more appropriate description of our intuitions.[15]

9 Other Inheritance Theories

Previous attempts to compare existing inheritance theories have proceeded largely on the basis of analysis of specific results on particular examples. Because existing theories either translate inheritance hierarchies into distinct target logics, or bury the ambiguity-resolving strategies (such as preemption) in complex path-based or translation criteria, little if any principled analysis has been possible. Touretzky, *et al.* [31] attempt to overcome this difficulty. They delimit a space of existing and potential inheritance theories, identifying several dimensions along which these theories vary. Nonetheless, this work still relies largely on examples rather than

[15] I was particularly surprised to reach this conclusion, myself. I'd been a long-standing opponent of off-path preemption, but when I'd gotten the principles worked out, I have to admit that I found myself convinced. I present this as anecdotal evidence that understanding the underlying principles really does help sort out the "right" intuitions—LAS

principles, and does not separate ambiguity-resolution from the space over which the ambiguities arise.

A second attempt to unify and compare existing inheritance theories is due to Haugh [12]. Haugh translates inheritance theories into circumscriptive meta-theories. His theory is modular with respect to the preemption axioms used, and indeed he presents preemption axioms for theories of credulous inheritance, skeptical inheritance both as presented by Horty, *et al.* [14] and in a corresponding ambiguity-propagating form, and various types of off-path preemption. However, he does not distinguish a set of interpretations corresponding to these preemption axioms, making it difficult to identify the intuitive principles to which his various axioms correspond. Haugh's confounding of meta-strategies for addressing ambiguity—credulous versus skeptical reasoning—with particular ambiguity-resolving heuristics—on- versus off-path, *etc.*—is a further symptom of this failure to identify the underlying space over which ambiguity quantifies.

The numerous preemption strategies in the literature result from differing interpretations of the notion of "subclass," or specificity. The underlying principle is that more specific information should override more general. But there is little agreement on a single definition of "more specific" at the level of the nodes and edges of an inheritance hierarchy. In the body of this paper, we present on-path preemption, which is conservative with respect to the acceptance of abnormality. In section 8, we present off-path preemption, which provides a more promiscuous approach to abnormality in inheritance.

The different treatments of *true* ambiguity are reflected in the debate over skeptical inheritance. In section 7, we describe previous approaches to skeptical inheritance, and demonstrate the shortcomings of path-based skeptical approaches.

By making various choices on the skeptical/credulous spectrum, and in on- *vs.*off-path preemption, we can describe many of the existing theories of inheritance. For example, Sandewall's theory [24] is off-path and credulous; Horty, *et al.*'s [14] is off-path and ambiguity-propagating. Touretzky's original theory [29, 30] is a downwards version of the on-path creduous theory presented in the body of this paper.

By couching other inheritance theories in the framework presented here, we obtain not merely soundness and completeness results, but a more profound understanding of the underlying principles. This allows us to make comparisons based not on *ad hoc* examples, but on intuitions that underly them.

References

[1] *Proceedings of the Sixth National Conference on Artificial Intelligence*, Seattle, Washington, July 1987. Morgan Kaufmann Publishers, Inc.

[2] *Proceedings of the Seventh National Conference on Artificial Intelligence*, St. Paul, Minnesota, August 1988. Morgan Kaufmann Publishers, Inc.

[3] Bacchus, Fahiem, A Modest, but Semantically Well Founded, Inheritance Reasoner, In *IJCAI-89* [15], 1104–1109.

[4] Boutilier, Craig, A Semantical Approach to Stable Inheritance Reasoning, In *IJCAI-89* [15], 1134–1139.

[5] de Kleer, Johan, An Assumption-based Truth Maintenance System, *Artificial Intelligence*, **28** (1986) 127–162.

[6] Doyle, Jon, A Truth-Maintenance System, *Artificial Intelligence*, **12** (1979) 231–272.

[7] Etherington, David W., More on Inheritance Hierarchies with Exceptions: Default Theories and Inferential Distance, In *AAAI-87* [1], 352–357.

[8] Etherington, David W., *Reasoning with Incomplete Information*, (Morgan Kaufmann Publishers, Inc, Los Altos, California, 1988).

[9] Etherington, David W. and Reiter, Raymond, On Inheritance Hierarchies with Exceptions, *Proceedings of the National Conference on Artificial Intelligence*, Washington, D.C., August 1983, 104–108, Morgan Kaufmann Publishers, Inc.

[10] Geffner, Hector and Verma, Tom, Inheritance = Chaining + Defeat, *Proceedings of the Fourth International Symposium on Methodologies for Intelligent Systems*, 1989.

[11] Gelfond, Michael and Przymusinska, Halina, Formalization of Inheritance Reasoning in Autoepistemic Logic, *Fundamenta Informaticae*, (To appear).

[12] Haugh, Brian A., Tractable Theories of Multiple Defeasible Inheritance in Ordinary Nonmonotonic Logics, In *AAAI-88* [2], 421–426.

[13] Horty, John F. and Thomason, Richmond H., Mixing Strict and Defeasible Inheritance, In *AAAI-88* [2], 427–432.

[14] Horty, John F., Thomason, Richmond H., and Touretzky, David S., A Skeptical Theory of Inheritance in Nonmonotonic Semantic Networks, In *AAAI-87* [1], 358–363.

[15] *Proceedings of the Eleventh International Joint Conference on Artificial Intelligence*, Detroit, Michigan, August 1989. Morgan Kaufmann Publishers, Inc.

[16] Kautz, Henry A. and Selman, Bart, Hard Problems for Simple Default Logics, Brachman, Ronald J., Levesque, Hector J., and Reiter, Raymond, (Eds.), *Proceedings of the First International Conference on Principles of Knowledge Representation and Reasoning*, Toronto, Ontario, Canada, May 1989, 189–197, Morgan Kaufmann Publishers, Inc.

[17] Krishnaprasad, T. and Kifer, Michael, An Evidence-based Framework for a Theory of Inheritance, In *IJCAI-89* [15], 1093–1098.

[18] Krishnaprasad, T., Kifer, Michael, and Warren, David S., On the Declarative Semantics of Inheritance Networks, In *IJCAI-89* [15], 1098–11103.

[19] Martins, João P. and Shapiro, Stuart C., A Model for Belief Revision, *Artificial Intelligence*, **35**(1) (1988) 25–79.

[20] McAllester, David, *An Outlook on Truth Maintenance*, MIT AI Memo 473, Massachusetts Institute of Technology Artificial Intelligence Laboratory, Cambridge, Massachusetts, 1980.

