Complete Axiomatisations of Properties of Finite Sets

THOMAS ÅGOTNES, *Bergen University College, Norway.*
E-mail: tag@hib.no

MICHAL WALICKI, *University of Bergen, Norway.*
E-mail: michal@ii.uib.no

Abstract

We study a logic whose formulae are interpreted as properties of a finite set over some universe. The language is propositional, with two unary operators *inclusion* and *extension*, both taking a finite set as argument. We present a basic Hilbert-style axiomatisation, and study its completeness. The main results are syntactic and semantic characterisations of complete extensions of the logic.

1 Introduction

Finite sets and, more generally, finite structures, play a vital role in many areas and applications (finite model theory, e.g., [10]; bounded arithmetics, e.g., [8, 9]; database theory, e.g., [15, 18]; reasoning about agents with bounded memory, e.g., [5, 3, 2]; reasoning about coalitions of agents, e.g., [4]). Unfortunately, many desirable properties which are obtained in the situation when all models are admitted, cease to be valid when the model class contains only finite structures. A paradigmatic example is Trahtenbrot's theorem according to which there is no sound and complete system for deducing all first-order properties valid in all finite structures. In this article we address the problem of complete reasoning about properties of finite sets in a propositional language with unary inclusion and extension operators taking (symbols representing) finite sets as arguments. Considering this restricted logical vocabulary we obtain a reasoning system and a characterisation of additional axioms for which the system is complete with respect to the corresponding class of finite sets. In particular, the empty theory satisfies the criteria, and so we obtain weak completeness for the class of all finite sets: the formulae of our language valid in all finite sets are recursively enumerable. Thus, the paper gives an example of a language and finite semantics for which (a counterpart of) Trahtenbrot's theorem does *not* apply.

One important proviso is the following. One can discuss possible axiomatisations of the finite sets among all sets. For instance, starting with the axioms of ZF and removing the axiom of infinity, one can ask what other adjustments are needed in order to obtain all and only finite sets. ([7, 23, 17] are examples of works in this direction. Reverse mathematics, originating with [12], gives numerous variations on this theme, all with the additional axioms for the number system and arithmetic operations.) Our purpose in this paper is very different. We do not aim at an axiomatisation of finite sets but at complete axiomatisations of their properties. We assume that such sets are given and use them as the semantic framework in which our logical theories are interpreted. Theories (in our restricted language)

Vol. 16 No. 3,
doi:10.1093/jigpal/jzn008

specify only lower and upper bounds on the contents of a set – their models are all finite sets satisfying the specification. It quickly turns out that there can be no finitary logical system which is strongly complete with respect to this semantics. Thus, the main question we study in this article is: for which theories is the reasoning system complete?

The general setting can be described as follows. We start with a, typically infinite, *universe* U, and let s be some finite subset of U. In a language with explicit symbols for finite sets, we could then express properties of s by relating it to other finite sets s', such as *inclusion* ($s \subseteq s'$) and *extension* ($s \supseteq s'$). If we wanted to express properties of an otherwise *unknown* finite set, we could introduce a formal variable x to allow statements such as $x \subseteq s'$ and $x \supseteq s'$. In this article we take another approach in order to simplify the presentation and focus on the essential issue of finitude. Instead of using formal variables, we express properties of *a single, given,* finite set s. We use unary versions of inclusion and extension, taking finite sets as arguments, to express properties of the set s, such as $\sqsubseteq(s')$ (s is included in s') and $\sqsupseteq(s')$ (s is an extension of s'). We use a propositional language with such expressions for each finite subset s' of U as atoms. It is interpreted as statements about finite sets s of U in the obvious way.

The language with the unary set operators was originally introduced in [1], from where most of the results presented here are adopted. Similar languages have been applied in several contexts. In [4], the language was used to express properties of groups of agents (*coalitions*) such as "having a as a member and not having any other agents than a, b or c as members". Incorporation of the language into a variant of *Coalition Logic* [24] provided a restricted form of quantification allowing succinct expressions of coalitional abilities. In [1, 5, 3, 2] the language was used for expressing properties of the beliefs of resource-bounded agents, such as "believing p and not believing anything else than p, q or r". In this article, some proof– and model theoretic properties of the language are presented independently from possible applications. The results are of primary interest when the universe U is infinite, such as in the second but not the first of the mentioned applications.

The paper is organised as follows. In the next section, the language and semantics are defined. In Section 3 we present a sound Hilbert-style logical system. Since the logic is not semantically compact, a strongly complete finitary axiomatisation is not possible. We set out instead, not only to show weak completeness, but also an intermediate notion of completeness between weak and strong: we characterise the sets of premises for which the logic is complete, i.e., the theories Γ for which $\Gamma \models \phi$ implies that $\Gamma \vdash \phi$ for any formula ϕ. Such theories are called *finitary*. As an intermediate step, Section 4 presents results for an alternative semantics in the form of general models, allowing also infinite sets. Section 5 provides a complete proof theoretic characterisation of finitary theories, while Section 6 provides several sufficient conditions in terms of algebraic properties of the model classes. The conditions imply, in particular, that the empty (and every finite) theory is finitary and thus that the system is weakly complete. Section 7 gives some examples illustrating application of the results in the context of beliefs of resource-bounded agents. Section 8 concludes.

2 Language and semantics

Assume a fixed (typically infinite) set U, henceforth called *the universe*. The set of all finite subsets of U is denoted $\wp^{fin}(U)$.

Formulae ϕ of the language $\mathcal{L}(U)$, or just $\mathcal{L}$ when U is clear from context, are defined over an alphabet including (names for) the elements of $\wp^{fin}(U)$ by the following grammar, where $X \in \wp^{fin}(U)$:

$$\phi ::= \sqsubseteq(X) \mid \sqsupseteq(X) \mid \neg\phi \mid \phi \wedge \phi \tag{1}$$

We use the usual derived propositional connectives, in addition to $\underline{=}(X)$ for $\sqsubseteq(X) \wedge \sqsupseteq(X)$. When $X = \{e\}$ is a singleton, we sometimes write $\underline{\in} e$ for $\sqsubseteq(X)$ and $\underline{\ni} e$ for $\sqsupseteq(X)$ for simplicity. Also, we write $\not\sqsubseteq(X)$, $\not\sqsupseteq(X)$, $\underline{\notin} e$ and $\underline{\not\ni} e$, for $\neg\sqsubseteq(X)$, $\neg\sqsupseteq(X)$, $\neg\underline{\in} e$ and $\neg\underline{\ni} e$, respectively. A *theory* is a set of formulae $\Gamma \subseteq \mathcal{L}$.

$\mathcal{L}$ is interpreted by a single finite subset of the universe, $s \in \wp^{fin}(U)$, called a *model set* or just a *model*. Formally, the satisfaction relation $\models \subseteq \wp^{fin}(U) \times \mathcal{L}$ is defined as follows:

$$\begin{array}{lcl} s \models \sqsubseteq(X) & \Leftrightarrow & s \subseteq X \\ s \models \sqsupseteq(X) & \Leftrightarrow & X \subseteq s \\ s \models \neg\psi & \Leftrightarrow & s \not\models \phi \\ s \models \psi_1 \wedge \psi_2 & \Leftrightarrow & s \models \psi_1 \text{ and } s \models \psi_2 \end{array} \tag{2}$$

Note that finite sets are used as symbols in the logical language, and that the interpretation of a set-symbol is the set itself.

$\mathcal{S}^{fin}$ denotes the class $\wp^{fin}(U)$ of all finite sets. We use the usual terminology: s satisfies a set of formulae iff it satisfies each formula in the set; a formula ϕ is *satisfiable* (in $\mathcal{S}^{fin}$) iff there is an $s \in \mathcal{S}^{fin}$ such that $s \models \phi$; it is *valid*, $\models \phi$, iff it is satisfied by every $s \in \mathcal{S}^{fin}$; it is a *logical consequence* of a set $\Gamma \subseteq \mathcal{L}$, $\Gamma \models \phi$, iff s satisfies ϕ whenever s satisfies Γ. The class of all finite sets satisfying a set of formulae Γ is denoted $\mathcal{S}^{fin}(\Gamma)$.

It is easy to see that the language is expressively complete, in the sense that for every finite set s there is a formula ϕ^s such that $s' \models \phi^s$ iff $s' = s$ (take $\phi^s = \sqsubseteq(s) \wedge \sqsupseteq(s)$).

Before we discuss the question of axiomatisation in the next section, observe that the logic is not semantically compact.

Example 2.1

The semantics of $\mathcal{L}$ is not compact. A counter example can be constructed from any infinite set. For instance:

$$\Gamma_1 = \{\underline{\ni} e : e \in U\}$$

Although every finite subset of Γ_1 is satisfiable, Γ_1 itself is not when U is infinite. A more specific example is given by the theory Γ_2 requiring membership of all Peano numerals, i.e., $\underline{\ni} 0$ and, whenever $\underline{\ni} n$ then also $\underline{\ni} s(n)$.

The analysis in the following sections is mostly interesting in the case that the universe is infinite.

