First-Order Default Logic

Howard A. Blair
EECS Department
Syracuse University
Syracuse, NY 13244-4100 USA
blair@top.cis.syr.edu

Victor W. Marek
Computer Science Department
University of Kentucky
Lexington, KY 40506 USA
marek@cs.uky.edu

Topic: **Nonmonotonic Logics**

1 Introduction

We propose a model theory for full first-order default logic that allows both closed and non-closed default theories. Beginning with first-order languages without logical equality, we note how Henkin's proof of the completeness theorem for first-order logic yields complete algebras; that is, algebras over which models of consistent theories may always be found. The uniformity is what is interesting here. The algebra is constructed independently of the theory for which a model is sought and depends *only* on the underlying first-order language.. With these observations in place, the model theory for first-order defaults can be treated. Reiter [Rei80] has already told us what the extensions of *closed* first-order default theories are. With these extensions as a guide we introduce models, and extensions, of first-order default theories (D, W) where these theories may be closed or non-closed[1]. Beginning with closed default theories, the principal issue is how to check for consistency of the justifications in the defaults. A justification is consistent with a set of structures iff it is satisfied by *some* structure in the set. Let $\mathbf{\Gamma}$ be a set of structures. A $\mathbf{\Gamma}$-model of (D, W) is a *set* of structures over an algebra $\mathcal{A}$ which individually satisfy W and collectively satisfy D with respect to using $\mathbf{\Gamma}$ to check the consistency of the justifications. We describe these notions in detail in Section 5. The family of $\mathbf{\Gamma}$-models of (D, W) is closed under arbitrary union. Hence, over a given algebra $\mathcal{A}$ there is a unique largest $\mathbf{\Gamma}$-model. $\mathbf{\Gamma}$ is a model of (D, W) if $\mathbf{\Gamma}$ is the unique largest $\mathbf{\Gamma}$-model of (D, W). A maximal model of (D, W) is a stable model of (D, W) and the theory of a stable model, over a complete algebra, of (D, W) is an extension of (D, W). All complete algebras determine the same set of extensions. What about non-closed default theories? If one is going to assign values to freely occurring variables in default rules, we assume that one has a domain in mind where these values are to be found, i.e. an algebra $\mathcal{A}$. One may then close (D, W) with respect to $\mathcal{A}$, fundamentally by adjoining the theory of $\mathcal{A}$ to W, and instantiating the freely occurring variables in the defaults in D, and

[1] Lifschitz, [Li90], uses the term *open* to denote defaults allowing both quantifiers and freely occurring variables. We use the term "non-closed" here as the term "open" is standardly used to refer to quantifier-free formulas.

then taking the extensions of the resulting closed default theory. Here however, one obtains different sets of extensions with respect to different algebras, even when the algebras are complete, because of the interaction between the augmentation of W by the theory of $\mathcal{A}$ and the names of elements of $\mathcal{A}$ that instantiate the freely occurring variables in D. If one takes non-isomorphic complete algebras $\mathcal{A}$ and $\mathcal{A}'$, default theories (D, W) can be obtained whose closures with respect to $\mathcal{A}$ and $\mathcal{A}'$ produce differing sets of extensions. Instantiations that produce this difficulty do not occur in the case of closed default theories.

Lifschitz, in [Li90] introduces a technique for handling defaults admitting variables. The idea is to have a set F of constants describing the desired domain and then instantiate the variables with these constants.

There problems with Lifschitz approach. First, there are problems with the treatment of equality (Lifschitz is aware of this problem). Second, there are no restrictions on the structure of the set F and at times the resulting extensions are counterintuitive. These difficulties are illustrated with the following example.

Example 1.1: Let $W = \emptyset$, $D = \{\frac{:\neg p(x)}{p(a)}\}$. With the Herbrand universe as the set of the constants that are used for grounding, (D, W) has no extensions. Yet, if we add to W a completely immaterial fact, $r(b)$ (thus making the larger theory aware of b), the new default theory possesses an extension: $\mathrm{Cn}(\{p(a), r(b)\})$. But even if this case, if we add to W the sentence $a = b$, we are destroying this extension.

Our model theory captures the following intuition about default reasoning. We want to reason about the world while making cautious guesses about a situation based on our most-likely incomplete knowledge of the situation. The knowledge we have is codified by a theory W. The models of W are all of the ways the situation could be that are consistent with our knowledge. We also have knowledge and beliefs about how matters *normally* stand in situations, codified by default rules D. We find that the possibilities for our situation group together; seldom do we find that all models of W coherently accord with our default rules. Wanting to be cautious, that is, wanting to minimally constrain the possibilities for how the situation is but still take account of the default rules, we take maximal groupings of models of W among groupings that accord with the default rules.