[21] McCarthy, John, Applications of Circumscription to Formalizing Common-Sense Knowledge, *Artificial Intelligence*, **28** (1986) 89–116.

[22] McDermott, Drew V., Contexts and Data Dependencies: A Synthesis, *IEEE Transactions on Pattern Analysis and Machine Intelligence*, **5** (1983) 237–246.

[23] Przymusinska, Halina and Gelfond, Michael, *Inheritance Hierarchies and Autoepistemic Logic*, Technical report, Computer Science Department, University of Texas at El Paso, 1988.

[24] Sandewall, Erik, Non-monotonic Inference Rules for Multiple Inheritance with Exceptions, *Proceedings of the IEEE*, **74** (1986) 1345–1353.

[25] Selman, Bart and Levesque, Hector, The Tractability of Path-Based Inheritance, In *IJCAI-89* [15], 1140–1145.

[26] Stein, Lynn Andrea, Skeptical Inheritance: Computing the Intersection of Credulous Extensions, In *IJCAI-89* [15].

[27] Stein, Lynn Andrea, Computing Skeptical Inheritance, Lenzerini, Maurizio, Nardi, Daniele, and Simi, Maria, (Eds.), *Inheritance Hierarchies in Knowledge Representation and Programming Languages*, (John Wiley and Sons, Ltd., London, England, 1990).

[28] Thomason, Richmond H., Horty, John F., and Touretzky, David S., *A Calculus for Inheritance in Monotonic Semantic Nets*, Technical Report CMU–CS–86–138, Department of Computer Science, Carnegie-Mellon University, Pittsburgh, PA, July 1986.

[29] Touretzky, David S., Implicit Ordering of Defaults in Inheritance Systems, *Proceedings of the National Conference on Artificial Intelligence*, Austin, Texas, August 1984, 322–325, Morgan Kaufmann Publishers, Inc.

[30] Touretzky, David S., *The Mathematics of Inheritance Systems*, (Morgan Kaufmann Publishers, Inc, Los Altos, California, 1986).

[31] Touretzky, David S., Horty, John F., and Thomason, Richmond H., A Clash of Intuitions: The Current State of Nonmonotonic Multiple Inheritance Systems, *Proceedings of the Tenth International Joint Conference on Artificial Intelligence*, Milan, Italy, August 1987, 476–482, Morgan Kaufmann Publishers, Inc.

A Proofs of Theorems

The proofs of all theorems are given for the positive cases only (e.g., $X_a^\Gamma \rhd a \to x$); the proofs of the negative cases (e.g., $X_a^\Gamma \rhd a \not\to x$) are similar.

Many of the proofs of theorems in this section are inductive. The following two definitions provide the basis for these inductive proofs.

Definition: A node $x \in V_\Gamma$ is a *leaf* of Γ if x has no children in Γ, i.e., $\mathcal{C}pos_\Gamma(y) = \mathcal{C}neg_\Gamma(y) = \emptyset$.

Definition: The *depth* of a node x in a hierarchy Γ, $\delta(x, \Gamma)$, is defined recursively:

1. If x is a leaf of Γ, $\delta(x, \Gamma) = 0$.
2. If x is not a leaf of Γ, $\delta(x, \Gamma) = 1 + \max(\{\delta(w, \Gamma) \mid w \in \mathcal{C}pos_\Gamma(x) \cup \mathcal{C}neg_\Gamma(x)\})$

Theorem 1 *Let Γ be an inheritance hierarchy, with $a, x \in V_\Gamma$. Let X_a^Γ be a credulous extension of Γ w.r.t. a, and let $\widehat{X_a^\Gamma}$ be the propositional theory corresponding to X_a^Γ. The following are equivalent:*

1. $X_a^\Gamma \rhd a \to x$ (resp., $a \not\to x$).

2. x is positively a-reachable in X_a^Γ(resp., negatively a-reachable).

3. $\widehat{X_a^\Gamma} \vdash \hat{x}$ (resp., $\sim\hat{x}$).

4. $\widehat{X_a^\Gamma} \models \hat{x}$ (resp., $\sim\hat{x}$).

Proof:

$1 \Leftrightarrow 2$: By definition, a hierarchy $X_a^\Gamma \rhd a \to x$ iff x is positively a-reachable in X_a^Γ by an *admissible* path; but since a credulous extension is unambiguous, every path is admissible.

$2 \Leftrightarrow 3$: $\Rightarrow$: Assume that x is positively a-reachable in X_a^Γ, i.e. $(\exists\, s_1 \cdots s_n)\ a \cdot s_1 \cdots s_n \cdot x \in E_{X_a^\Gamma}$. Then

$$\hat{a} \quad \wedge \bigwedge_{1 \le i \le n-1} (\hat{s}_i \supset \hat{s}_{i+1}) \quad \wedge \quad (\hat{s}_n \supset \hat{x}) \tag{1}$$

is a sub-theory of $\widehat{X_a^\Gamma}$; since a theory validates any subtheory, $\widehat{X_a^\Gamma} \vdash$ (1) and its components. So $\widehat{X_a^\Gamma} \vdash \hat{a}$; $\widehat{X_a^\Gamma} \vdash \hat{a} \supset \hat{s}_1$, therefore $\widehat{X_a^\Gamma} \vdash \hat{s}_1$; similarly $\widehat{X_a^\Gamma} \vdash \hat{s}_1 \supset \hat{s}_2$, so $\widehat{X_a^\Gamma} \vdash \hat{s}_2$; *etc.* So $\widehat{X_a^\Gamma} \vdash \hat{x}$.