3 An axiomatisation

Definition 3.1

FS is the logical system over the language $\mathcal{L}$ consisting of the following axiom schemata:

All substitution instances of tautologies of propositional calculus		Prop
$\sqsupseteq(\varnothing)$		E1
$(\sqsupseteq(X)\wedge\sqsupseteq(Y))\rightarrow\sqsupseteq(X\cup Y)$		E2
$\neg(\sqsupseteq(X)\wedge\sqsubseteq(Y))$	when $X\not\subseteq Y$	E3
$(\sqsubseteq(Y\cup\{\gamma\})\wedge\not\sqsupseteq\gamma)\rightarrow\sqsubseteq(Y)$		E4
$\sqsupseteq(X)\rightarrow\sqsupseteq(Y)$	when $Y\subseteq X$	S
$\sqsubseteq(X)\rightarrow\sqsubseteq(Y)$	when $X\subseteq Y$	G

and the following deduction rule

$$\frac{\Gamma\vdash\phi,\Gamma\vdash\phi\rightarrow\psi}{\Gamma\vdash\psi} \qquad \text{MP}$$

$\Gamma\vdash\phi$ means that ϕ is derivable from Γ in *FS*. $\vdash\phi$ denotes that $\varnothing\vdash\phi$. *FS* is *sound* (wrt. $\mathcal{S}^{fin}$) iff $\Gamma\vdash\phi$ implies that $\Gamma\models\phi$, for any theory Γ and formula ϕ. A set of formulae Γ is *consistent* if it is not the case that both $\Gamma\vdash\phi$ and $\Gamma\vdash\neg\phi$ for some formula ϕ; Γ is *maximal* if either $\phi\in\Gamma$ or $\neg\phi\in\Gamma$ for any $\phi\in\mathcal{L}$. Γ is *maximal consistent* if it is both maximal and consistent.

Note that the system *FS* cannot be *strongly* complete, in the sense that $\Gamma\vdash\phi$ whenever $\Gamma\models\phi$ for every Γ and ϕ, because the semantics is not compact. For a given Γ, by Γ-*completeness* (wrt. $\mathcal{S}^{fin}$), we mean the property that $\Gamma\vdash\phi$ whenever $\Gamma\models\phi$ for any ϕ. Note that in this terminology, *weak* completeness, i.e., the property that $\models\phi$ implies that $\vdash\phi$ for all ϕ, is the same as $\varnothing$-completeness. Γ-completeness is useful if we want to axiomatise a subclass of all finite sets: if *FS* is Γ-complete then the system obtained by adding Γ as axioms to *FS* is weakly complete with respect to $\mathcal{S}^{fin}(\Gamma)$. While *FS* is not Γ-complete for all Γ (i.e., not strongly complete), we now set out to investigate for which theories Γ *FS* *is* Γ-complete. Of particular interest, of course, is weak completeness. While it is not too difficult to prove weak completeness of *FS* directly, we shall prove a more general result from which weak completeness, as well as other forms of Γ-completeness, follow as special cases.

We first relax the semantic assumption about finiteness, in order to establish some intermediary results.

4 A detour to infinite sets

Consider a variant of the semantics where formulae are interpreted by *arbitrary* subsets of the universe, and not only by finite subsets. The satisfaction relation is defined exactly as in Section 2 for finite sets. Now, the theory Γ_1 from Example 2.1 *is* satisfiable. However, the semantics is still not compact.

Example 4.1
Consider an infinite U and let $e \in U$ be arbitrary. The following theory states that $e \in s \neq \{e\}$ while for all $e' \neq e$, $e' \notin s$:

$$\Gamma_3 = \{\underline{\ni} e, \underline{\not\in} e\} \cup \{\underline{\not\ni} e' : e' \in U, e' \neq e\}$$

Every finite subset of Γ_3 is satisfiable (in the infinite sets semantics), but Γ_3 is not.

The problem with respect to compactness with the above Γ_3 is that, while the first part requires some element $e' \neq e$ to be included in the model, the second part forbids any other element to be included.

However, there is a slightly different compact variant of the semantics. Let $* \notin U$ be a new element which is not in the universe, and let

$$\mathcal{S}^* = \wp(U) \cup \wp^{fin}(U \cup \{*\})$$

The sets $\mathcal{S}^*$ – finite subsets of the universe, infinite subsets of the universe, or finite subsets of the universe with the element $*$ added – are henceforth called *general models*. We interpret the language in general models by defining a satisfaction relation $\models_* \subseteq \mathcal{S}^* \times \mathcal{L}$ in exactly the same way, i.e., using (2), as for the satisfaction relation $\models$ between finite sets and $\mathcal{L}$. It is important to note that while we extend the set of models, we do not change the language $\mathcal{L}$: sets X occurring in formulae must still be finite subsets of the universe – in particular, they cannot contain the element $*$.

We now establish some results for this semantics. These will be useful for studying the finite sets semantics in the following sections.

In the following, we will use $\models_*$ to denote satisfiability and logical consequence with respect to $\mathcal{S}^*$, while we keep $\models$ for the case of only finite sets. A formula ϕ is satisfiable in $\mathcal{S}^*$ iff there is an $s \in \mathcal{S}^*$ such that $s \models_* \phi$; ϕ is valid wrt. $\mathcal{S}^*$ ($\models_* \phi$) iff every $s \in \mathcal{S}^*$ satisfies ϕ; ϕ is a logical consequence of Γ wrt. $\mathcal{S}^*$ ($\Gamma \models_* \phi$) iff $s \models \Gamma$ implies that $s \models \phi$ for every $s \in \mathcal{S}^*$; *FS* is sound wrt. $\mathcal{S}^*$ iff $\Gamma \vdash \phi$ implies that $\Gamma \models_* \phi$ for any theory Γ and formula ϕ; *FS* is strongly complete wrt. $\mathcal{S}^*$ iff $\Gamma \models_* \phi$ implies that $\Gamma \vdash \phi$ for any theory Γ and formula ϕ. Given a theory Γ, we write $\mathcal{S}^*(\Gamma)$ for the class of its general models.

Theorem 4.2
FS is sound and strongly complete with respect to $\mathcal{S}^*$.

In order to prove Theorem 4.2, we first establish some lemmas. The proof of the following lemma is standard.

Lemma 4.3
For any consistent set of formulae Γ, there is a maximal consistent set Γ' such that $\Gamma \subseteq \Gamma'$.

Lemma 4.4
Let Γ be a maximal consistent theory. If there is an X' such that $\sqsubseteq(X') \in \Gamma$, then for every $X \in \wp^{fin}(U)$,

$$\sqsubseteq(X) \in \Gamma \Leftrightarrow \{e : \underline{\ni} e \in \Gamma\} \subseteq X$$

Proof. Assume that $\sqsubseteq(X') \in \Gamma$. For the direction to the right, let $\sqsubseteq(X) \in \Gamma$, and let $\underline{\ni} e \in \Gamma$. If $e \notin X$, then $\neg \sqsubseteq(X) \in \Gamma$ by axiom E3, which contradicts the consistency of Γ. For the direction to the left, let $\{e : \sqsupseteq e \in \Gamma\} \subseteq X$. If $X' \subseteq X$ then $\sqsubseteq(X) \in \Gamma$ by axiom G and we are done. Assume that $X' \not\subseteq X$, i.e., that $X' \setminus X = \{e_1, \ldots, e_k\}$ for some $k \geq 1$. Because $e_j \notin X$ for each j, $e_j \notin \{e : \sqsupseteq e \in \Gamma\}$. Thus, for each j, $\sqsupseteq e_j \notin \Gamma$. By maximality, $\neg \sqsupseteq e_j \in \Gamma$. By axiom E4, $\sqsubseteq(X' \setminus \{e_1, \ldots, e_k\}) \in \Gamma$, and because $X' \setminus \{e_1, \ldots, e_k\} \subseteq X$, $\sqsubseteq(X) \in \Gamma$ by axiom G. ■

Lemma 4.5
Every maximal consistent theory is satisfiable in $\mathcal{S}^*$.

Proof. Given a maximal consistent theory Γ, define the following general model $s \in \mathcal{S}^*$:

$$s = \begin{cases} \{e : \sqsupseteq e \in \Gamma\} & \text{if } \sqsubseteq(X') \in \Gamma \text{ for some } X', \text{ or } \{e : \sqsupseteq e \in \Gamma\} \text{ is infinite} \\ \{e : \sqsupseteq e \in \Gamma\} \cup \{*\} & \text{otherwise} \end{cases}$$

We show that

$$s \models_* \phi \Leftrightarrow \phi \in \Gamma$$

for any formula ϕ, by structural induction. When $\phi = \sqsupseteq(X)$ and $X = \varnothing$ we are done by E1. When $\phi = \sqsupseteq(X)$ and $X \neq \varnothing$, $s \models_* \phi$ iff $X \subseteq s$ iff, since $* \notin X$, $X \subseteq \{e : \sqsupseteq e \in \Gamma\}$ iff, by E2 in one direction and S in the other, $\phi \in \Gamma$.

Let $\phi = \sqsubseteq(X)$. First assume that there is a X' such that $\sqsubseteq(X') \in \Gamma$. $s = \{e : \sqsupseteq e \in \Gamma\}$, and by Lemma 4.4, $\phi \in \Gamma$ iff $s \subseteq X$, which holds iff $s \models_* \phi$. Second, assume that $\sqsubseteq(X') \notin \Gamma$ for all X'. In particular, $\phi \notin \Gamma$. But then we also have that $s \not\models_* \phi$, since either $s = \{e : \sqsupseteq e \in \Gamma\} \cup \{*\}$ and $* \notin X$, or s is infinite (and X is finite). Induction passes trivially through the propositional connectives. ■

Proof of theorem 4.2. Prop, E1–G are valid, and MP preserves logical consequence, giving soundness. Strong completeness follows directly from Lemmas 4.3 and 4.5. ■

5 Finitary theories and completeness

We now consider soundness and completeness of *FS* with respect to $\mathcal{S}^{fin}$. The following follows immediately from Theorem 4.2.