In order to find which domains actually lead to results consistent with Reiter concept of extension we carefully investigate algebras for first-order language in which default theory is formulated. These algebras have, for each formula $\varphi(x)$ an element named by a constant "$\exists x \varphi(x)$" in the extension of the language $\mathbf{L}$ that assigns a *tentative example*, that is that the implication

$$\exists x \varphi(x) \rightarrow (\text{“}\exists x \varphi(x)\text{”})$$

holds. Notice the subtle difference from Skolemization – the constants "$\exists x \varphi(x)$" are added regardless whether the existential statement holds or not. The essence of the Henkin completeness proof is that the algebra constructed in that fashion is complete, that is every consistent theory in $\mathbf{L}$ has a model over that algebra (denoted below by $\mathbf{Hk_L}$). In fact the Henkin algebra is entirely effective as long as the language $\mathbf{L}$ possesses an effective presentation. A Henkin algebra which, of course, contains the Herbrand universe of the theory is uniquely determined by the language. Henkin's proof of the completeness theorem can then be interpreted as saying that the Henkin algebra has precisely the same properties with respect to the all the sentences of the language as its Herbrand universe plays with respect to the set of universal setences of $\mathbf{L}$.

2 Preliminaries

We assume that readers are familiar with first-order logic and understand what is meant by the terms *first-order language* and *structure*. Any of [Sh67], [Mo76] and [Kl67] for example are useful texts on mathematical logic where discussion of terminology and basic background results, not explicitly treated in this paper, may be found. We now clarify some of our basic terms and notation.

Let Σ be the set of function and predicate symbols of a first order language $\mathbf{L}$. σ is the *signature* of $\mathbf{L}$. Let Σ_1 and Σ_2 be the signatures of languages $\mathbf{L}_1$ and $\mathbf{L}_2$, respectively. Σ_1 is a *subsignature* of Σ_2 iff every function symbol of Σ_1 of arity n is a function symbol of Σ_2 of arity n and similarly for each predicate symbol. $\mathbf{L}_1$ is a *restriction* of $\mathbf{L}_2$ iff the signature of $\mathbf{L}_1$ is a subsignature of the signature of $\mathbf{L}_2$. Conversely, $\mathbf{L}_2$ is an *extension* of $\mathbf{L}_1$ iff $\mathbf{L}_1$ is a restriction of $\mathbf{L}_2$. We denote that $\mathbf{L}_1$ is a restriction of $\mathbf{L}_2$ by $\mathbf{L}_1 \preceq \mathbf{L}_2$. A *theory* T is a pair $(\mathbf{L}, \Gamma)$ where $\mathbf{L}$ is a first-order language and Γ is a set of formulas of $\mathbf{L}$. The *language of* T, denoted by $L(T)$, is $\mathbf{L}$ and the *nonlogical axioms* of T are the formulas in Γ. (Γ may be empty, and any nonlogical axiom may be a logical truth.) We denote Γ by $\mathbf{NLAx}(T)$. If T_1 and T_2 are theories such that $L(T_1) \preceq L(T_2)$ and every theorem of T_1 is a theorem of T_2, then T_1 is a *restriction* of T_2 and T_2 is an *extension* of T_1. Further, if every theorem of T_2 which is a formula of $L(T_1)$ is a theorem of T_1, then T_2 is a *conservative extension* of T_1.

Concerning the elements of a signature we adopt the following convention. If σ is a symbol in a signature Σ which is a function [resp. predicate] symbol of arity n, then σ is a function [resp. predicate] symbol of arity n in *every* signature. Alternatively, one may think of the type of a symbol as a part of the symbol, the actual character used to depict the symbol being a metalinguistic consideration. The convention is necessary for the following notion to be sensible.

Let $\mathbf{L}_1$ and $\mathbf{L}_2$ be languages with signatures Σ_1 and Σ_2, respectively. The language determined by the signature $\Sigma_1 \bigcup \Sigma_2$ is denoted by $\mathbf{L}_1 \bigcup \mathbf{L}_2$. If T_1 and T_2 are theories then we denote the theory $(\mathbf{L}_1 \bigcup \mathbf{L}_2,\ \mathbf{NLAx}(T_1) \bigcup \mathbf{NLAx}(T_2))$ by $T_1 \bigcup T_2$. The definition of the union of theories extends to infinite collections of theories in the obvious way.

The domain of a structure, often called the *universe* of a structure, the set of individuals, or, synonymously, elements, of the structure. Let $\mathbf{L}_1$, $\mathbf{L}_2$ be languages such that $\mathbf{L}_1 \preceq \mathbf{L}_2$, and let $\mathcal{U}$ be structure for $\mathbf{L}_2$. $\mathcal{U} \upharpoonright \mathbf{L}_1$ is the structure for $\mathbf{L}_1$ obtained from $\mathcal{U}$ be deleting the interpretations of the function and predicate symbols not in the signature of $\mathbf{L}_1$. (Note that the *domains* of $\mathcal{U}$ and $\mathcal{U} \upharpoonright \mathbf{L}_1$ are equal.) We say that $\mathcal{U} \upharpoonright \mathbf{L}_1$ is a *restriction* of $\mathcal{U}$ to $\mathbf{L}_1$ and $\mathcal{U}$ is *expansion* of $\mathcal{U} \upharpoonright \mathbf{L}_1$ to $\mathbf{L}_2$. If $\mathbf{\Gamma}$ is a class of structures for $\mathbf{L}_2$ then $\mathbf{\Gamma} \upharpoonright \mathbf{L}_1 = \{\mathcal{U} \upharpoonright \mathbf{L}_1\}$.