$\Leftarrow$: Assume x is not positively a-reachable in X^Γ_a, but $\widehat{X^\Gamma_a} \vdash \widehat{x}$. Since $\widehat{X^\Gamma_a} \vdash \widehat{x}$, $\widehat{X^\Gamma_a} \cup \{\sim\widehat{x}\} \vdash \bot$. In particular, $\widehat{X^\Gamma_a} \cup \{\sim\widehat{x}\} \vdash \sim\widehat{a}$; and by the deduction theorem, $\widehat{X^\Gamma_a} \vdash \sim\widehat{x} \supset \sim\widehat{a}$. Since $\widehat{X^\Gamma_a}$ is a conjunction of clauses of the form $\widehat{s} \supset \widehat{t}$ (and the singleton $\widehat{a}$), $\widehat{X^\Gamma_a} \vdash \sim\widehat{x} \supset \sim\widehat{a}$ means that either $\widehat{x} = \widehat{a}$ or there is some conjunct $\sim\widehat{x} \supset \widehat{s}_1$ (or $\widehat{t_1} \supset \widehat{x}$, where $\widehat{t_1} = \sim\widehat{s}_1$) in $\widehat{X^\Gamma_a}$ such that $\widehat{X^\Gamma_a} \vdash \widehat{s}_1 \supset \sim\widehat{a}$. By a similar argument, we can show that either $\widehat{s}_1 = \widehat{a}$ or there is some $\widehat{s}_2$ such that $\widehat{s}_1 \supset \widehat{s}_2$ is in $\widehat{X^\Gamma_a}$ and $\widehat{X^\Gamma_a} \vdash \widehat{s}_2 \supset \sim\widehat{a}$, etc., so that there must be a chain of conjuncts $\widehat{x} \supset \widehat{s}_1, \widehat{s}_1 \supset \widehat{s}_2, ..., \widehat{s}_{n-1} \supset \widehat{s}_n, \widehat{s}_n \supset \sim\widehat{a}$ in $\widehat{X^\Gamma_a}$. But these are just $\widehat{a} \supset \widehat{t_n}, \widehat{t_n} \supset \widehat{t_{n-1}}, \ldots, \widehat{t_2} \supset \widehat{t_1}, \widehat{t_1} \supset \widehat{x}$, and since every conjunct of $\widehat{X^\Gamma_a}$ corresponds to an edge of $E_{X^\Gamma_a}$, this implies a sequence of edges $a \cdot t_1 \cdots t_n \cdot x$—making x a-reachable in X^Γ_a; contradiction.

$3 \Leftrightarrow 4$: This follows from the soundness and completeness of propositional logic: a formula is derivable from a propositional theory iff every model of that theory entails that formula.

The following lemma is used in the proof of several theorems, below. It says that any unambiguous a-connected subhierarchy of Γ can be extended to make it maximal—i.e., a credulous extension of Γ w.r.t. a.

Lemma 1 *Let Γ be an inheritance hierarchy, and let $\{v_i \cdot x_i\}$ be an unambiguous a-connected subhierarchy of Γ. with $a \cdot s_1 \cdots s_n \cdot x \in E_\Gamma$. Then there is an extension (of Γ w.r.t. a) $X^\Gamma_a = \mathit{Extend}(\{v_i \cdot (\neg)x_i\}, \Gamma)$ containing $\{v_i \cdot (\neg)x_i\}$ (i.e. $\{v_i \cdot (\neg)x_i\} \in E_{X^\Gamma_a}$) Further, if $\{v_i \cdot (\neg)x_i\}$ are all admissible in Γ w.r.t. a, then there is a preferred extension of Γ w.r.t. a containing $\{v_i \cdot x_i\}$.*

Proof: Construct $X^\Gamma_a = \mathit{Extend}(\{v_i \cdot x_i\}, \Gamma)$ as follows:

1. Let $\{v_i \cdot (\neg)x_i\} \in E_{X^\Gamma_a}$.
2. For each node x, in topological order, add the edge $v \cdot (\neg)x \in E_\Gamma$ to $E_{X^\Gamma_a}$ if
   (a) v is positively a-reachable in $\mathit{Extend}(\{v_i \cdot x_i\}, \Gamma)$ (as constructed so far),
   (b) $v \cdot (\neg)x$ is admissible in Γ w.r.t. a, and
   (c) adding $v \cdot (\neg)x$ won't make $\mathit{Extend}(\{v_i \cdot x_i\}, \Gamma)$ (as constructed so far) ambiguous.

By construction, $X^\Gamma_a = \mathit{Extend}(\{v_i \cdot (\neg)x_i\}, \Gamma)$ is a credulous extension of Γ w.r.t. a: it is unambiguous and a-connected, and any edge that could be added without contradicting one of these two properties has been.

Further, if every edge in $\{v_i \cdot (\neg)x_i\}$ is admissible in Γ w.r.t. a, we claim that $X_a^\Gamma = \mathit{Extend}(\{v_i \cdot x_i\}, \Gamma)$ is a preferred credulous extension of Γ w.r.t. a.

Let $Y_a^\Gamma \neq X_a^\Gamma$ be a credulous extension of Γ w.r.t. a. Assume by way of contradiction that $Y_a^\Gamma \preceq X_a^\Gamma$. Then there are some nodes $s, t \in V_{X_a^\Gamma}$ such that

1. X_a^Γ and Y_a^Γ agree on all edges whose endpoints topologically precede t in any topological order,
2. $s \cdot (\neg)t$ is inadmissible in Γ w.r.t. a,
3. $X_a^\Gamma \rhd a \cdot s_1 \cdots s_m \cdot s \cdot (\neg)t$, $(m \geq 0)$, and
4. $Y_a^\Gamma \not\rhd a \cdot s_1 \cdots s_m \cdot s \cdot (\neg)t$

If $s \cdot (\neg)t$ is inadmissible in Γ w.r.t. a, then either

1. $s \cdot (\neg)t$ is in $\{v_i \cdot (\neg)x_i\}$. But then it must be admissible, since $\{v_i \cdot (\neg)x_i\}$ is admissible in Γ; or
2. it was added during the construction. But then it is admissible in Γ w.r.t. a, by condition 2b.

So X_a^Γ does not contain an inadmissible edge, and it is is therefore a preferred extension of Γ w.r.t. a.

The following lemma says that a non-redundant unsupported path is not in any preferred credulous extension.

Lemma 2 *Let Γ be an inheritance hierarchy, and let $a \cdot s_1 \cdots s_n \cdot (\neg)x \in E_\Gamma$, with none of $a \cdot s_1 \cdots s_n$ redundant in Γ w.r.t. a. If $\Gamma \not\rhd a \cdot s_1 \cdots s_n \cdot (\neg)x$, then no preferred extension of Γ w.r.t. a—$X_a^\Gamma \in \mathcal{P}ref(\Gamma, a)$—supports $a \cdot s_1 \cdots s_n \cdot (\neg)x$.*

Proof: The proof proceeds by induction on n:

Base Case: If $n = 0$, $\Gamma \rhd a \cdot x$: $a \cdot x$ is trivially admissible in Γ w.r.t. a, since it contains no intermediate nodes to preempt it. Therefore, the theorem is trivially true for $n = 0$.

Induction Hypothesis: Assume that whenever $\Gamma \not\rhd a \cdot s_1 \cdots s_n \cdot x$, $n < k$, no preferred extension of Γ w.r.t. a supports $a \cdot s_1 \cdots s_n \cdot x$.