Corollary 5.1
FS is sound wrt. $\mathcal{S}^{fin}$.

As discussed in Section 3, we want to characterise the theories Γ for which *FS* is Γ-complete, i.e., for which $\Gamma \models \phi$ implies that $\Gamma \vdash \phi$ for any ϕ. In this section we provide such a characterisation. We define the concept of a *finitary theory*, and show that the finitary theories are exactly the theories for which *FS* is Γ-complete. The proof builds upon the completeness result for the more general semantics described in the previous section.

Definition 5.2 (Finitary and Finitarily Open Theories)
A theory Γ is *finitary* iff it is consistent and for any ϕ:

$$\begin{array}{c} \Gamma \vdash \sqsubseteq(X) \to \phi \text{ for all } X \in \wp^{fin}(U) \\ \Downarrow \\ \Gamma \vdash \phi \end{array}$$

A theory Γ is *finitarily open* iff there exists a finite set X such that $\Gamma \nvdash \not\sqsubseteq(X)$.

Informally speaking, a theory is finitary if provability of a formula under arbitrary upper bounds on the model set implies provability of the formula itself.

As an illustration, the following are examples of non-finitary theories; the claims will be proved after the following lemmas.

Example 5.3

- *Not finitarily open and not finitary:* The theories Γ_1 and Γ_3 in examples 2.1 and 4.1, respectively
- *Finitarily open, but not finitary:* $\Gamma_4 = \{\not\supseteq p_i \to \supseteq a : i \in N\}$ (where $a \neq p_i$ for all i)

Lemma 5.4

1. A finitary theory is finitarily open
2. If Γ is a finitary theory and $\Gamma \nvdash \phi$, then $\Gamma \cup \{\neg\phi\}$ is finitarily open

Proof.

1. If Γ is finitary and not finitarily open, $\Gamma \vdash \not\sqsubseteq(X)$ for all X. Thus, $\Gamma \vdash \sqsubseteq(X) \to \phi$ for any ϕ, and by finitariness $\Gamma \vdash \phi$. But by the same argument $\Gamma \vdash \neg\phi$, contradicting consistency of Γ.
2. Let Γ be a finitary theory such that $\Gamma \nvdash \phi$. There must be an X such that $\Gamma \nvdash \sqsubseteq(X) \to \phi$. By Prop, $\Gamma \nvdash \neg\phi \to \not\sqsubseteq(X)$, and thus $\Gamma \cup \{\neg\phi\} \nvdash \not\sqsubseteq(X)$ which shows that $\Gamma \cup \{\neg\phi\}$ is finitarily open.

■

Lemma 5.5

A theory Γ is finitarily open iff it is satisfiable in $\mathcal{S}^{fin}$.

Proof. Γ is finitarily open iff there exists an $X \in \wp^{fin}(U)$ such that $\Gamma \nvdash \not\sqsubseteq(X)$; iff, by Theorem 4.2, there exists an X such that $\Gamma \not\models_* \not\sqsubseteq(X)$ iff there exists an X and a model set $s \in \mathcal{S}^*$ such that $s \models_* \Gamma$ and $s \models \sqsubseteq(X)$ iff there is a finite set $s \in \wp^{fin}(U)$ such that $s \models \Gamma$. ■

Example 5.6 (Example 5.3 continued)

- Γ_1 and Γ_3 are not satisfiable by a finite set, and thus not finitarily open or finitary.
- $\Gamma_4 = \{\not\supseteq p_i \to \supseteq a : i \in N\}$ ($a \neq p_i$ for all i) has a finite model, for instance the set $\{a\}$. However, the theory $\Gamma_4 \cup \{\not\supseteq a\}$ has no finite model and, by the above lemma, is not finitarily open. $\Gamma_4 \nvdash \supseteq a$ because otherwise there would be a finite set $\Delta \subseteq \Gamma_4$ such that $\Delta \vdash \supseteq a$, and by soundness (Corollary 5.1) $\Delta \models \supseteq a$. To see that the latter is not true, take $s = \{p_i : \not\supseteq p_i \to \supseteq a \in \Delta\}$ (a finite set); we have that $s \models \Delta$ but $s \not\models \supseteq a$. Hence, by lemma 5.4, Γ_4 is not finitary. The formula $\supseteq a$ also provides a direct witness to nonfinitariness of Γ_4: for every finite set X, we do have $\Gamma_4 \vdash \sqsubseteq(X) \to \supseteq a$, while $\Gamma_4 \nvdash \supseteq a$ as just observed.

It is, however, more difficult to prove *finitariness* of a theory. Indeed, much of the following is concerned with that problem. It is particularly interesting because of the following theorem.

Theorem 5.7

FS is Γ-complete wrt. $\mathcal{S}^{fin}$ iff Γ is finitary.

Proof. Let Γ be a finitary theory and let $\Gamma \models \phi$. By Lemma 5.4.1 Γ is finitarily open and thus satisfiable by Lemma 5.5. $\Gamma \cup \{\neg\phi\}$ is unsatisfiable in $\wp^{fin}(U)$, and thus not finitarily open, and it follows from Lemma 5.4.2 that $\Gamma \vdash \phi$.

For the other direction, let *FS* be Γ-complete wrt. $\mathcal{S}^{fin}$. Assume that $\Gamma \nvdash \phi$. By Γ-completeness, $\Gamma \not\models \phi$, that is, there is an $s \in \wp^{fin}(U)$ such that $s \models \Gamma$ and $s \not\models \phi$. We have that $s \models \sqsubseteq(s)$, and so $s \not\models \sqsubseteq(s) \rightarrow \phi$. By soundness (Corollary 5.1) it follows that $\Gamma \nvdash \sqsubseteq(s) \rightarrow \phi$, showing that Γ is finitary. ■

The following corollary sums up characterisations of finitary theories. Note point 5 which is a special form of a finite model property.

Corollary 5.8
Let Γ be a theory. The following statements are equivalent:

1. Γ is finitary
2. $\Gamma \nvdash \phi \Rightarrow \Gamma \cup \{\neg\phi\}$ is finitarily open, for any ϕ
3. $\Gamma \models \phi \Rightarrow \Gamma \vdash \phi$, for any ϕ
4. $\Gamma \models \phi \Rightarrow \Gamma \models_* \phi$, for any ϕ
5. $\mathcal{S}^*(\{\Gamma,\phi\}) \neq \varnothing \Rightarrow \mathcal{S}^{fin}(\{\Gamma,\phi\}) \neq \varnothing$, for any ϕ

We thus have a complete characterisation of theories for which *FS* is complete. However, the use of the above corollary to actually check whether a given theory is finitary may be rather difficult and impractical. We therefore inquire now into other ways of checking finitariness of a theory.

6 Algebraic conditions for completeness

In this section we study finitariness from a semantic perspective. The main results are algebraic conditions on the general models of a theory, which will be sufficient to ensure finitariness of the theory. Note that these are conditions on the *general* models introduced in Section 4, but they are used to ensure completeness with respect to *finite* models. Thus we use general models as a tool to prove completeness with respect to finite models. Given a theory Γ, we construct its class of general models, and if that class satisfies the algebraic conditions given below, *FS* is Γ-complete with respect to finite models.

The following Definition 6.1 will form the basis for checking finitariness of a theory by checking the respective properties of its class of general models. The involved conditions of this definition are strengthened and simplified in Lemma 6.9 and Lemma 6.10. The main result is Theorem 6.2. Some examples are found at the end of this section, and in the next section. The remainder of this section is concerned with intermediate, and rather technical, results for proving Theorem 6.2. First some terminology:

Directed Set A set A with a reflexive and transitive relation $\leq$ is *directed* iff for every finite subset B of A, there is an element $a \in A$ such that $b \leq a$ for every $b \in B$. In the following directedness of a set of sets is implicitly taken to be with respect to subset inclusion.

Cover A family of subsets of a set A whose union includes A is a *cover* of, or *covers*, A.

DEFINITION 6.1 (Finitary Set of General Models)
If $S \subseteq \mathcal{S}^*$ is a set of general models and $s \in \wp(U)$, then the set of finite subsets of s included in S is denoted

$$S|_s^f = S \cap \wp^{fin}(s).$$

$S \subseteq \mathcal{S}^*$ is *finitary* iff both:

1. For every infinite set $s \in S$:
   (a) $S|_s^f$ is directed
   (b) $S|_s^f$ is a cover of s
2. $\forall s \uplus \{*\} \in S : \forall s' \in \wp^{fin}(U) : \exists \alpha \in U \setminus s'$:
   (a) $\exists s^f \in S \cap \wp(s \cup \{\alpha\}) : s' \cap s \subseteq s^f$
   (b) $\exists s^f \in S \cap \wp(s \cup \{\alpha\}) : s^f \not\subseteq s'$
   (c) $S \cap \wp(s \cup \{\alpha\})$ is directed

The definition specifies conditions for each infinite set in S, condition 1, and each finite set in S containing $*$, condition 2. (The notation $s \uplus \{*\}$ stands for $s \cup \{*\}$ where $* \notin s$). Condition 1 requires every infinite model to be a limit of approximations by finite models. Condition 2 is similar, but it is complicated by the fact that, informally speaking, the existence of a "proper" element of the universe α to "replace" $*$ is needed. In practice, the simplified (and stronger) conditions presented in Lemma 6.10 below can often be used.

