Advanced Logic: Lecture Notes

Theories, Models, and the Completeness of the Tableaux Method

Pascal Ludwig
Sorbonne Université

Contents

1 Axiomatised Theories

1.1 Deductive closure and theories

We begin by making precise the notion of a *theory* in first-order logic (FOL).

Definition 1.1 (Semantic and syntactic closure)**.** A set Γ of sentences is **semantically closed** if and only if, for every sentence φ, whenever $\Gamma \models \varphi$ then $\varphi \in \Gamma$. It is **syntactically closed** if and only if, whenever $\Gamma \vdash \varphi$ then $\varphi \in \Gamma$.

Because first-order logic is complete (a fact we shall prove in Section 5), these two notions coincide: a set of sentences is semantically closed if and only if it is syntactically closed. A deductively closed set of formulas is called a **theory**.

Definition 1.2 (Closure)**.** Given a set of sentences Γ, its **semantic closure** is $\{\varphi \mid \Gamma \models \varphi\}$, and its **syntactic closure** is $\{\varphi \mid \Gamma \vdash \varphi\}$. By completeness, these two sets are identical.

Definition 1.3 (Axiomatisation)**.** When a theory Γ is the closure of some set Δ (so that $\Delta \subseteq \Gamma$ and $\Gamma = \{\varphi \mid \Delta \models \varphi\}$), we say that Γ is a theory **axiomatised by** Δ, and that Δ is an **axiomatisation** of Γ.

First-order logic provides an excellent tool for expressing a domain of knowledge in an axiomatised way. Given a first-order language whose non-logical symbols (the *signature*) stand for the primitives of the science we wish to study, together with a set of sentences expressing the fundamental laws of that science, we can think of the resulting theory in two complementary ways. The first is *proof-theoretic*: the theory consists of all the sentences entailed by the axioms, that is, the deductive closure of the axiom set. The second is *model-theoretic*: the theory is the set of sentences satisfied by a given $\mathcal{L}$-structure.

1.2 Non-trivial axiomatisation

Every closed theory is trivially its own axiomatisation; but when the theory is infinite this observation is not very informative.

Definition 1.4 (Non-trivial axiomatisation)**.** A theory is **non-trivially axiomatisable** if and only if it is axiomatised by a set Δ that is **decidable**, meaning that there exists an effective procedure to determine whether a given sentence of the language belongs to Δ or not.

We take a method to be *effective* if and only if it is algorithmic: it can be mechanically carried out in a finite number of steps.

2 Theory of a Model

2.1 Definition

Definition 2.1 (Theory of a structure)**.** Let $\mathcal{M}$ be an $\mathcal{L}$-structure. The **theory of** $\mathcal{M}$, written $\mathrm{Th}(\mathcal{M})$, is the set

$$\mathrm{Th}(\mathcal{M}) \;=\; \{\varphi \in \mathcal{L} \mid \mathcal{M} \models \varphi\}.$$

Proposition 2.2. $\mathrm{Th}(\mathcal{M})$ *is a theory (i.e. it is deductively closed).*

Proof. Suppose $\mathrm{Th}(\mathcal{M}) \models \psi$. Then every model satisfying all formulas in $\mathrm{Th}(\mathcal{M})$ also satisfies ψ. In particular $\mathcal{M}$ satisfies all formulas in $\mathrm{Th}(\mathcal{M})$ (by definition), so $\mathcal{M} \models \psi$, hence $\psi \in \mathrm{Th}(\mathcal{M})$. □

The definition generalises to classes of structures.

Definition 2.3 (Theory of a class of structures)**.** Let $\mathcal{K}$ be a class of $\mathcal{L}$-structures. The **theory of** $\mathcal{K}$ is

$$\mathrm{Th}(\mathcal{K}) \;=\; \{\varphi \in \mathcal{L} \mid \forall \mathcal{M} \in \mathcal{K},\ \mathcal{M} \models \varphi\}.$$

Exercise 2.4. Prove that $\mathrm{Th}(\mathcal{K})$ is deductively closed.