Induction Step: Let $a \cdot s_1 \cdots s_n \cdot x \in E_\Gamma$, with none of $a \cdot s_1 \cdots s_k$ redundant in Γ w.r.t. a; let $\Gamma \not\triangleright a \cdot s_1 \cdots s_k \cdot x$, and let $a \cdot s_1 \cdots s_n \cdot x \in E_{\mathrm{X}^\Gamma_a}$. If any of $a \cdot s_1 \cdots s_k$ is inadmissible in Γ w.r.t. a, then $\Gamma \not\triangleright a \cdot s_1 \cdots s_k$, and by the induction hypothesis, $\mathrm{X}^\Gamma_a \notin \mathcal{P}ref(\Gamma, a)$. But if all of $a \cdot s_1 \cdots s_k$ are admissible in Γ, then some intermediate node $a, s_1, ..., s_{k-1}$ must be a preemptor of $s_k \cdot x$. Call it s_i. Then there is an edge $s_i \cdot \neg x \in E_\Gamma$. Let $\mathrm{Y}^\Gamma_a = \mathit{Extend}(\mathrm{X}^\Gamma_a|_x \cup \{s_i \cdot \neg x\}, \Gamma)$, where $\mathrm{X}^\Gamma_a|_x$ is

$$\{s_i \cdot (\neg)t_i \mid s_i \cdot (\neg)t_i \in E_{\mathrm{X}^\Gamma_a} \text{ and } t_i \text{ precedes } x \text{ in some topological sort of } \Gamma\}$$

Then Y^Γ_a is a credulous extension of Γ w.r.t. a (by lemma 1, since $\mathrm{X}^\Gamma_a|_x \cup \{s_i \cdot \neg x\}$ is an unambiguous subhierarchy of Γ w.r.t. a). Further, $\mathrm{Y}^\Gamma_a \preceq \mathrm{X}^\Gamma_a$:

1. X^Γ_a and Y^Γ_a agree on all edges whose endpoints topologically precede x in any topological order,
2. $s_k \cdot x$ is inadmissible in Γ w.r.t. a,
3. $\mathrm{X}^\Gamma_a \triangleright a \cdot s_1 \cdots s_k \cdot x$, and
4. $\mathrm{Y}^\Gamma_a \not\triangleright a \cdot s_1 \cdots s_k \cdot x$.

So $\mathrm{X}^\Gamma_a \notin \mathcal{P}ref(\Gamma, a)$.

Theorem 2 *Let Γ be an inheritance hierarchy, with $a, x \in V_\Gamma$, and let $\mathcal{P}ref(\Gamma, a)$be the set of preferred extensions of Γ w.r.t. a. Then $\Gamma \triangleright a \rightarrow x$ iff there is some preferred extension $\mathrm{X}^\Gamma_a \in \mathcal{P}ref(\Gamma, a)$ such that $\widehat{\mathrm{X}^\Gamma_a} \models \widehat{x}$ (resp., $a \not\rightarrow x$ and $\sim\widehat{x}$).*

Proof:

$\Rightarrow$: This follows directly from lemma 1 and theorem 1, conditions 1 and 4.

$\Leftarrow$: Assume that $\Gamma \not\triangleright a \rightarrow x$, i.e. that there is no admissible path from a to x in Γ. We demonstrate that x is not positively a-reachable in any preferred credulous extension X^Γ_a of Γ w.r.t. a; by By theorem 1, conditions 2 and 4, this means that there is no preferred credulous extension of Γ w.r.t. a for which $\mathrm{X}^\Gamma_a \models \widehat{x}$.

1. Assume that there is *no* path from a to x in Γ. Then there is no path from a to x in *any* credulous extension of Γ w.r.t. a (since every credulous extension is a subhierarchy).
2. Assume that x is a-reachable in Γ, but not by an admissible path. Let $a \cdot s_1 \cdots s_n \cdot x$ be an inadmissible path in Γ w.r.t. a; i.e., either some edge of $a \cdot s_1 \cdots s_n$ is redundant in Γ w.r.t. a, or some edge of the path as a whole is inadmissible. Let X^Γ_a be a credulous extension of Γw.r.t. a, and let $a \cdot s_1 \cdots s_n \cdot x \in E_{\mathrm{X}^\Gamma_a}$. Assume by way of contradiction that $\mathrm{X}^\Gamma_a \in \mathcal{P}ref(\Gamma, a)$.

If no edge of $a \cdot s_1 \cdots s_n$ is redundant, then $\mathrm{X}_a^\Gamma \notin \mathcal{P}ref(\Gamma, a)$, by lemma 2. So some edge must be redundant; call it $s_i \cdot s_{i+1}$. Then there is a sequence of edges $s_i \cdot q_1 \cdots q_m \cdot s_{i+1} \in E_\Gamma$, $(m > 0)$, such that $\Gamma \not\rhd a \not\rightarrow q_j$, $(1 \leq j \leq n)$. Now either

(a) $s_i \cdot q_1 \cdots q_m \cdot s_{i+1} \in E_{\mathrm{X}_a^\Gamma}$. Then by lemma 2, $\mathrm{X}_a^\Gamma \notin \mathcal{P}ref(\Gamma, a)$; or

(b) $s_i \cdot q_1 \cdots q_m \cdot s_{i+1} \notin E_{\mathrm{X}_a^\Gamma}$. Then q_j must be negatively a-reachable in X_a^Γ, for some j—or we could add this path without making X_a^Γ ambiguous, so it is not maximal and therefore not a credulous extension. But $\Gamma \not\rhd a \not\rightarrow q_j$. Here, we argue inductively that because q_j is negatively a-reachable in X_a^Γ, $\mathrm{X}_a^\Gamma \notin \mathcal{P}ref(\Gamma, a)$.

Theorem 3 *Let Γ be an inheritance hierarchy, with $a, x \in V_\Gamma$, and let Σ_a^Γ be the specificity extension of Γ w.r.t. a. Then the following are equivalent:*

1. *$\Sigma_a^\Gamma \rhd a \rightarrow x$ (resp., $a \not\rightarrow x$).*
2. *x is positively a-reachable in Σ_a^Γ (resp., negatively a-reachable).*
3. *There is an extension $\mathrm{X}_a^{\Sigma_a^\Gamma}$ of Σ_a^Γ w.r.t. a, such that $\mathrm{X}_a^{\Sigma_a^\Gamma} \rhd a \rightarrow x$ (resp., $a \not\rightarrow x$).*

Proof:

$1 \Leftrightarrow 2$: $\Rightarrow$: If $\Sigma_a^\Gamma \rhd a \rightarrow x$, then x is positively a-reachable in Σ_a^Γ by an admissible path, so x is certainly positively a-reachable in Σ_a^Γ.