A *sentence* of the language $\mathbf{L}$ is a closed formula of $\mathbf{L}$; i.e. a formula without free occurrences of variables. A variable-free term is a called *ground* and a closed atomic formula is also called *ground*.

We will consider only languages without logical equality. This restriction is not a loss of generality because, when needed, we can restore equality nonlogically by adding a binary predicate symbol $\cong$ (which we write in infix position) to the language of a theory and extending the theory by including nonlogical axioms that say that $\cong$ is a congruence. Model-theoretically, we can then take the quotient of a given structure $\mathcal{S}$ that satisfies the congruence axioms for $\cong$ by the congruence relation that interprets $\cong$ in $\mathcal{S}$. The reason for this restriction will become apparent when we discuss Henkin's proof of the completeness theorem, below.

Definition 2.1: Let $\mathbf{L}$ be a first-order language (without equality). We denote by $\mathbf{L_f}$ the result of deleting *all* predicate symbols from $\mathbf{L}$ (including $\cong$, if present). An *algebra* for $\mathbf{L}$ is a structure for the language $\mathbf{L_f}$. A structure $\mathcal{S}$ for $\mathbf{L}$ is said to be *over* algebra $\mathcal{A}$ for $\mathbf{L}$ iff the restriction of $\mathcal{S}$ to $\mathbf{L_f}$ is $\mathcal{A}$. $\mathcal{A}$ is said to be the algebra *underlying* structure $\mathcal{S}$ iff $\mathcal{S}$ is over $\mathcal{A}$.

The principal applications of the theory of models over fixed algebra will be to default logic. Specifically we will be dealing with default theories, as introduced by Reiter in [Rei80].

Definition 2.2: Let $\mathbf{L}$ be a countable first-order language. A *default* over $\mathbf{L}$ is a syntactic object

$$\frac{\alpha\colon \beta_1, \ldots, \beta_m}{\gamma}$$

where the α, $\beta_j,\ 1 \leq j \leq m$ and γ are formulas of $\mathbf{L}$. A *default theory* is a pair (D, W) where D is a set of defaults over $\mathbf{L}$ and W is a theory whose language is $\mathbf{L}$. When all formulas occurring in defaults are sentences, we call (D, W) *closed*. Otherwise, (D, W) is called non-closed.

3 Henkin Algebras

The restriction of models of a consistent theory to considering only those models over a quotient of a fixed algebra is one of the main ideas in Henkin's proof of the strong classical completeness theorem for first-order logic. We shall exploit this idea to provide a single algebra over which we can construct a suitable notion of *extension* of full first-order default theories both with and without free occurrences of variables in defaults rule. We need to be able to ground-instantiate formulas over a fixed algebra $\mathcal{A}$ without having to worry about equality between ground terms, and we must be able to guarantee that whenever a set of formulas is satisfiable, the set is satisfiable over $\mathcal{A}$. Having banished logical equality (i.e. logical equality is a congruence relation that must be interpreted as the identity relation in any structure), we will be able to restrict ourselves to considering models over fixed algebras, without the need for the quotients to which we alluded above.

Let $\mathbf{L}$ be a first-order language, and let $\mathbf{L}_0 = \mathbf{L}$. Suppose languages $\mathbf{L}_0 \preceq \ldots \preceq \mathbf{L}_k$ have been obtained. To obtain $\mathbf{L}_{k+1}$ we introduce a new constant symbol e_φ for each sentence φ of $\mathbf{L}_k$ which is not a formula of $\mathbf{L}_i$, $i = 0, \ldots k-1$ and which has the form $\exists y\psi$. We call $k+1$ the *rank* of each constant symbol introduced to obtain $\mathbf{L}_{k+1}$. The constant symbols of $\mathbf{L}$ have *rank* 0. Let $\mathbf{L}_\omega$ be the extension of $\mathbf{L}$ whose function symbols consist of all constant symbols of finite rank, introduced above together with the function symbols of $\mathbf{L}$. Thus

$$\text{signature}(\mathbf{L}_\omega) = \bigcup_{k=0}^{\infty} \text{signature}(\mathbf{L}_k)\,.$$

Let $\mathcal{U}$ be the Herbrand universe of $\mathbf{L}_\omega$, and let $\mathcal{A}$ be the restriction of $\mathcal{U}$ to $\mathbf{L}$. We hereafter denote the constant symbol $\mathrm{e}_{\exists \mathrm{y}\psi}$ by "$\exists y\psi$". We refer to these constant symbols as *Henkin constants.* This is a notation we use in the metalanguage. Syntactically, within $\mathbf{L}_\omega$, the Henkin constant "$\exists y\psi$" has no internal structure.