The following connects the syntactic (Definition 5.2) and the semantic (Definition 6.1) notions of finitariness.

THEOREM 6.2
A theory Γ is finitary if $\mathcal{S}^*(\Gamma)$ is finitary.

The following definitions and intermediate results are needed in the main proof of Theorem 6.2, which then follows.

Given a set of general models $S \subseteq \mathcal{S}^*$, we henceforth let S^f denote the subset of *finite* model sets (without the $*$ element) it contains:

$$S^f = S \cap \wp^{fin}(U)$$

Condition 1 of Definition 6.1 can be expressed equivalently as follows.

LEMMA 6.3
Let $S \subseteq \mathcal{S}^*$ be a set of general models. Definition 6.1.1 holds iff for every infinite $s \in S$

$$\forall s' \in \wp^{fin}(s) : \exists s^f \in S^f : s' \subseteq s^f \subseteq s \tag{3}$$

PROOF.

$\Rightarrow$) Assume that Def. 6.1.1 holds, and let $s' \in \wp^{fin}(s)$. s' is finite, say $s' = \{\beta_1, \ldots, \beta_k\}$. Because $s' \subseteq s$, by Def. 6.1.1.b) $s' \subseteq \bigcup(S \cap \wp^{fin}(s))$, so for each β_j there is a $t_j \in (S \cap \wp^{fin}(s))$ such that $\beta_j \in t_j$. By Def. 6.1.1.a), there is an $s^f \in (S \cap \wp^{fin}(s))$ such that $\bigcup_{1 \le j \le k}\{t_j\} \subseteq s^f$. Then $s' \subseteq s^f$. Since $s^f \in (S \cap \wp^{fin}(s))$, $s' \in S^f$ and $s^f \subseteq s$.

$\Leftarrow$) Assume that (3) holds.

Def. 6.1.1.a) Let S' be a finite subset of $S\cap\wp^{fin}(s)$. Clearly, $s'=\bigcup S'\in\wp^{fin}(s)$, and by (3) there is an $s^f\in S\cap\wp^{fin}(s)$ such that $s'\subseteq s^f$.

Def. 6.1.1.b) Let $\alpha\in s$. Because $\{\alpha\}\in\wp^{fin}(s)$, by (3) there is an $s^f\in S\cap\wp^{fin}(s)$ such that $\alpha\in s^f$. Then, $\alpha\in\bigcup(S\cap\wp^{fin}(s))$. Thus, $s\subseteq\bigcup(S\cap\wp^{fin}(s))$, and Def. 6.1.1.b holds. ■

Lemma 6.4
Let $S\subseteq\mathcal{S}^*$ be a set of general models. If S is finitary then:

1. For every infinite $s\in S$:
   (a) $\exists s^f\in S^f : s^f\subseteq s$
   (b) $\forall s'\in\wp^{fin}(U):\exists s^f\in S^f:(s^f\subseteq s \text{ and } s^f\not\subseteq s')$
2. $\forall s\uplus\{*\}\in S:\forall s'\in\wp^{fin}(U):\exists\alpha\in U\setminus s'$:
   (a) $\forall s''\subseteq s': s''\subseteq s\Rightarrow\exists s^f\in S\cap\wp(s\cup\{\alpha\}): s''\subseteq s^f$
   (b) $\forall s''\subseteq s':\exists s^f\in S\cap\wp(s\cup\{\alpha\}): s^f\not\subseteq s''$
   (c) $S\cap\wp(s\cup\{\alpha\})$ is directed

Proof.

1. Let $s\in S$ be infinite.
   (a) Follows from Lemma 6.3 by letting $s'=\varnothing$.
   (b) Assume that 1b) does *not* hold for s, i.e. that

$$\exists s'\in\wp^{fin}(U):\forall s^f\in S^f:(s^f\subseteq s\Rightarrow s^f\subseteq s')$$

   That is, there is an $s'\in\wp^{fin}(U)$ such that

$$\forall s^f\in S\cap\wp^{fin}(s): s^f\subseteq s'$$

   in other words

$$(\bigcup(S\cap\wp^{fin}(s)))\subseteq s'$$

   Since s is infinite and s' is finite, $s\not\subseteq s'$ and thus

$$s\not\subseteq(\bigcup(S\cap\wp^{fin}(s)))$$

   which contradicts the fact that $\bigcup(S\cap\wp^{fin}(s))$ covers s. Thus, 1b) must hold.
2. Let $s\cup\{*\}\in S$ $(*\notin s)$ and $s'\in\wp^{fin}(U)$, and let α be as defined in Def. 6.1.2.
   (a) Let $s''\subseteq s'$ and $s''\subseteq s$. By Def. 6.1.2.a there is an $s^f\in S\cap\wp(s\cup\{\alpha\})$ such that $s'\cap s\subseteq s^f$. Since $s''\subseteq s'\cap s$, $s''\subseteq s^f$ which proves 2a).
   (b) Let $s''\subseteq s'$. By Def. 6.1.2.b there is an $s^f\in S\cap\wp(s\cup\{\alpha\})$ such that $s^f\not\subseteq s'$. Then, $s^f\not\subseteq s''$ which proves 2b).
   (c) Def. 6.1.2.c. ■

In the proof of the main theorem below, we show the existence of a finite model from the assumption of the existence of a general model. This requires replacing a possible occurrence of $*$ by a "proper" element of the universe and we begin by showing the existence of such an element in Lemma 6.6. When ϕ, ψ are formulae, we use the notation $\psi \leq \phi$ to denote that ψ is a (not necessarily proper) subformula of ϕ.

Definition 6.5 (s^ϕ)
Given a formula $\phi \in \mathcal{L}$,

$$s^\phi = \bigcup_{\sqsupseteq(X)\leq\phi \text{ or } \sqsubseteq(X)\leq\phi} X$$

Observe that s^ϕ is always finite, since all X are finite and a formula only has a finite number of subformulae.

Lemma 6.6
If $\phi \in \mathcal{L}$, $S \subseteq \mathcal{S}^*$ is a finitary set of general models and $s = \hat{s} \uplus \{*\} \in S$, then there exists a

$$\alpha_S^{\phi,s} \in U \setminus s^\phi \tag{4}$$

where s^ϕ is defined in Def. 6.5, such that

1. $\forall s'' \subseteq s^\phi : s'' \subseteq \hat{s} \Rightarrow \exists s^f \in S \cap \wp(\hat{s} \cup \{\alpha_S^{\phi,s}\}) : s'' \subseteq s^f$
2. $\forall s'' \subseteq s^\phi : \exists s^f \in S \cap \wp(\hat{s} \cup \{\alpha_S^{\phi,s}\}) : s^f \not\subseteq s''$
3. $S \cap \wp^{fin}(\hat{s} \cup \{\alpha_S^{\phi,s}\})$ is directed

Note that, given S, s and ϕ, there may exist more than one $\alpha \in U \setminus s^\phi$ satisfying the three properties above, but we select one of them (arbitrarily) and call it $\alpha_S^{\phi,s}$.

Proof. Follows from Lemma 6.4.2, since S is finitary, $\hat{s} \cup \{*\} \in S$ and $s^\phi \in \wp^{fin}(U)$. ■

Definition 6.7 ($\tilde{s}_S^{\phi,s}$)
Let $\phi \in \mathcal{L}$, S be a finitary set of general models and $s \in S$.
Let

$$\tilde{s}_S^{\phi,s} = \begin{cases} s & \text{if } * \notin s \\ (s \setminus \{*\}) \cup \{\alpha_S^{\phi,s}\} & \text{if } * \in s \end{cases}$$

($\tilde{s}_S^{\phi,s}$ is s possibly with the asterisk replaced by $\alpha_S^{\phi,s}$).

Proof of theorem 6.2. Let Γ be a theory such that $\mathcal{S}^*(\Gamma)$ is finitary. Henceforth, let $S = \mathcal{S}^*(\Gamma)$. Let ϕ be an arbitrary formula, and assume that there is a general model $s \in S$ satisfying ϕ. We will show that there is a model $s^f \in S^f$ satisfying ϕ, which proves the theorem by Corollary 5.8.

Either $* \in s$ or $* \notin s$. When $* \in s$, the following shorthand notation is used:

$$\alpha = \alpha_S^{\phi,s}$$

where $\alpha_S^{\phi,s}$ is defined in Lemma 6.6, where

$$\alpha \notin s^\phi$$

Similarly, the following shorthand notation is used (see Def. 6.7):

$$\tilde{s} = \tilde{s}_S^{\phi,s}$$

If $\psi \in \mathcal{L}$, let $L(\psi)$ be the following statement

$$L(\psi) : \exists s^f \begin{cases} \text{a) } s^f \in S^f \\ \text{b) } s^f \subseteq \tilde{s} \\ \text{c) } s' \models_* \psi \text{ for all } s' \text{ s. t. } s^f \subseteq s' \subseteq \tilde{s} \end{cases}$$

and let $P(\psi)$ be the following statement

$$P(\psi) : \psi \leq \phi \Rightarrow \begin{cases} \text{1) } s \models_* \psi \Rightarrow L(\psi) \\ \text{2) } s \models_* \neg\psi \Rightarrow L(\neg\psi) \end{cases}$$

If $P(\phi)$ holds, then, since $\phi \leq \phi$, $L(\phi)$ holds by 1) and the fact that $s \models_* \phi$. $P(\phi)$ is a stronger statement than the theorem, and is needed for the inductive structure of the proof. By taking $s' = s^f$ in c) (in $L(\phi)$), we get that $s^f \models \phi$ for $s^f \in S^f$, which proves the theorem.