$\Leftarrow$: Let $\Sigma_a^\Gamma \not\rhd a \rightarrow x$. Then every non-redundant path $a \cdot s_1 \cdots s_n \cdot v \cdot x \in E_{\Sigma_a^\Gamma}$ contains a preemptor of $v \cdot x$. But if $v \cdot x \in E_{\Sigma_a^\Gamma}$, v is positively a-reachable by a non-redundant path containing no preemptors of $v \cdot x$ (or $v \cdot x$ would have been removed from Σ_a^Γ); so $v \cdot x \notin E_{\Sigma_a^\Gamma}$. Since our choice of v was arbitrary, the argument holds for any v; and x is not positively a-reachable in Σ_a^Γat all.

$2 \Leftrightarrow 3$: $\Rightarrow$: If x is positively a-reachable in Σ_a^Γ, then there is some sequence of edges $a \cdot s_1 \cdots s_n \cdot x \in E_{\Sigma_a^\Gamma}$. By Lemma 1, there is an extension of Σ_a^Γ, $\mathrm{X}_a^{\Sigma_a^\Gamma}$, such that $\mathrm{X}_a^{\Sigma_a^\Gamma}$ contains $a \cdot s_1 \cdots s_n \cdot x$; by Theorem 1, this means that $\mathrm{X}_a^{\Sigma_a^\Gamma} \rhd a \rightarrow x$.

$\Leftarrow$: $\mathrm{X}_a^{\Sigma_a^\Gamma} \rhd a \rightarrow x$ means that there is a path $a \cdot s_1 \cdots s_n \cdot x \in E_{\mathrm{X}_a^{\Sigma_a^\Gamma}}$. But $E_{\mathrm{X}_a^{\Sigma_a^\Gamma}} \subseteq E_{\Sigma_a^\Gamma}$, so $a \cdot s_1 \cdots s_n \cdot x \in E_{\Sigma_a^\Gamma}$, and x is positively a-reachable in Σ_a^Γ.

Theorem 4 *Let Γ be an inheritance hierarchy, with $a, x \in V_\Gamma$, and let Σ_a^Γ be the specificity extension of Γ w.r.t. a. Then $\Gamma \rhd a \rightarrow x$ iff x is positively a-reachable in Σ_a^Γ (resp., $a \not\rightarrow x$ and negatively a-reachable).*

Proof: The proof proceeds by induction on $\delta(x, \Gamma)$:

Base Cases: If x is a leaf of Γ, then either $x = a$—and $\Gamma \rhd a \rightarrow a$ and a is positively a-reachable in Σ_a^Γ—or $x \neq a$—and $\Gamma \not\rhd x \rightarrow a$, nor is x positively a-reachable in Σ_a^Γ.

Induction Hypothesis: Assume that, for all nodes $x \in V_\Gamma$ with $\delta(x, \Gamma) < n$, $\Gamma \rhd a \rightarrow x$ iff x is positively a-reachable in Σ_a^Γ.

Induction Step: Let $y \in V_\Gamma$, $\delta(y, \Gamma) = n$.

$\Rightarrow$: Let $\Gamma \rhd a \rightarrow y$. Then there is some non-redundant sequence of admissible edges $a \cdot s_1 \cdots s_n \cdot y \in E_\Gamma$ containing no preemptors of $s_n \cdot y$. In particular, $a, s_1, \ldots, s_n$ are positively a-reachable in Γ by a sequence of non-redundant admissible edges—$\Gamma \rhd a \rightarrow a, s_1, \ldots, s_n$. Further, $\delta(a, \Gamma) < \delta(s_1, \Gamma) < \ldots < \delta(s_n, \Gamma) < \delta(y, \Gamma) = n$, so by the induction hypothesis, $a, s_1, \ldots, s_n$ are positively a-reachable in Σ_a^Γ. Then if we remove the preemptors of $s_n \cdot y$ from Σ_a^Γ, the path $a \cdot s_1 \cdots s_n \cdot y$ will remain; it contains no redundant edges, so we do not delete the edge $s_n \cdot y$ from Σ_a^Γ, and y is positively a-reachable in Σ_a^Γ.

$\Leftarrow$: Let $\Gamma \not\rhd a \rightarrow y$. Then every non-redundant sequence of admissible edges $a \cdot s_1 \cdots s_n \cdot y \in E_\Gamma$ contains a preemptors of $s_n \cdot y$.

1. Assume that y is not positively a-reachable in Γ. Then certainly, y is not positively a-reachable in Σ_a^Γ.
2. Assume that y is positively a-reachable in Γ, but for every positive child p of y, $\Gamma \not\rhd a \rightarrow p$. Then by the induction hypothesis, since $\delta(p, \Gamma) < \delta(y, \Gamma) = n$, p is not positively a-reachable in Σ_a^Γ; so y is not positively a-reachable in Σ_a^Γ.
3. Assume that there is some node $p \in Cpos_\Gamma(y)$ such that $\Gamma \rhd a \rightarrow p$. Then by the induction hypothesis, p is positively a-reachable in Σ_a^Γ—say, by $a \cdot s_1 \cdots s_n \cdot p$, with all of these edges admissible and none of them, except possibly $s_n \cdot p$ redundant in Γ w.r.t. a (there must be some such admissible non-redundant path, or $s_n \cdot p$ would have been pruned from Σ_a^Γ). Also, if $s_n \cdot p$ is redundant in Γ w.r.t. a, then $p \cdot y$ is pruned from Σ_a^Γ, and $p \notin Cpos_{\Sigma_a^\Gamma}(y)$; either there is some other $p' \in Cpos_\Gamma(y) \cup Cpos_{\Sigma_a^\Gamma}(y)$—w.l.o.g. , we consider this p' instead—or there is not, and we revert to one of the previous cases.

We claim that one of $a, s_1, ..., s_n$ is a preemptor of $p \cdot y$ in Γ. Assume by way of contradiction that none of $a, s_1, ..., s_n \in \mathcal{C}neg_\Gamma(y)$. Then y is positively a-reachable in Σ_a^Γ by a non-redundant (in Γ w.r.t. a) sequence of admissible (in Γ w.r.t. a) edges containing no preemptors (in Γ) of the final edge. But then $\Gamma \rhd a \rightarrow y$; contradiction. So one of $a, s_1, ..., s_n$ is a preemptor of $p \cdot y$ in Γ; and when the preemptors of $p \cdot y$ are removed from Σ_a^Γ, p is not positively a-reachable by a non-redundant (in Σ_a^Γ w.r.t. a) path in Σ_a^Γ, so $p \cdot y$ is pruned. By the same argument, every edge $p_i \cdot y \in E_\Gamma$ is pruned from Σ_a^Γ, and y is not positively a-reachable in Σ_a^Γ.