Definition 3.1: We call $\mathcal{A}$ the *Henkin algebra* of $\mathbf{L}$, and denote it by HA($\mathbf{L}$). A *Henkin structure* for $\mathbf{L}$ is structure over $\mathcal{A}$.

Note that $\mathcal{A}$ is unique up to isomorphism by the choice of each new constant symbol of finite rank while forming $\mathbf{L}_\omega$. Also note that the domain of $\mathcal{A}$ contains infinitely many elements not named by any ground term in $\mathbf{L}$.

Proposition 3.1: There is a one-to-one correspondence between the subsets of the set of all ground atomic sentences of $\mathbf{L}_\omega$, and Henkin structures for $\mathbf{L}$. ∎

The preceding proposition says that Henkin structures are in one-to-one correspondence with Herbrand interpretations for $\mathbf{L}_\omega$.

4 The Rôle of Henkin's Proof of the Completeness Theorem

We recall one of the forms of the classical strong completeness theorem for first-order logic:

Completeness Theorem: A theory T is consistent iff T has a model.

Our purpose in recalling the Completeness Theorem is to exploit for somewhat different purposes a part of the construction given in Henkin's proof. We follow the presentation in [Sh67].

[Sh67] begins by first defining the *canonical structure* for a theory T. The interest in canonical structures, which we will define below, lies in the fact that under special circumstances, namely, for complete consistent Henkin theories, they are the models that we seek in proving the completeness theorem.

Informally, a Henkin theory is a theory in which every existential sentence $\exists x\, \varphi(x)$ has a 'witness'. Formally, T is a Henkin theory iff $(\exists x\, \varphi(x)) \rightarrow \varphi(t))$ is a theorem of T, for some variable-free term t depending on $\varphi(x)$, for each sentence $\exists x\, \varphi(x)$ in $L(T)$. Let us formally call such a ground term t a *Henkin witness* for $\exists x\, \varphi(x)$.

The main property of Henkin theories on which we focus is:

> The canonical structure of T is a model of T, if T is itself a complete Henkin theory.

This property is an immediate corollary of:

> Let T be a complete Henkin theory. For each sentence φ of $L(T)$, φ is true in the canonical structure for T iff φ is a theorem of T.

The canonical structure of an arbitrary theory T is defined as follows [Sh67]. For all variable-free terms t_1, t_2 of $L(T)$, put $t_1 \sim t_2$ iff $t_1 = t_2$ is a theorem of T. It follows that $\sim$ is a congruence relation. The individuals of the canonical structure are the congruence classes. Denote the congruence class of a variable-free term t by t°. The interpretation of an n-ary ($n \geq 0$) function symbol f of $L(T)$ is the function that maps an n-tuple $(t_1^\circ, \ldots, t_n^\circ)$ to the congruence class $f(t_1, \ldots, t_n)^\circ$. The n-ary relation which is the interpretation of an n-ary ($n \geq 1$) predicate symbol is the relation which holds of an n-tuple $(t_1^\circ, \ldots, t_n^\circ)$ iff $p(t_1, \ldots, t_n)$ is a theorem of T.

Notice that since logical equality does not occur in our languages the congruence classes induced by $\sim$ are simply singleton classes. Hence the canonical structure is an Herbrand interpretation, and the underlying algebra is of course the Herbrand universe of $\mathbf{L}$. If it happens that we are given a consistent complete Henkin theory with language $\mathbf{L}$, then the theory's canonical structure is in the unique Herbrand model of the theory.

Recall the construction of the language $\mathbf{L}_\omega$ in the construction of HA($\mathbf{L}$). Let T be a theory and let T_0 be T. Paralleling the construction of the sequence of languages $\{\mathbf{L}_i\}_{i=0}^\infty$ we construct a sequence of theories $\{T_i\}_{i=0}^\infty$ such that $L(T_i) = \mathbf{L}_i$. The nonlogical axioms $\mathbf{NLAx}(T_{k+1})$, $k \geq 0$, which hereafter we call *Henkin* axioms, are obtained by including the sentence

$$\exists y \psi(y) \rightarrow \psi(\text{“}\exists y \psi(y)\text{”})$$

for each sentence $\exists y\psi(y)$ of $\mathbf{L}_k$ such that "$\exists y\psi(y)$" has rank $k+1$. Let $T_H = \bigcup_{i=0}^{\infty} T_i$. Let $\mathbf{Hk_L}$ be the theory with language $\mathbf{L}$ and with nonlogical axioms consisting of the Henkin axioms introduced in the preceding construction. Notice that the Henkin axioms are independent of T. Therefore, $T_H = T \bigcup \mathbf{Hk_L}$. Hereafter we call $\mathbf{Hk_L}$ the *initial* Henkin theory with language $\mathbf{L}$.