Before the main proof of $P(\phi)$, one property of S is shown:

$$\exists s^f \in S^f : s^f \subseteq \tilde{s} \tag{5}$$

To see that (5) holds, first consider the case that s is finite. If $* \notin s$, then $s^f = s \in S^f$ and $s^f \subseteq \tilde{s} = s$. If $* \in s$, then there is a $s^f \subseteq \tilde{s}$ by Lemma 6.6.2 (take, e.g., $s'' = s^\phi$) and $s^f \in S^f$ since $\tilde{s}$ is finite. Second, in the case that s is infinite then $\tilde{s} = s$ and (5) holds by Lemma 6.4.1.a.

We now prove $P(\psi)$ for all formulae ψ (including ϕ), by induction over the structure of ψ.[1]

$\psi = \sqsupseteq(X)$**:** Assume that $\sqsupseteq(X) \leq \phi$.

1. Assume that $s \models_* \sqsupseteq(X)$, i.e., $X \subseteq s$. We show $L(\sqsupseteq(X))$ in the following three cases:
   i) $* \notin s$ and s finite: let $s^f = s$, then a), b) and c) hold trivially.
   ii) $* \notin s$ and s infinite: by Lemma 6.3, since X is a finite subset of s, there is an $s^f \in S \cap \wp^{fin}(s)$, giving a) and b) ($\tilde{s} = s$), such that $X \subseteq s^f$. If $s^f \subseteq s' \subseteq \tilde{s}$, then $X \subseteq s'$ and $s' \models_* \sqsupseteq(X)$ giving c).
   iii) $* \in s$: by Lemma 6.6.1, since $X \subseteq s^\phi$ and $X \subseteq (s \setminus \{*\})$, there is an $s^f \in S \cap \wp^{fin}(\tilde{s})$, giving a) and b), such that $X \subseteq s^f$. If $s^f \subseteq s' \subseteq \tilde{s}$, then $X \subseteq s'$ and $s' \models_* \sqsupseteq(X)$ giving c).
2. Assume that $s \models_* \neg\sqsupseteq(X)$, i.e., $X \not\subseteq s$. We show $L(\neg\sqsupseteq(X))$. By (5) there exists an s^f such that $s^f \in S^f$ and $s^f \subseteq \tilde{s}$, giving a) and b). Let s' be such that $s^f \subseteq s' \subseteq \tilde{s}$. First, consider that $* \notin s$. Then $\tilde{s} = s$, and since $X \not\subseteq s$, $X \not\subseteq s'$. Second, consider that $* \in s$. Assume that $X \subseteq s'$. Then $X \subseteq \tilde{s}$ but, since $\sqsupseteq(X) \leq \phi$, $\alpha \notin X$ by definition of α, so $X \subseteq (s \setminus \{*\})$. But this is a contradiction, since $(s \setminus \{*\}) \subset s$ and $X \not\subseteq s$, so the assumption that $X \subseteq s'$ is impossible. Thus in either case, $X \not\subseteq s'$, and $s' \models_* \neg\sqsupseteq(X)$ giving c).

$\psi = \sqsubseteq(X)$**:** Assume that $\sqsubseteq(X) \leq \phi$.

1. Assume that $s \models_* \sqsubseteq(X)$, i.e., $s \subseteq X$. We show $L(\sqsubseteq(X))$. Simply choosing $s^f = s$ suffice: a) holds since $s \in S$ and s is finite since $s \subseteq X$. b) holds since $s \subseteq X \Rightarrow * \notin s \Rightarrow \tilde{s} = s = s^f \Rightarrow s^f \subseteq \tilde{s}$. Let s' be such that $s^f \subseteq s' \subseteq \tilde{s}$. Since $s^f = \tilde{s}$, $s' = \tilde{s} = s$. $s' \subseteq X$, and $s' \models_* \sqsubseteq(X)$.

[1] Keep in mind that $s, \phi, \tilde{s}$, and α whenever $* \in s$, are fixed *before* the inductive proof of P.

2. Assume that $s \models_* \neg\sqsubseteq(X)$, i.e., $s \not\subseteq X$. We show $L(\neg\sqsubseteq(X))$ in the following three cases:
   i) $* \notin s$ and s finite: let $s^f = s$, then a), b) and c) hold trivially.
   ii) $* \notin s$ and s infinite: $\tilde{s} = s$. By Lemma 6.4.1.b, since $X \in \wp^{fin}(U)$, there is an s^f such that $s^f \in S^f$ and $s^f \subset s = \tilde{s}$ giving a) and b), and such that $s^f \not\subseteq X$. If s' is such that $s^f \subseteq s' \subseteq \tilde{s} = s$, then $s' \not\subseteq X$ and $s' \models_* \neg\sqsubseteq(X)$ giving c).
   iii) $* \in s$: Since $X \subseteq s^\phi$, by Lemma 6.6.2 there is an s^f such that $s^f \in S^f$ and $s^f \subset \tilde{s}$, giving a) and b), and such that $s^f \not\subseteq X$. If s' is such that $s^f \subseteq s' \subseteq \tilde{s}$, then $s' \not\subseteq X$ and $s' \models_* \neg\sqsubseteq(X)$ giving c).

$\psi = \neg\psi_1$**:** The induction hypothesis is $P(\psi_1)$. Assume that $\neg\psi_1 \leq \phi$; then also $\psi_1 \leq \phi$.
1. Assume that $s \models_* \neg\psi_1$. Then, since $\psi_1 \leq \phi$, $L(\neg\psi_1) = L(\psi)$ holds by $P(\psi_1)$ 2).
2. Assume that $s \models_* \neg\neg\psi_1$. Then, $s \models_* \psi_1$ and since $\psi_1 \leq \phi$, $L(\psi_1)$ holds by $P(\psi_1)$ 1). By the definition of L, $L(\psi_1)$ holds iff $L(\neg\psi) = L(\neg\neg\psi_1)$ holds.

$\psi = \psi_1 \wedge \psi_2$**:** The induction hypotheses are $P(\psi_1)$ and $P(\psi_2)$. Assume that $\psi_1 \wedge \psi_2 \leq \phi$.
1. Assume that $s \models_* \psi_1 \wedge \psi_2$. $\psi_1 \wedge \psi_2 \leq \phi$ implies that $\psi_1 \leq \phi$, and $s \models_* \psi_1 \wedge \psi_2$ implies that $s \models_* \psi_1$, and thus, by $P(\psi_1)$, $L(\psi_1)$ holds. That is, there exists $s_1^f \in S^f$ such that $s_1^f \subseteq \tilde{s}$ and for all s_1' such that $s_1^f \subseteq s_1' \subseteq \tilde{s}$, $s_1' \models_* \psi_1$. Similarly, by $P(\psi_2)$, $L(\psi_2)$ holds; there exist, $s_2^f \in S^f$ such that $s_2^f \subseteq \tilde{s}$ and for all s_2' such that $s_2^f \subseteq s_2' \subseteq \tilde{s}$, $s_2' \models_* \psi_2$. We show $L(\psi_1 \wedge \psi_2)$. Since $S \cap \wp^{fin}(\tilde{s})$ is directed (by Def. 6.1.1.a when $* \notin s$ and by Lemma 6.6.3 when $* \in s$ (recall that s is finite when $* \in s$)) and $s_1^f, s_2^f \in S \cap \wp^{fin}(\tilde{s})$, there exists an $s^f \in S \cap \wp^{fin}(\tilde{s})$ such that $s_1^f, s_2^f \subseteq s^f$. a) holds since $s^f \in S$ is finite, and b) holds since $s^f \in \wp^{fin}(\tilde{s})$. Let s' be such that $s^f \subseteq s' \subseteq \tilde{s}$. Because $s_1^f \subseteq s^f \subseteq s' \subseteq \tilde{s}$, $s_1^f \subseteq s' \subseteq \tilde{s}$ and, by $L(\psi_1)$, $s' \models_* \psi_1$. Similarly, because $s_2^f \subseteq s^f \subseteq s' \subseteq \tilde{s}$, $s_2^f \subseteq s' \subseteq \tilde{s}$ and, by $L(\psi_2)$, $s' \models_* \psi_2$. Thus, $s' \models_* \psi_1 \wedge \psi_2$, and c) holds.
2. Assume that $s \models_* \neg(\psi_1 \wedge \psi_2)$; $s \models_* \neg\psi_1 \vee \neg\psi_2$; $s \models_* \neg\psi_1$ or $s \models_* \neg\psi_2$. Assume the first case (the proof in the second case is symmetrical). $\psi_1 \wedge \psi_2 \leq \phi$ implies that $\psi_1 \leq \phi$ and since $s \models_* \neg\psi_1$, $L(\neg\psi_1)$ holds by $P(\psi_1)$. That is, there exist $s^f \in S^f$ such that $s^f \subseteq \tilde{s}$ and for all s' such that $s^f \subseteq s' \subseteq \tilde{s}$, $s' \models_* \neg\psi_1$. But then we also have that $s' \models_* \neg(\psi_1 \wedge \psi_2)$ (i.e., s^f, the witness in $L(\neg\psi_1)$, is also a witness in $L(\neg(\psi_1 \wedge \psi_2))$). ■

Corollary 6.8
FS is Γ-complete if $\mathcal{S}^*(\Gamma)$ is finitary.