The next lemma says that if x is (positively or negatively) a-reachable, then $[\![x]\!]_a^\Gamma \equiv (\sim[\![\overline{x}]\!]_a^\Gamma)$.

Lemma 3 *Let Γ be an inheritance hierarchy, with $a, x \in V_\Gamma$. Let $\mathcal{L}(\Gamma, a)$ be the ATMS labeling of Γ w.r.t. a, with $[\![x]\!]_a^\Gamma, [\![\overline{x}]\!]_a^\Gamma \in \mathcal{L}(\Gamma, a)$. Then*

1. $[\![x]\!]_a^\Gamma = (\sim[\![\overline{x}]\!]_a^\Gamma) \wedge ([\![\mathcal{C}pos_\Gamma(x)]\!]_a^\Gamma \vee [\![\mathcal{C}neg_\Gamma(x)]\!]_a^\Gamma)$
2. $[\![\overline{x}]\!]_a^\Gamma = (\sim[\![x]\!]_a^\Gamma) \wedge ([\![\mathcal{C}pos_\Gamma(x)]\!]_a^\Gamma \vee [\![\mathcal{C}neg_\Gamma(x)]\!]_a^\Gamma)$.

Proof:

$[\![\mathcal{C}pos_\Gamma(x)]\!]_a^\Gamma$	$[\![\mathcal{C}neg_\Gamma(x)]\!]_a^\Gamma$	$\widehat{x}$	$[\![x]\!]_a^\Gamma$	$(\sim[\![\overline{x}]\!]_a^\Gamma) \wedge ([\![\mathcal{C}pos_\Gamma(x)]\!]_a^\Gamma \vee [\![\mathcal{C}neg_\Gamma(x)]\!]_a^\Gamma)$
$\perp$	$\perp$	$\perp$	$\perp$	$\perp$
$\perp$	$\perp$	$\top$	$\perp$	$\perp$
$\perp$	$\top$	$\perp$	$\perp$	$\perp$
$\perp$	$\top$	$\top$	$\perp$	$\perp$
$\top$	$\perp$	$\perp$	$\top$	$\top$
$\top$	$\perp$	$\top$	$\top$	$\top$
$\top$	$\top$	$\perp$	$\top$	$\top$
$\top$	$\top$	$\top$	$\perp$	$\perp$

Similarly,

$[\![\mathcal{C}pos_\Gamma(x)]\!]_a^\Gamma$	$[\![\mathcal{C}neg_\Gamma(x)]\!]_a^\Gamma$	$\widehat{x}$	$[\![\overline{x}]\!]_a^\Gamma$	$(\sim[\![x]\!]_a^\Gamma) \wedge ([\![\mathcal{C}pos_\Gamma(x)]\!]_a^\Gamma \vee [\![\mathcal{C}neg_\Gamma(x)]\!]_a^\Gamma)$
$\perp$	$\perp$	$\perp$	$\perp$	$\perp$
$\perp$	$\perp$	$\top$	$\perp$	$\perp$
$\perp$	$\top$	$\perp$	$\top$	$\top$
$\perp$	$\top$	$\top$	$\top$	$\top$
$\top$	$\perp$	$\perp$	$\perp$	$\perp$
$\top$	$\perp$	$\top$	$\perp$	$\perp$
$\top$	$\top$	$\perp$	$\perp$	$\perp$
$\top$	$\top$	$\top$	$\top$	$\top$

Theorem 5 *Let Γ be an inheritance hierarchy, with $a, x \in V_\Gamma$. Then there are mappings between the set of credulous extensions of Γ w.r.t. a, and the truth-assignments to the labels of Γ w.r.t. a, such that the truth-assignment assigns $[\![x]\!]^\Gamma_a = [\top]$ iff the corresponding extension supports $a \to x$ (resp., $[\![\overline{x}]\!]^\Gamma_a$ and $a \not\to x$).*

Proof:

Let $\mathbf{T}$ be the set of all possible truth-assignments to the variable corresponding to the nodes of V_Γ. If $\mathcal{T} \in \mathbf{T}$ is a truth-assignment to the variables corresponding to V_Γ, and $[\![x]\!]^\Gamma_a$ is a propositional formula over these variables, then $\mathcal{T}([\![x]\!]^\Gamma_a)$ is the truth-value assigned to the formula $[\![x]\!]^\Gamma_a$ under $\mathcal{T}$ in the traditional sense.

Let $\{X^\Gamma_a\}$ be the set of credulous extensions of Γ w.r.t. a.

Define $\varphi : \mathbf{T} \to \{X^\Gamma_a\}$ as

$$\begin{aligned} E_{\varphi(\mathcal{T})} &= \{s \cdot t \in E_\Gamma \mid \mathcal{T}([\![s]\!]^\Gamma_a) = \mathcal{T}([\![t]\!]^\Gamma_a) = [\top]\} \\ &\quad \cup \{s \cdot \neg t \in E_\Gamma \mid \mathcal{T}([\![s]\!]^\Gamma_a) = [\top] \textit{ and } \mathcal{T}([\![\overline{t}]\!]^\Gamma_a) = [\top]\} \\ V_{\varphi(\mathcal{T})} &= \{t \mid s \cdot (\neg)t \in E_\Gamma \text{ or } t \cdot (\neg)v \in E_\Gamma\} \end{aligned}$$

We claim that $\varphi(\mathcal{T})$ is a credulous extension of Γ w.r.t. a.

1. Certainly, $\varphi(\mathcal{T})$ is a subhierarchy of Γ.

2. $\varphi(\mathcal{T})$ is unambiguous: by lemma 3, $[\![t]\!]^\Gamma_a \wedge [\![\overline{t}]\!]^\Gamma_a$ is unsatisfiable:

$$\begin{aligned} [\![t]\!]^\Gamma_a \wedge [\![\overline{t}]\!]^\Gamma_a &= [\![t]\!]^\Gamma_a \wedge ((\sim [\![t]\!]^\Gamma_a) \wedge ([\![\mathcal{C}pos_\Gamma(t)]\!]^\Gamma_a \vee [\![\mathcal{C}neg_\Gamma(t)]\!]^\Gamma_a)) \\ &= ([\![t]\!]^\Gamma_a \wedge (\sim [\![t]\!]^\Gamma_a)) \wedge ([\![\mathcal{C}pos_\Gamma(t)]\!]^\Gamma_a \vee [\![\mathcal{C}neg_\Gamma(t)]\!]^\Gamma_a) \\ &= \perp \wedge ([\![\mathcal{C}pos_\Gamma(t)]\!]^\Gamma_a \vee [\![\mathcal{C}neg_\Gamma(t)]\!]^\Gamma_a) \\ &= \perp \end{aligned}$$

   So it is not the case that $s_1 \cdot t$ and $s_2 \cdot \neg t$ are both in $E_{\varphi(\mathcal{T})}$, for any s_1, s_2, t.