Lemma 4.1: For each $k \geq 0$, T_{k+1} is a conservative extension of T_k. ∎

The preceding lemma is essential to the Henkin construction, and it is insightful to see how the Henkin axioms yield conservative extensions. The basic observation is that if e is a constant that does not occur in ψ or in $\varphi(x)$, and T is a theory whose nonlogical axioms do not contain e and $T \vdash (\exists x\varphi(x) \rightarrow \varphi(e)) \rightarrow \psi$ then $T \vdash \psi$ by prenex operations and the theorem on constants [Sh67].

Theorem 4.1: T_H is a conservative extension of T.

Corollary 4.1: If T is consistent then T_H is consistent.

Assume theory T is consistent. Henkin's proof of the Completeness Theorem is finished by applying some version of the axiom of choice, to obtain a maximal consistent extension T' of T_H with language $L(T_H)$. T' is then a complete theory, and hence a complete consistent Henkin theory. In case when the language $\mathbf{L}$ is denumerable, axiom of choice is not, actually, needed. The canonical structure of T', whose underlying algebra, when restricted to $L(T)$ is the Henkin algebra for $L(T)$, is a model of T', hence its restriction to $L(T)$ is a model of T. Thus, we can state the following more detailed form of the Completeness Theorem.

Theorem 4.2: A theory T is consistent iff T has a model over the Henkin algebra for $L(T)$.

Theorem 4.2 corresponds to the well-known completeness property of Herbrand algebra, that is a universal theory in $\mathbf{L}$ is consistent then it has a model over Herbrand algebra. Theorem 4.2 also shows that the concept of a *complete algebra* given by the next definition is not vacuous.

Definition 4.1: An algebra $\mathcal{A}$ for language $\mathbf{L}$ is *complete* iff every consistent theory with language $\mathbf{L}$ has a model over $\mathcal{A}$.

Is Henkin algebra the only complete algebra for $\mathbf{L}$? In fact it is easy to construct other complete algebras for $\mathbf{L}$.

Example 4.1: Consider the set of formulas $\mathbf{False} = \{\varphi : \forall x \neg\varphi(x) \text{ is a tautology}\}$. Thus **False** consists of formulas that are provably false. Let $\equiv$ be the congruence that "glues together" all elements with names "$\exists x\varphi(x)$" for $\varphi \in \mathbf{False}$, but nothing else. The algebra $\mathbf{HA}/\equiv$ is complete. The reason is that the constants "$\exists x\varphi(x)$" for $\varphi \in \mathbf{False}$ cannot be used as examples. It may be of interest that the algebra constructed here is no more effective, as the set **False** is Π^0_1-complete.

With the notion of a complete algebra in hand we can state an another equivalent form of the completeness theorem.

Theorem 4.3: Let $\mathcal{A}$ be a complete algebra for language $\mathbf{L}$ and let T be a theory with language $\mathbf{L}$. $T \vdash \varphi$ iff φ is valid in every model of T over $\mathcal{A}$.

Is one complete algebra as good as another? Among complete algebras, the Henkin algebra enjoys an important property given by the next theorem.

Theorem 4.4: Let T be a theory with language $\mathbf{L}$. Then there is a propositional theory (i.e. a theory whose nonlogical axioms are closed and quantifier-free) P with langauge $\mathbf{L}_\omega$ such that $T \bigcup \mathbf{Hk_L}$ is logically equivalent to $P \bigcup \mathbf{Hk_L}$.

Informally stated, the preceding theorem says that every theory is logically equivalent, modulo the initial Henkin theory of its language, to a propositional theory. Note both the similarity to Herband's theorem and the uniformity in theorem 4.4. The preceding theorem allows one to similarly reduce all default theories to propositional default theories if one so chooses.

Also, the notion of an algebra's not being complete allows us to introduce another idea, *domain-consistency*.

Definition 4.2: Let T be a theory, $\mathcal{A}$ an algebra for $L(T)$, not necessarily complete. T is *domain-consistent*, with respect to $\mathcal{A}$ iff T has a model over $\mathcal{A}$. We also say that T is $\mathcal{A}$-consistent if T is domain-consistent with respect to $\mathcal{A}$.

Domain consistency will be useful in the model theory of first-order default logic by allowing us to introduce a tighter notion of extension when the domain is fixed. Domain-consistency is stronger than consistency. Every theory which is domain-consistent is consistent. If the algebra is complete, then domain-consistency is just consistency.

We point out that while the Henkin algebra for $\mathbf{L}$ is a complete algebra for $\mathbf{L}$, it is not, as the Herbrand universe of $\mathbf{L}_\omega$ complete for $\mathbf{L}_\omega$. One needs only to give a language $\mathbf{L}'$ with infinitely many constants and a consistent theory T with language $\mathbf{L}'$ that has no Herbrand model in order to see this.

To treat the model theory of first-order default logic we need a few ideas and results from classical model theory. The reader interested in greater detail is referred to [Mo76]. One of the fundamental ideas underlying the concepts given by the next definition is to obtain classes of structures that stand in for theories.