Theorem 6.2 shows that the conditions in Definition 6.1 on the class of general models of a theory are sufficient to conclude that the axioms are finitary. In the following lemmas, we present several simpler but stronger sufficient conditions (the proofs are straightforward and left for the reader).

Lemma 6.9
A set of general models $S \subseteq \mathcal{S}^*$ is finitary if

1. For every infinite $s \in S$:
   (a) $S|_s^f$ is directed
   (b) $S|_s^f$ is a cover of s
2. $\forall s \uplus \{*\} \in S : \forall s' \in \wp^{fin}(U) : \exists \alpha \in U \setminus s'$:
   (a) $S|_{s \cup \{\alpha\}}^f$ is directed
   (b) $S|_{s \cup \{\alpha\}}^f$ is a cover of $s \cup \{\alpha\}$

The following conditions are even stronger than Lemma 6.9.

Lemma 6.10
$S \subseteq \mathcal{S}^*$ is finitary if either one of the following three conditions hold:

1. For every $s \subseteq U$:
   (a) $S|_s^f$ is directed
   (b) $S|_s^f$ is a cover of s
2. (a) $S|_s^f$ is directed for every $s \subseteq U$
   (b) $\{\alpha\} \in S$ for every $\alpha \in U$
3. (a) $S|_s^f$ is directed for every infinite $s \in S$
   (b) $\{\alpha\} \in S$ for every $\alpha \in U$
   (c) $\forall s \uplus \{*\} \in S : \forall s' \in \wp^{fin}(U) : \exists \alpha \in U \setminus s' : s \cup \{\alpha\} \in S$

Thus, given a theory Γ, we now have a tool for proving Γ-completeness: take the *general* model class $\mathcal{S}^*(\Gamma)$ and check whether any of the conditions given above hold. If the answer is positive, we have Γ-completeness: $\Gamma \vdash \phi$ for all ϕ such that $\Gamma \models \phi$.

The sufficient conditions for (semantic) finitariness above can be used to show that the empty theory is (syntactically) finitary, and thus that *FS* is weakly complete, i.e., that $\models \phi$ implies that $\vdash \phi$ for all ϕ.

Corollary 6.11 (Weak Completeness)
FS is weakly complete wrt. $\mathcal{S}^{fin}$.

Proof. It suffices to show that $\mathcal{S}^*(\varnothing) = \mathcal{S}^*$ is finitary. Then, it follows by Theorem 6.2 that $\varnothing$ is a finitary theory, and thus Theorem 5.7 implies that *FS* is $\varnothing$-complete, i.e., weakly complete.

We make use of Lemma 6.10.1. Let $s \subseteq U$. $\mathcal{S}^*|_s^f = \mathcal{S}^* \cap \wp^{fin}(s) = \wp^{fin}(s)$. $\wp^{fin}(s)$ is directed, because for every finite subset $B \subset \wp^{fin}(s)$, $\cup_{s' \in B} s' \in \wp^{fin}(s)$. $\wp^{fin}(s)$ is a cover of s, because $s \subseteq \bigcup \wp^{fin}(s)$. ■

Since $\models \phi \to \psi$ iff $\forall s \in \mathcal{S}^{fin} : s \models \phi$ implies $s \models \psi$, the corollary implies also that for every finite Γ, *FS* is Γ-complete. Although the particular case of weak completeness can be proved directly in a less complicated manner, the above proofs and statements reflect the interest in the more general properties of finitariness. The next section gives some examples.

7 Example: finite epistemic states

Epistemic or *doxastic* logics [14, 11, 19] are used to reason about knowledge and belief. Most often such logics are based on propositional modal logic, with a modality B with the intended meaning of $B\alpha$ that α is believed or known. Epistemic logics vary according to which principles of knowledge they accept, some common ones being:

$B(\alpha \to \beta) \to (B\alpha \to B\beta)$	Distribution	**K**
$B\alpha \to \neg B \neg \alpha$	Consistency	**D**
$B\alpha \to BB\alpha$	Positive Introspection	**4**
$\neg B\alpha \to B\neg B\alpha$	Negative Introspection	**5**

All (normal) modal epistemic logics agree on one principle, namely that belief is closed under logical consequence: if α is believed and ψ is a logical consequence of α, then ψ is also believed. It follows, e.g., that it is assumed that all tautologies are believed, that an inconsistent belief entails belief in everything, and that an agent always has infinitely many beliefs. Thus, such logics give an account of *implicit* beliefs, or belief after an indefinite amount of time, or the beliefs of ideal reasoners, but do not model the explicit beliefs of real resource-bounded agents at a certain point in time very well. This is commonly known as the *logical omniscience problem* [13].

In particular, modal epistemic logic does not seem to be very well suited to model the explicitly computed beliefs, at some point in time, of reasoners who store their beliefs syntactically as logical formulae (e.g., in a database or written on a piece of paper). In practice, closure conditions such as closure under logical consequence on such belief states can not be assumed, because a certain inference might not yet have been made. Often no consistency condition can be assumed either – real agents often hold contradictory beliefs. Finally, a real reasoner can only store a finite number of formulae simultaneously. The logical closure assumption in modal epistemic logics makes them inadequate for modeling finite sets of beliefs.

A logic for reasoning about belief as arbitrary finite sets of formulae can be obtained by interpreting formulae $B\alpha$ using syntactic assignments instead of Kripke structures (see, e.g., [11, Ch. 9]). Syntactic assignments are generalisations of both Kripke structures and Montague-Scott structures [20, 21, 28]. The literature contains numerous proposed solutions to the logical omniscience problem, see, e.g., [22, 29, 11] for reviews. Wansing [30] shows that many of these approaches can be modeled using Rantala models [25, 26], and that Rantala models can be seen as the most general models of knowledge. It is easy to see that syntactic structures are as general as Rantala models. However, the logic one obtains, assuming the traditional propositional belief language, is simply propositional logic – this notion of belief has no non-trivial properties expressible in the language. By using our logic of finite sets, however, we obtain a more expressive language for reasoning about finite syntactic belief states. Thus, take the universe U to be some *object language* OL – the language we assume that the agents store their beliefs in (propositional logic, traditional epistemic logic, even our language $\mathcal{L}$ itself, etc.). Formulae are thus interpreted as statements about a finite set of beliefs, and we can now read the $\supseteq$ operator as B: $\supseteq\alpha$ means that α is included in the belief set. Furthermore, $\sqsupseteq(X)$, where X is a finite set of object formulae, means that *at least* X is believed (stored). Correspondingly, $\sqsubseteq(X)$ means that *at most* X is believed. While $\sqsupseteq(X)$ is already definable in the traditional propositional belief language, $\sqsupseteq(X) \equiv \bigwedge_{\alpha\in X} B\alpha$, $\sqsubseteq(X)$ is not – it corresponds to a conjunction $\sqsubseteq(X) \equiv \bigwedge_{\alpha\in OL\setminus X} \neg B\alpha$ in the case that OL is finite, but the conjunction would be infinite in the typical case of an infinite object language (e.g., an object language closed under propositional connectives). This notion of believing at most can be seen as a syntactic version of the notion of *only knowing* [16] in modal epistemic logic. Thus, FS is a sound and weakly complete axiomatisation for reasoning about finite syntactic belief states in the language of believing at least and believing at most.

FS axiomatises unrestricted finite belief sets. We might want to restrict the possible belief sets s by imposing closure conditions which do not entail infiniteness, or consistency conditions. For example, we may want to axiomatise the class of belief sets not containing both α and $\neg\alpha$ for any α – i.e., the class of model sets $\mathcal{S}^{fin}(\mathbf{D})$ for the **D** axiom above. Or the class $\mathcal{S}^{fin}(\mathbf{K})$ of finite sets containing $B\beta$ whenever they contain $B(\alpha \to \beta)$ and $B\alpha$; $\mathcal{S}^{fin}(\mathbf{4})$ of finite sets containing $BB\alpha$ whenever they contain $B\alpha$; or $\mathcal{S}^{fin}(\mathbf{5})$ of finite sets containing $B\neg B\alpha$ whenever they contain $\neg B\alpha$. The results in the previous sections can help us investigate axiomatisations of such classes.[2] The following are some examples. Let $FS\mathbf{A}$ denote the system obtained by adding axiom schema $\mathbf{A}$ to FS (note that we view an axiom schema as an (infinite) set of formulae). Henceforth, assume that the object language OL is closed under propositional connectives and the B ($\supseteq$) operator.