3. $\varphi(\mathcal{T})$ is maximal: Assume by way of contradiction that there is some edge $s \cdot (\neg)t \in E_\Gamma - E_{\varphi(\mathcal{T})}$, and $\langle V_{\varphi(\mathcal{T})}, E_{\varphi(\mathcal{T})} \cup \{s \cdot (\neg)t\}\rangle$ is unambiguous and a-connected. Then either $\mathcal{T}([\![s]\!]^\Gamma_a) \neq [\top]$, or $\mathcal{T}([\![t]\!]^\Gamma_a)$ (resp., $\mathcal{T}([\![\overline{t}]\!]^\Gamma_a)) \neq [\top]$.

   $\mathcal{T}([\![s]\!]^\Gamma_a) \neq \top$: Then s is not positively a-reachable in $\varphi(\mathcal{T})$, and adding $s \cdot (\neg)t$ would not preserve a-connectedness.

   $\mathcal{T}([\![t]\!]^\Gamma_a)$ (resp., $\mathcal{T}([\![\overline{t}]\!]^\Gamma_a)) \neq \top$: But $\mathcal{T}([\![t]\!]^\Gamma_a)$ is either $[\top]$ or $[\perp]$—$\mathcal{T}$ assigns a truth-value to every variable corresponding to a node in V_Γ, and $[\![t]\!]^\Gamma_a$ is a propositional formula over these variables. So $\mathcal{T}([\![t]\!]^\Gamma_a) = [\perp]$. But

$$\begin{aligned} \mathcal{T}([\![t]\!]^\Gamma_a) &= (\mathcal{T}([\![\mathcal{C}pos_\Gamma(t)]\!]^\Gamma_a) \wedge (\sim \mathcal{T}([\![\mathcal{C}neg_\Gamma(t)]\!]^\Gamma_a))) \\ &\quad \vee (\mathcal{T}([\![\mathcal{C}pos_\Gamma(t)]\!]^\Gamma_a) \wedge \mathcal{T}([\![\mathcal{C}neg_\Gamma(t)]\!]^\Gamma_a) \wedge \mathcal{T}(\hat{t})) \\ &= \mathcal{T}([\![\mathcal{C}pos_\Gamma(t)]\!]^\Gamma_a) \wedge ((\sim \mathcal{T}([\![\mathcal{C}neg_\Gamma(t)]\!]^\Gamma_a)) \vee \mathcal{T}(\hat{t})) \\ &= [\perp] \end{aligned}$$

In particular,

$$\mathcal{T}([\![\mathcal{C}pos_\Gamma(t)]\!]_a^\Gamma) \;=\; [\bot]$$

This means that $\forall p \in \mathcal{C}pos_\Gamma(t)$, $\mathcal{T}([\![p]\!]_a^\Gamma) = [\bot]$. In particular, $\mathcal{T}([\![s]\!]_a^\Gamma) = [\bot] \neq [\top]$, and s is not positively a-reachable in $\varphi(\mathcal{T})$, so adding $s{\cdot}t$ would not preserve a-connectedness. By a similar argument, $\mathcal{T}([\![\bar{t}]\!]_a^\Gamma) \neq [\top]$ means that $\mathcal{T}([\![\mathcal{C}neg_\Gamma(t)]\!]_a^\Gamma) = [\bot]$; $\mathcal{T}([\![s]\!]_a^\Gamma) = [\bot] \neq [\top]$; s is not positively a-reachable in $\varphi(\mathcal{T})$; and adding $s{\cdot}\neg t$ would not preserve a-connectedness.

Then $\mathcal{T}([\![x]\!]_a^\Gamma) = [\top]$ iff x is positively a-reachable in $\varphi(\mathcal{T})$, so by theorem 1, $\mathcal{T}([\![x]\!]_a^\Gamma) = [\top]$ iff $\varphi(\mathcal{T}) \rhd a \rightarrow x$.

Define $\psi : \{\mathrm{X}_a^\Gamma\} \rightarrow \mathsf{T}$ as

$$\mathcal{T}(\hat{x}) \;=\; \begin{cases} \top, & \mathit{if}\ \mathrm{X}_a^\Gamma \rhd a \rightarrow x \\ \bot, & \mathit{if}\ \mathrm{X}_a^\Gamma \not\rhd a \rightarrow x \end{cases}$$

Then $(\psi(\mathrm{X}_a^\Gamma))([\![x]\!]_a^\Gamma) = [\top]$ iff $\mathrm{X}_a^\Gamma \rhd a \rightarrow x$. The proof proceeds by induction on $\delta(x, \Gamma)$:

Base Cases: If x is a leaf of Γ, then either

1. $x = a$. Then $\mathrm{X}_a^\Gamma \rhd a \rightarrow a$. Also, $[\![a]\!]_a^\Gamma = [\top]$, so $(\psi(\mathrm{X}_a^\Gamma))([\![x]\!]_a^\Gamma) = (\psi(\mathrm{X}_a^\Gamma))([\top]) = [\top]$, or
2. $x \neq a$. Then $\Gamma \not\rhd x \rightarrow a$, and $[\![a]\!]_a^\Gamma = [\bot]$, so $(\psi(\mathrm{X}_a^\Gamma))([\![x]\!]_a^\Gamma) = (\psi(\mathrm{X}_a^\Gamma))([\bot]) = [\bot]$.

Induction Hypothesis: Assume that, for all nodes $x \in V_\Gamma$ with $\delta(x, \Gamma) < n$, $(\psi(\mathrm{X}_a^\Gamma))([\![x]\!]_a^\Gamma) = [\top]$ iff $\mathrm{X}_a^\Gamma \rhd a \rightarrow x$.