Definition 4.3: The class of models of a theory T, called an *elementary class* is denoted by $\mathbf{Mod}(T)$. We denote the set of models of T over algebra $\mathcal{A}$ by $\mathbf{Mod}_{\mathcal{A}}(T)$. We call such a set of models an $\mathcal{A}$-elementary set. Let $\mathbf{\Gamma}$ be a class of structures for a language $\mathbf{L}$. The theory of $\mathbf{\Gamma}$, denoted by $\mathrm{Th}(\mathbf{\Gamma})$, is the theory whose nonlogical axioms consist of the set of sentences of $\mathbf{L}$ that are true in *every* structure in $\mathbf{\Gamma}$. If $\mathbf{\Gamma}$

is a singleton $\{\mathcal{S}\}$ then we write $\mathrm{Th}(\mathcal{S})$ for $\mathrm{Th}(\{\mathcal{S}\})$. Two structures $\mathcal{S}_1$ and $\mathcal{S}_2$ are *elementarily equivalent* iff $\mathrm{Th}(\mathcal{S}_1)$ is logically equivlent to $\mathrm{Th}(\mathcal{S}_2)$, i.e. they have the same theorems. A class of structures $\mathbf{\Gamma}$ for a language $\mathbf{L}$ is *compact* iff $\mathbf{\Gamma}$ satisfies: for every set S of sentences of $\mathbf{L}$ if every finite subset of S has a model in $\mathbf{\Gamma}$ then S has a model in $\mathbf{\Gamma}$. A class of structures $\mathbf{\Gamma}$ for $\mathbf{L}$ is elementarily closed iff for every structure $\mathcal{S}_1$ in $\mathbf{\Gamma}$, if $\mathcal{S}_2$ is elementarily equivalent to $\mathcal{S}_1$ then $\mathcal{S}_2$ is in $\mathbf{\Gamma}$. Similarly, A class of structures $\mathbf{\Gamma}$ ove $\mathcal{A}$ is elementarily closed over $\mathcal{A}$ iff for every structure $\mathcal{S}_1$ in $\mathbf{\Gamma}$ if $\mathcal{S}_2$ is elementarily equivalent to $\mathcal{S}_1$ and in $\mathcal{A}$ then $\mathcal{S}_2$ is in $\mathbf{\Gamma}$. Finally, the *compact elementary closure* of a class of structures $\mathbf{\Gamma}$ is $\mathrm{Mod}(\mathrm{Th}(\mathbf{\Gamma}))$. The *compact elementary $\mathcal{A}$-closure* of a set of structures $\mathbf{\Gamma}$ over algebra $\mathcal{A}$ is $\mathrm{Mod}_{\mathcal{A}}(\mathrm{Th}(\mathbf{\Gamma}))$

The following results show how elementary classes stand in for theories. *cf.* [Mo76].

Proposition 4.1:

A class of structures $\mathbf{\Gamma}$ for a language $\mathbf{L}$ is an elementary class iff $\mathbf{\Gamma}$ is elementarily closed and compact.

T is logically equivalent to $\mathrm{Th}(\mathbf{Mod}(T))$.

If $\mathbf{\Gamma}$ is elementarily closed and compact, then $\mathbf{\Gamma} = \mathrm{Mod}(\mathrm{Th}(\mathbf{\Gamma}))$.

Corollary 4.2: The smallest class of structures containing $\mathbf{\Gamma}$ which is compact and elementarily closed is $\mathrm{Mod}(\mathrm{Th}(\mathbf{\Gamma}))$.

There is a proposition and corollary for complete algebras that parallels the preceding proposition and its corollary.

Proposition 4.2: Let $\mathcal{A}$ be a complete algebra for a language $\mathbf{L}$. Let $\mathbf{\Gamma}$ be a class of structures over algebra $\mathcal{A}$.

$\mathbf{\Gamma}$ is an $\mathcal{A}$-elementary class iff $\mathbf{\Gamma}$ is elementarily closed over $\mathcal{A}$ and compact.

T is logically equivalent to $\mathrm{Th}(\mathbf{Mod}_{\mathcal{A}}(T))$.

If $\mathbf{\Gamma}$ is elementarily closed over $\mathcal{A}$ and compact, then $\mathbf{\Gamma} = \mathrm{Mod}_{\mathcal{A}}(\mathrm{Th}(\mathbf{\Gamma}))$.

Corollary 4.3: Let $\mathbf{\Gamma}$ be a class of structures over algebra $\mathcal{A}$. The smallest class of structures containing $\mathbf{\Gamma}$ which is compact and elementarily closed over $\mathcal{A}$ is $\mathrm{Mod}_{\mathcal{A}}(\mathrm{Th}(\mathbf{\Gamma}))$.

We need two more notions and a proposition that relates them to elementary classes in order to treat the model theory of first-order default logic.