Theorem 7.1 (Completeness Results)

1. $FS\mathbf{K}$ is sound and weakly complete with respect to $\mathcal{S}^{fin}(\mathbf{K})$
2. $FS\mathbf{D}$ is sound and weakly complete with respect to $\mathcal{S}^{fin}(\mathbf{D})$
3. $FS\mathbf{4}$ is not complete with respect to $\mathcal{S}^{fin}(\mathbf{4})$
4. $FS\mathbf{5}$ is not complete with respect to $\mathcal{S}^{fin}(\mathbf{5})$

Proof. Consider first the two former parts of the theorem. Soundness follows immediately from Corollary 5.1 and the fact that $\mathbf{K}$ and $\mathbf{D}$ are valid in $\mathcal{S}^{fin}(\mathbf{K})$ and $\mathcal{S}^{fin}(\mathbf{D})$, respectively. For the completeness claims, we have the following classes of *general* models, for the two axioms, respectively:

$$\begin{aligned} \mathcal{S}^*(\mathbf{K}) &= \mathcal{S}^* \setminus \{s \in \mathcal{S}^* : \exists \alpha, \beta \in OL : \alpha \to \beta, \alpha \in s, \beta \notin s\} \\ \mathcal{S}^*(\mathbf{D}) &= \mathcal{S}^* \setminus \{s \in \mathcal{S}^* : \exists \alpha \in OL : \alpha, \neg\alpha \in s\} \end{aligned}$$

It suffices to show that these classes of general models are finitary. It then follows by Corollary 6.8 that FS is $\mathbf{K}$-complete and $\mathbf{D}$-complete. Thus, if ϕ is valid on $\mathcal{S}^{fin}(\mathbf{K})$ then ϕ is derivable from the combination of FS and $\mathbf{K}$, and similarly for $\mathbf{D}$.

For the two latter parts of the theorem, we show that $\mathbf{4}$ and $\mathbf{5}$ are not finitary theories; it follows by Theorem 5.7 that FS is not $\mathbf{4}$-complete or $\mathbf{5}$-complete. Thus, there exists formulae ϕ such that ϕ is valid on $\mathcal{S}^{fin}(\mathbf{4})$ but ϕ is not derivable from the combination of FS and $\mathbf{4}$, and similarly for $\mathbf{5}$.

1. Lemma 6.10.3 holds for $\mathcal{S}^*(\mathbf{K})$:

   Lemma 6.10.3.(a): It must be shown that $\mathcal{S}^*(\mathbf{K})|_s^f$ is directed for infinite $s \in \mathcal{S}^*(\mathbf{K})$. Let $s', s'' \in \mathcal{S}^*(\mathbf{K}) \cap \wp^{fin}(s)$, and let:

   $$\begin{aligned} s_0 &= s' \cup s'' \\ s_j &= s_{j-1} \cup \{\beta : \alpha \to \beta, \alpha \in s_{j-1}\} \\ s^f &= \bigcup_k s_k \end{aligned}$$

   It is easy to show that $s^f \in \mathcal{S}^*(\mathbf{K})$, each s_j is a finite subset of s, and s^f is finite.

   Lemma 6.10.3.(b): Clearly, $\{\alpha\} \in \mathcal{S}^*(\mathbf{K})$ for every $\alpha \in OL$.

   Lemma 6.10.3.(c): Let $s \uplus \{*\} \in \mathcal{S}^*(\mathbf{K})$ and $s' \in \wp^{fin}(OL)$. Let $\alpha \in OL$ be s. t.:

   - $\alpha \to \beta \notin s$ for any $\beta \in OL$
   - $\alpha \notin s'$
   - The main connective in α is not implication

   It is easy to see that there exist infinitely many α satisfying these three conditions; there are infinitely many $\alpha \in OL$ without implication as main connective, and both s and s' are finite. It can easily be shown that $s \cup \{\alpha\} \in \mathcal{S}^*(\mathbf{K})$.

[2] Note that, for instance, $\mathcal{S}^{fin}(\mathbf{K})$ denotes the class of models of this one axiom *only*, not of the modal logic (also called) $\mathbf{K}$.

2. Lemma 6.10.3 holds for $\mathcal{S}^*(\mathbf{D})$:
   Lemma 6.10.3.(a): It must be shown that $\mathcal{S}^*(\mathbf{D})|_s^f$ is directed for infinite $s \in \mathcal{S}^*(\mathbf{D})$. Let $s', s'' \in \mathcal{S}^*(\mathbf{D}) \cap \wp^{fin}(s)$, and let $s^f = s' \cup s''$. It can easily be shown that $s^f \in \mathcal{S}^*(\mathbf{D})$, and $s^f \in \wp^{fin}(s)$ trivially.
   Lemma 6.10.3.(b): Clearly, $\{\alpha\} \in \mathcal{S}^*(\mathbf{D})$ for every $\alpha \in OL$.
   Lemma 6.10.3.(c): Let $s \uplus \{*\} \in \mathcal{S}^*(\mathbf{D})$ and $s' \in \wp^{fin}(OL)$. Let $\alpha \in OL$ be s. t.:

   - $\neg\alpha \notin s$
   - $\alpha \notin s'$
   - α does not start with negation

   It is easy to see that there exist infinitely many α satisfying these three conditions; there are infinitely many $\alpha \in OL$ without negation as main connective, and both s and s' are finite. It can easily be shown that $s \cup \{\alpha\} \in \mathcal{S}^*(\mathbf{D})$.
3. Let $s \in \mathcal{S}^{fin}$ be such that $s \models \mathbf{4}$. s must be the empty set – otherwise it would not be finite. Thus, $\mathbf{4} \models \sqsubseteq(\varnothing)$. $\mathbf{4}$ does, however, have *infinite* models, so $\mathbf{4} \not\models_* \sqsubseteq(\varnothing)$. Corollary 5.8.4 gives that $\mathbf{4}$ is not finitary.
4. It is easy to see that $\mathbf{5}$ is not satisfiable in $\mathcal{S}^{fin}$ (i.e., that a model for $\mathbf{5}$ must be infinite). By Lemma 5.5 and Lemma 5.4, $\mathbf{5}$ is not finitary. ■

Of course, it is intuitively clear that **K** and **D** are compatible with finite beliefs and that **4** and **5** are not, but the examples above show how this (in)compatibilty is manifested as (in)completeness and illustrate how, e.g., completeness can be *proved*. Note that in the case of *FS***4**, unlike *FS***5**, there *are* finite models (namely the empty set).

The above examples show the possibilities of reasoning about *static* syntactic belief sets at a given point in time. In [3] the *dynamics* of such states are studied by adding ‘believing at least’ and ‘believing at most’ operators to modal logics.

A remark

Most epistemic axioms require a certain relationship between the object language OL and the language $\mathcal{L}$. In the example above we required, for the **4** and **5** axioms to be well-formed, that OL is closed under $B(\supseteq)$.

Consider now the veridicality axiom $\mathbf{T}: B\alpha \to \alpha$. This axiom requires that $OL \subseteq \mathcal{L}$. However, that is not possible in the typical case when OL contains primitive propositions, such as when OL is the traditional language of epistemic logic, because $\mathcal{L}$ does not contain such propositions. This problem can be circumvented by slightly modifying the definition of the meta language $\mathcal{L}$ to allow primitive propositions.[3] An interesting consequence of a finitary semantics is that such an axiom precludes knowing the (finite) limit of one's knowledge. If knowledge implies truth, $\mathbf{T}$, then there is no (finite) s and X such that $s \models \sqsupseteq(\{\sqsubseteq(X)\})$. For assume that there are such an s and X. Then, $s \models \sqsupseteq(\{\sqsubseteq(X)\})$ iff $\sqsubseteq(X) \in s$, while by $\mathbf{T}$, also $s \models \sqsubseteq(X)$, i.e., $s \subseteq X$. Thus $\sqsubseteq(X) \in X$ which is impossible, at least as long as we are working with the usual well-founded syntax and set theory.

[3] At the same time, $\mathcal{L}$ is defined as a function of $U = OL$. $\mathcal{L}$ and OL can in this case more precisely be defined simultaneously as follows. We take as the basis some, possibly empty, set of primitive atoms $\mathcal{L}_0 = A$, and define the formulae ϕ_{i+1} of the language $\mathcal{L}_{i+1}(\mathcal{L}_i)$ as in (1) with the additional clause $\ldots | \phi_i$. $\mathcal{L}$ is then the least fixed point, relative to a given A, which exists by the Knaster-Tarski theorem. With $A \neq \varnothing$, the notion of a model is extended with the valuation of the elements of A to handle their propositional combinations. Details can be found in [1].

Finally we remark that, as discussed in the beginning of this section, the interpretation of the belief operator used here is very different from the interpretation in classical epistemic logic, and thus that the results obtained here are not directly transferable to, or comparable to results in, modal logic.

8 Discussion and conclusions

We introduced a propositional logic over expressions $\sqsubseteq(X)$ and $\sqsupseteq(X)$, where X is a finite set, interpreted as statements about some (other) finite set.

The main results are an axiomatisation of the logic, characterisation of finitely satisfiable theories (finitarily open theories), and two characterisations of additional axioms for which the resulting logic is complete. The first, the notion of finitary theories, is proof-theoretic. The second, in form of algebraic conditions on the class of general models, is a semantic one, but gives only sufficient conditions for finitariness. The latter was used to show finitariness of the empty theory and thus weak completeness of the system. The semantic conditions give a general completeness proof, of which weak completeness is a special case, and the complexity of the proof is due to this generality. In general, we get the following methodology. We can prove that a theory *is* finitary by taking its class of general models, and check that it has one of the algebraic properties discussed in Section 6. Conversely, we can prove that a theory *is not* finitary, by using the finitariness properties discussed in Section 5. We showed some examples of both methods in Section 7. It remains for future work to show whether or not the semantic conditions in Definition 6.1 are also *necessary* for finitariness and, if not, tighten them to obtain conditions which are both sufficient and necessary. The idea of characterising a class of logics for which we have completeness is also used in Sahlqvist's completeness theorem for modal logic [27], but the characterisation does not seem to be directly comparable.