Induction Step: Let $y \in V_{\mathrm{X}_a^\Gamma}$, $\delta(y, \mathrm{X}_a^\Gamma) = n$. Then

$$\begin{aligned}(\psi(\mathrm{X}_a^\Gamma))([\![t]\!]_a^\Gamma) \;=\; & ((\psi(\mathrm{X}_a^\Gamma))([\![\mathcal{C}pos_\Gamma(t)]\!]_a^\Gamma) \wedge ({\sim}(\psi(\mathrm{X}_a^\Gamma))([\![\mathcal{C}neg_\Gamma(t)]\!]_a^\Gamma))) \\ & \vee((\psi(\mathrm{X}_a^\Gamma))([\![\mathcal{C}pos_\Gamma(t)]\!]_a^\Gamma) \wedge (\psi(\mathrm{X}_a^\Gamma))([\![\mathcal{C}neg_\Gamma(t)]\!]_a^\Gamma) \wedge (\psi(\mathrm{X}_a^\Gamma))(\hat{t}))\end{aligned}$$

Since $\forall w \in \mathcal{C}pos_\Gamma(x) \cup \mathcal{C}neg_\Gamma(x)$, $\delta(w, \Gamma) < n$, the induction hypothesis applies and $(\psi(\mathrm{X}_a^\Gamma))([\![w]\!]_a^\Gamma) = [\top]$ iff $\mathrm{X}_a^\Gamma \rhd a \rightarrow w$. Now, either

1. $\mathrm{X}_a^\Gamma \rhd a \rightarrow y$. Then y is positively a-reachable in X_a^Γ, but not negatively a-reachable: $\exists p \in \mathcal{C}pos_\Gamma(y)$, $\mathrm{X}_a^\Gamma \rhd a \rightarrow p$, but $\forall n \in \mathcal{C}neg_\Gamma(y)$, $\mathrm{X}_a^\Gamma \not\rhd a \rightarrow n$. Then $\exists p \in \mathcal{C}pos_\Gamma(y)$, $(\psi(\mathrm{X}_a^\Gamma))([\![p]\!]_a^\Gamma) = [\top]$, but $\forall n \in \mathcal{C}neg_\Gamma(y)$, $(\psi(\mathrm{X}_a^\Gamma))([\![n]\!]_a^\Gamma) = [\bot]$; and $(\psi(\mathrm{X}_a^\Gamma))([\![y]\!]_a^\Gamma) = [\top]$

2. $X_a^\Gamma \not\triangleright a \to y$. Then y is not positively a-reachable in X_a^Γ: either y is negatively a-reachable, or y is not a-reachable at all. In either case, $\forall p \in \mathcal{C}pos_\Gamma(y)$, $X_a^\Gamma \not\triangleright a \to p$, and so $(\psi(X_a^\Gamma))([\![p]\!]_a^\Gamma) = [\bot]$. But then $(\psi(X_a^\Gamma))([\![y]\!]_a^\Gamma) = [\bot]$ as well.

Corollary 5.1 *Let Γ be an inheritance hierarchy, with $a, x \in V_\Gamma$. Then $[\![x]\!]_a^\Gamma$ is satisfiable iff x is positively a-reachable in Γ. (resp., $[\![\overline{x}]\!]_a^\Gamma$ and negatively a-reachable).*

Proof: If $[\![x]\!]_a^\Gamma$ is satisfiable, then there is a truth-assignment $\mathcal{T}$ to the variables corresponding to the nodes of Γ such that $\mathcal{T}([\![x]\!]_a^\Gamma) = [\top]$. $\varphi(\mathcal{T})$ is a credulous extension of Γ w.r.t. a such that $\mathcal{T}([\![y]\!]_a^\Gamma) = [\top]$ iff $\varphi(\mathcal{T}) \triangleright a \to y$; so $\varphi(\mathcal{T}) \triangleright a \to x$. By theorem 1, this means that x is positively a-reachable in $\varphi(\mathcal{T})$, and since $\varphi(\mathcal{T})$ is a subhierarchy of Γ, x is positively a-reachable in Γ as well. Conversely, if x is positively a-reachable in Γ, then by lemma 1, there is an extension X_a^Γ of Γ w.r.t. a containing a positive path from a to x; by theorem 1, $X_a^\Gamma \triangleright a \to x$; and $\psi(X_a^\Gamma)$ is a valuation of the variables corresponding to the nodes of Γ that assigns $[\![x]\!]_a^\Gamma = [\top]$.

Corollary 5.2 *Let Γ be an inheritance hierarchy, with $a, x \in V_\Gamma$, and let Σ_a^Γ be the specificity extension of Γ w.r.t. a. Then $[\![x]\!]_a^{\Sigma_a^\Gamma}$ is satisfiable iff $\Gamma \triangleright a \to x$ (resp., $[\![\overline{x}]\!]_a^{\Sigma_a^\Gamma}$ and $a \not\to x$).*

Proof: By corollary 5.1, $[\![x]\!]_a^{\Sigma_a^\Gamma}$ is satisfiable iff x is positively a-reachable in Σ_a^Γ. By theorem 4, x is positively a-reachable in Σ_a^Γ iff $\Gamma \triangleright a \to x$.

Corollary 5.3 *Let Γ be an inheritance hierarchy, with $a, x \in V_\Gamma$, and let Σ_a^Γ be the specificity extension of Γ w.r.t. a. Then $[\![x]\!]_a^\Gamma$ is valid (tautological) iff $a \to x$ holds in every credulous extension of Γ, and $[\![x]\!]_a^{\Sigma_a^\Gamma}$ is valid iff $a \to x$ holds in every preferred credulous extension of Γ (resp., $[\![\overline{x}]\!]_a^\Gamma$ and negatively a-reachable).*

Proof: If $[\![x]\!]_a^\Gamma$ is valid, then every truth-assignment to the variables corresponding to the nodes of Γ satisfies $[\![x]\!]_a^\Gamma$. In particular, if X_a^Γ is a credulous extension of Γ w.r.t. a, then $\psi(X_a^\Gamma)$ satisfies $[\![x]\!]_a^\Gamma$—$\psi(X_a^\Gamma)([\![x]\!]_a^\Gamma) = [\top]$. But by theorem 5, $\psi(X_a^\Gamma)$ assigns $\psi(X_a^\Gamma)([\![x]\!]_a^\Gamma) = [\top]$ iff $X_a^\Gamma \triangleright a \to x$. Since the choice of X_a^Γ was arbitrary, this holds for *every* credulous extension of Γ w.r.t. a, and every credulous extension of Γ w.r.t. a supports $a \to x$.

Conversely, assume that $[\![x]\!]_a^\Gamma$ is not valid. Then there is some truth-assignment $\mathcal{T}$ that falsifies $[\![x]\!]_a^\Gamma$: $\mathcal{T}([\![x]\!]_a^\Gamma) = [\bot]$. Consider the credulous extension $\varphi(\mathcal{T})$; by theorem 5, $\varphi(\mathcal{T}) \not\triangleright a \to x$. So there is some credulous extension of Γ w.r.t. a that does not support $a \to x$.