Definition 4.4: Let $\mathcal{A}$ be an algebra for a language $\mathbf{L}$. Extend $\mathbf{L}$ to a language $\mathbf{L}(\mathcal{A})^{\cong}$ by including a name $\mathbf{i}_a$ for each element a in the universe of $\mathcal{A}$. Denote by $\hat{\mathcal{A}}$ the result of expanding $\mathcal{A}$ to be the algebra for the language $\mathbf{L}(\mathcal{A})$ in which each constant symbol $\mathbf{i}_a$ is interpreted as a and $\cong$ is interprested as the identity relation. The *theory of* $\mathcal{A}$, denoted $\mathrm{Th}(\mathcal{A})$ is the theory whose language is $\mathbf{L}(\mathcal{A})^{\cong}$ and whose nonlogical axioms are the congruence axioms for $\cong$ together with all sentences of $\mathbf{L}(\mathcal{A})^{\cong}_{\mathbf{f}}$ that are true in $\hat{\mathcal{A}}$.

Lemma 4.2:

$$\mathrm{Mod}_{\mathcal{A}}(T) = \mathrm{Mod}_{\hat{\mathcal{A}}}(T \bigcup \mathrm{Th}(\mathcal{A})) \upharpoonright \mathbf{L}\,.$$

5 Default Theories and their Model Theory

In this section we will apply the our results to study of default theories. Recall that a default theory [Rei80] is a pair (D, W) where D is a set of defaults, and W a set of formulas of $\mathbf{L}$. We will say that a default d is *closed* if all formulas occurring in d are sentences (that is have no free variables). Often one talks about *open* defaults. Those are defaults with formulas that may contain variables, but no quantifiers. We will also consider the most general type of defaults, that we will call *partially open*. Those will put no restrictions on the formulas appearing in d. Let us recall that Reiter, in [Rei80] assigns an extension to a closed default theories as follows:

Definition 5.1
1. Let (D, W) be a closed default theory. Given a set of sentences S, $\Gamma^{D,W}(S)$ is the least set U satisfying these conditions: (a) $W \subseteq U$, (b) U is closed under consequence, (c) Whenever $d = \frac{\alpha:\beta_1,\ldots,\beta_k}{\gamma}$ belongs to D, $\alpha \in U$, β_j is consistent with S for all $j \leq k$ then $\gamma \in S$.

2. A set of sentences Γ is an extension of (D, W) if $\Gamma^{D,W}(S) = S$.

Lifschitz [Li90] (see also [MT93]) shows how this notion can be defined semantically in the propositional case.

We will see how the notion of extension can be defined semantically. Our technique will involve constructing a set of extensions over a fixed algebra. The construction that we assign has the advantage to be independent of the algebra under consideration, as long as that algebra is complete over $\mathbf{L}$.

Thus let $\mathcal{A}$ be any algebra, and let us consider the set of all structures for $\mathbf{L}$ over $\mathcal{A}$. The elements of $|\mathcal{A}|$ are called *individuals*.

Now, fix a class Γ of structures and let $\mathcal{K}$ be a class of structures. A formula φ is *consistent with* $\mathbf{\Gamma}$ iff it is valid in some structure in $\mathbf{\Gamma}$. φ is $\mathcal{K}$-valid iff φ is valid in all structures in $\mathcal{K}$. Then, for closed defaults, a default $\frac{\alpha:\beta_1,\ldots,\beta_k}{\gamma}$ is $\mathcal{K}$-valid with respect to $\mathbf{\Gamma}$ iff α is $\mathcal{K}$-valid and each justification is consistent with $\mathbf{\Gamma}$ implies γ is $\mathcal{K}$-valid.

So, $\mathcal{K}$ is a $\mathbf{\Gamma}$-model iff all of the nonlogical axioms of W are $\mathcal{K}$-valid and all defaults in D are $\mathcal{K}$-valid with respect to $\mathbf{\Gamma}$.

Notice that here, a $\mathbf{\Gamma}$-model of (D, W) is not a single structure, but a family of structures.

We now have the following result:

Proposition 5.1: Let $\mathbf{\Gamma}$ be a fixed class of structures. Then the family of models of closed (D, W) is closed under arbitrary unions. Hence, over a given algebra $\mathcal{A}$ there is a unique largest $\mathbf{\Gamma}$-model.

With this proposition, we now define a *stable* model (over an algebra $\mathcal{A}$ of a default theory (D, W) as a class of structures $\mathbf{\Gamma}$ which is its own largest $\mathbf{\Gamma}$-model.

The following result tells us something about what the stable models look like and shows that they do stand in for theories without having to pad them with extra structures. The theorem is a consequence of proposition 4.1.

Proposition 5.2: Every stable model of a closed default theory is elementarily closed and compact. ∎

We say that T is an $\mathcal{A}$-extension of (D, W) if the set of all structures over $\mathcal{A}$ satisfying T is a $\mathcal{A}$-stable model of (D, W).

Could stable models of (D, W) nest? It turns out, that this is not the case.

Proposition 5.3: Let $\mathcal{A}$ be a complete algebra. Then $\mathcal{A}$-Stable models of (D, W) form an antichain, that is, if $\mathcal{K}_1$ and $\mathcal{K}_2$ are different $\mathcal{A}$-stable models of (D, W) then $\mathcal{K}_1$ is not included in $\mathcal{K}_2$.