Most of the results presented here originate from [1], in which they are presented in the context of resource-bounded reasoners discussed in Section 7. Some related works on the two unary set operators in the same context have since appeared. There are some minor differences between variants of the language; the variant of the language we have used in the current article is identical to the language used in [4] (see further discussion below).[4] In [6] we study axiomatisations of the language with the inclusion and exclusion operators interpreted over general (possibly infinite) sets. In [3] modal logics extended with the two unary set operators are studied. The extended modal language is interpreted over Kripke structures where a finite set is associated with each state. The finite sets are assumed to be finite syntactic belief sets, as in Section 7, and the modal operators are used to express properties about how belief sets can evolve as a result of reasoning or communication. For example, $(\underline{\ni} p \wedge \underline{\ni}(p \rightarrow q)) \rightarrow \diamond \underline{\ni} q$ means that the agent can reason with modus ponens. In [2] a propositional language with nullary operators $min(n)$ and $max(n)$ interpreted in the context of a set, meaning, respectively, that the set has *at least* n elements and *at most* n

[4]In [1] and the mentioned follow-up works in the same context the symbols $\triangle$ and $\triangledown$ are used for the exclusion and inclusion operators $\sqsupseteq$ and $\sqsubseteq$, respectively. In these works, the finite sets represent sets of formulae in an agent's belief base. Two minor differences between the set operators in [1, 5], on the one hand, and the other mentioned related works and the current article, on the other are, first, that instead of an atomic symbol for each finite set, terms representing sets are built from symbols standing for individual elements and set-building operators; and, second, that several pairs of the unary set operators are allowed, each pair interpreted over a designated finite set representing the belief base of one of several agents.

elements, is studied, and the relative expressiveness of the operators $min(n)$, $max(n)$, $\sqsupseteq(X)$ and $\sqsubseteq(X)$ are compared.

In a separate strand of research, a language identical to the one discussed in this article was used by Ågotnes, van der Hoek and Wooldridge [4] in order to succinctly express the ability of coalitions of agents in a variant of *Coalition Logic* [24] called *Quantified Coalition Logic (QCL)*.[5] The main construct of QCL is of the form $\langle P\rangle\phi$ where ϕ is a formula of the QCL language and the *coalition predicate* P is a formula over the language we have discussed in this article taken over the universe of all agents (or agent *names*). Intuitively, $\langle P\rangle\phi$ means that there exists some coalition (finite set of agents) satisfying P which has the power to bring about ϕ (by acting in a certain way). For example, $\neg\langle\neg\sqsupseteq(\{a\})\rangle p$ means that agent a is a *weak veto player for* p: it is not the case that there exists a coalition C which does not include agent a and which can bring about p. The system FS presented in the current paper was in fact used as an axiomatisation of coalition predicates in [4]. However, in [4] the universe of all agents (the *grand coalition*) is taken to be finite, and most of the results in the current article are interesting only in the case of an infinite universe. Completeness in the case of a finite universe is a rather trivial special case of the general result since the inclusion operator then can be expressed in terms of the exclusion operator ($\sqsubseteq(X) \equiv \bigwedge_{e\in U\setminus X} \neg\sqsupseteq e$). An interesting opportunity for future work is to apply the results presented here to a setting of [4] with an infinite universe of agents.

Acknowledgments

This work has benefited very much from discussions with Natasha Alechina. The first author would also like to thank Wiebe van der Hoek and Michael Wooldridge for discussions on closely related issues. The NFR funded project SHIP provided a partial financial support. Finally, we thank the anonymous reviewers for helpful remarks.

References

[1] Thomas Ågotnes. *A Logic of Finite Epistemic States*. PhD thesis, Department of Informatics, University of Bergen, 2004.

[2] Thomas Ågotnes and Natasha Alechina. Knowing minimum/maximum n formulae. In Gerhard Brewka, Silvia Coradeschi, Anna Perini, and Paolo Traverso, editors, *Proceedings of the 17th European Conference on Artificial Intelligence (ECAI 2006)*, pages 317–321. IOS Press, 2006.

[3] Thomas Ågotnes and Natasha Alechina. The dynamics of syntactic knowledge. *Journal of Logic and Computation*, 17(1), 2007.

[4] Thomas Ågotnes, Wiebe van der Hoek, and Michael Wooldridge. Quantified coalition logic. In M. M. Veloso, editor, *Proceedings of the Twentieth International Joint Conference on Artificial Intelligence (IJCAI 2007)*, pages 1181–1186, California, 2007. AAAI Press.

[5] Thomas Ågotnes and Michal Walicki. Complete axiomatizations of finite syntactic epistemic states. In Matteo Baldoni, Ulle Endriss, Andrea Omicini, and Paolo Torroni,

[5] In [4] the symbols *supseteq* and *subseteq* are used for the exclusion and inclusion operators $\sqsupseteq$ and $\sqsubseteq$, respectively. In this work, the finite sets represent sets of agents.

editors, *Declarative Agent Languages and Technologies III: Third International Workshop, DALT 2005, Utrecht, The Netherlands, July 25, 2005, Selected and Revised Papers*, volume 3904 of *Lecture Notes in Computer Science (LNCS)*, pages 33–50. Springer Berlin / Heidelberg, 2006.

[6] Thomas Ågotnes and Michal Walicki. Strongly complete axiomatizations of "knowing at most" in syntactic structures. In Francesca Toni and Paolo Torroni, editors, *CLIMA VI*, volume 3900 of *Lecture Notes in Computer Science (LNCS)*, pages 57–76, London, UK, 2006. Springer Berlin / Heidelberg.

[7] K. R. Brown and H. Wang. Finite set theory, number theory and axioms of limitation. *Math. Annalen*, 164:26–29, 1966.

[8] S. R. Buss. *Bounded Arithmetic*. PhD thesis, Department of Mathematics, University of Princeton, 1985.

[9] Peter Clote and Jan Krajícek, editors. *Arithmetic, Proof Theory and Computational Complexity*. Oxford University Press, 1993.

[10] Heinz-Dieter Ebbinghaus and Jörg Flum. *Finite Model Theory*. Springer, 1999.

[11] Ronald Fagin, Joseph Y. Halpern, Yoram Moses, and Moshe Y. Vardi. *Reasoning About Knowledge*. The MIT Press, Cambridge, Massachusetts, 1995.

[12] Harvey Friedman. Some systems of second order arithmetic and their use. In *Proceedings of 1974 International Congress of Mathematicians*, 1975.

[13] J. Hintikka. Impossible possible worlds vindicated. *Journal of Philosophical Logic*, 4:475–484, 1975.

[14] Jaakko Hintikka. *Knowledge and Belief*. Cornell University Press, Ithaca, New York, 1962.

[15] Neil Immerman, Sushant Patnaik, and David Stemple. The expressiveness of a family of finite set languages. *Theoretical Computer Science*, 155(1):111–140, 1996.

[16] H. J. Levesque. All I know: a study in autoepistemic logic. *Artificial Intelligence*, 42:263–309, 1990.

[17] Shaughan Lewine. *Understanding the Infinite*. Cambridge University Press, 1994.

[18] Leonid Libkin. Locality of queries and transformations. In Ruy de Queiroz, Angus Macintyree, and Guilherme Bittencourt, editors, *Proceedings of WoLLIC'05*, 2005.

[19] J.-J. Ch. Meyer and W. van der Hoek. *Epistemic Logic for AI and Computer Science*. CUP, 1995.

[20] Richard Montague. Pragmatics. In R. Klibansky, editor, *Contemporary Philosophy: A Survey. I*, pages 102–122. La Nuova Italia Editrice, Florence, 1968.

[21] Richard Montague. Universal grammar. *Theoria*, 36:373–398, 1970.

[22] Antonio Moreno. Avoiding logical omniscience and perfect reasoning: a survey. *AI Communications*, 11:101–122, 1998.

[23] J. Mycielski. Locally finite theories. *Journal of Symbolic Logic*, 51:59–62, 1986.

[24] M. Pauly. A modal logic for coalitional power in games. *Journal of Logic and Computation*, 12(1):149–166, 2002.

[25] V. Rantala. Impossible worlds semantics and logical omniscience. *Acta Philosophica Fennica*, 35:106–115, 1982.

[26] V. Rantala. Quantified modal logic: non-normal worlds and propositional attitudes. *Studia Logica*, 41:41–65, 1982.

[27] H. Sahlqvist. Completeness and correspondence in the first and second order semantics for modal-logic. In S. Kanger, editor, *Proceedings of the Third Scandinavian Logic Symposium*, pages 110–143, Uppsala, 1975. North-Holland Publishing Company.

[28] Dana S. Scott. Advice on modal logic. In Karel Lambert, editor, *Philosophical Problems in Logic*, pages 143–173. D. Reidel Publishing Co., Dordrecht, 1970.

[29] Kwang Mong Sim. Epistemic logic and logical omniscience: A survey. *International Journal of Intelligent Systems*, 12:57–81, 1997.

[30] H. Wansing. A general possible worlds framework for reasoning about knowledge and belief. *Studia Logica*, 49(4):523–539, 1990.

Received 18 August 2007. Revised 16 April 2008