It is natural to ask how the set of $\mathcal{A}$-extensions depends on $\mathcal{A}$. It turns out that for complete algebras we do not get much dependence.

Theorem 5.1 All complete algebras determine the same set of semantic extensions. That is, given two complete algebras $\mathcal{A}_1$, and $\mathcal{A}_2$, the sets of $\mathcal{A}_1$-extensions of a default theory (D, W) and $\mathcal{A}_2$-extensions of (D, W) coincide. Moreover, these extensions coincide with Reiter extensions.

Corollary 5.1: If (D, W) is closed and normal default theory, W is consistent and $\mathcal{A}$ is a complete algebra, then (D, W) possesses at least one stable model over $\mathcal{A}$ in particular a consistent extension (all such extensions coincide with Reiter extensions).

Clearly this corollary does not hold for incomplete algebras. Indeed, if W has no model over $\mathcal{A}$ then the empty class is the stable model and so we get as its theory the inconsistent theory.

There is a version of theorem 4.4 for defualt logic which has essentially the same proof.

Theorem 5.2: Let (D, W) be a closed default theory with language $\mathbf{L}$. There is a set of propositional default rules D' and a propositional theory W' with language $\mathbf{L}_\omega$ such that $(D, W \bigcup \mathbf{Hk_L}))$ has precisely the same stable models over $\widehat{\mathrm{HA}(\mathbf{L})}$ as $(D', W' \bigcup \mathbf{Hk_L})$. ∎

We will investigate the case of partially open defaults, that is defaults of the form $d = \frac{\alpha(x):\beta_1(x),\ldots,\beta_k(x)}{\gamma(x)}$, where $\alpha, \beta_1, \ldots, \beta_k, \gamma$ are formulas of $\mathbf{L}$ and x is a set of variables. These formulas allow quantifiers.

First, we need to understand what happens when we assign values to variables occurring free in our formulas. This means that we have a domain consisting of values that we are using. Specifically, we must have an algebra $\mathcal{A}$ whose elements we are using. The properties of $\mathcal{A}$ are given **a priori** (as $\mathrm{Th}(\mathcal{A})$) and therefore must be included in the initial conditions. This means that we have to add $\mathrm{Th}(\mathcal{A})$ to W. Once this is done, we are now able to transform a default theory (D, W) to a *closed* default theory (but over a given algebra $\mathcal{A}$) as follows.

Definition 5.2: Let (D, W) be a default theory (possibly partially open), and let $\mathcal{A}$ be an algebra. A default theory $(D, W)^{\mathcal{A}}$ is a closed default theory (formulated in $\mathbf{L}_{\mathcal{A}}$), (D_1, W_1) where:
(1) $W_1 = W \cup \mathrm{Th}(\mathcal{A})$
(2) D_1 consists of all instantiation of defaults in D to ground terms of $\mathcal{A}$.

Since the default theory $(D, W)^{\mathcal{A}}$ is closed, we know how to construct extensions of $(D, W)^{\mathcal{A}}$. Those will be called extensions of (D, W) *over* $\mathcal{A}$. Notice that our construction is similar to that of Lifschitz [Li90] but differs from it in that we add theory of $\mathcal{A}$ to W. This is motivated by the fact we need to incorporate the facts true in $\dashv$ in the initial data.

The extensions of $(D, W)^{\mathcal{A}}$ are sets of sentences of $\mathbf{L}_{\mathcal{A}}$ rather than $\mathbf{L}$. Thus it is natural to ask for the correctness of our construction with respect to Reiter extensions. That is, we need to investigate what happens when we restrict to the original language $\mathbf{L}$ and closed theories. Lemma 4.2 implies the following crucial property:

Theorem 5.3: Let $\mathcal{A}$ be a complete algebra for $\mathbf{L}$ and let (D, W) be a closed default theory in $\mathbf{L}$. Then the restrictions (to $\mathbf{L}$) of extensions of (D, W) over $\mathcal{A}$ are precisely Reiter extensions of (D, W).

References

[GL88] Gelfond, M. and Lifschitz V., The stable semantics for logic programs. In R. Kowalski and K. Bowen, editors, *Proceedings of the 5th International Symposium on Logic Programming*, pages 1070–1080, Cambridge, MA, 1988. MIT Press.

[Kl67] Kleene, S. C., *Mathematical Logic*. Wiley, 1967.

[Li90] Lifschitz, V., On Open Defaults, In J. Lloyd, editor, *Proceedings of the Symposium on Computational Logic,* Berlin: Springer-Verlag, pages 80–95, 1990.

[MT93] Marek, W. and Truszczyński, M., *Nonmonotonic Logics; Context-Dependent Reasoning*. Berlin: Springer-Verlag, 1993.

[Mo76] Monk, J.D., *Mathematical Logic*. Springer-Verlag, 1976.

[Rei80] Reiter, R.. A logic for Default Reasoning. *Artificial Intelligence*, 13:81–132, 1980.

[Sh67] Shoenfield, J.R., *Mathematical Logic*. Addison-Wesley, 1967.