2.2 Complete theories

Definition 2.5 (Completeness of a theory)**.** A theory Γ is **complete** if and only if, for every sentence φ in the language of the theory, either $\varphi \in \Gamma$ or $\neg\varphi \in \Gamma$.

Proposition 2.6. *For any structure $\mathcal{M}$, the theory* $\mathrm{Th}(\mathcal{M})$ *is complete.*

Proof. Suppose for contradiction that $\mathrm{Th}(\mathcal{M})$ is not complete. Then there is a sentence φ such that $\varphi \notin \mathrm{Th}(\mathcal{M})$ and $\neg\varphi \notin \mathrm{Th}(\mathcal{M})$. From $\varphi \notin \mathrm{Th}(\mathcal{M})$ we get $\mathcal{M} \not\models \varphi$, hence $\mathcal{M} \models \neg\varphi$, which gives $\neg\varphi \in \mathrm{Th}(\mathcal{M})$. This contradicts our assumption. □

Warning 2.7. Do not confuse the completeness of a *theory* (every sentence or its negation belongs to the theory) with the completeness of a *formal system* or logic (every semantic consequence is provable). These are distinct concepts.

Example 2.8 (A trivially complete theory)**.** Consider the inconsistent theory axiomatised by $\exists x\,(x \neq x)$. This axiomatisation is non-trivial (it consists of a single decidable axiom). The theory has no model and contains every sentence of the language. It is therefore complete, but trivially so.

2.3 Proving incompleteness of an axiomatisation

To show that an axiomatisation Δ is *incomplete* (i.e. that the theory it generates is not complete), it suffices to exhibit a sentence φ that can be neither proved nor refuted from Δ. Concretely:

(i) Find a model of Δ that falsifies φ (showing $\Delta \nvdash \varphi$).

(ii) Find a model of Δ that falsifies $\neg\varphi$ (showing $\Delta \nvdash \neg\varphi$).

Here we use the contrapositive of soundness: $\Delta \not\models \varphi \Rightarrow \Delta \nvdash \varphi$.

2.4 Proving the independence of an axiom

Given a set of axioms Δ and a candidate new axiom φ, we show that φ is **independent** of Δ (i.e. $\Delta \nvdash \varphi$ and $\Delta \nvdash \neg\varphi$) by the same model-theoretic method: find one structure satisfying Δ but not φ, and another satisfying Δ but not $\neg\varphi$.

2.5 Proving consistency

Recall that Γ is **consistent** if and only if $\Gamma \nvdash \bot$.

Proposition 2.9. *If Γ is satisfiable, then Γ is consistent.*

Proof. By the contrapositive of soundness: if $\Gamma \not\models \bot$ then $\Gamma \nvdash \bot$. Now suppose Γ is satisfiable, i.e. some structure $\mathcal{M}$ satisfies Γ. Since no structure satisfies $\bot$, we have $\Gamma \not\models \bot$, and therefore $\Gamma \nvdash \bot$. □

Warning 2.10. The converse statement—that if Γ is consistent then it is satisfiable—has *not* been proved at this stage. It is in fact equivalent to the completeness theorem, and proving it requires substantial work (see Section 6).

3 Examples of First-Order Theories

3.1 Group theory

The language of group theory has signature $\{1, \cdot\}$, where 1 is a constant symbol and $\cdot$ is a binary function symbol. The axioms are:

(G1) $\forall x\, \exists y\, (x \cdot y = 1)$ (inverse)

(G2) $\forall x\, (x \cdot 1 = x)$ (identity)

(G3) $\forall x\, \forall y\, \forall z\, (x \cdot (y \cdot z) = (x \cdot y) \cdot z)$ (associativity)

Group theory is *not* complete. Consider the sentence expressing commutativity:

$$\forall x\, \forall y\, (x \cdot y = y \cdot x).$$

There exist abelian groups (which satisfy commutativity) and non-abelian groups (which falsify it). Both kinds of structure satisfy the group axioms (G1)–(G3), so commutativity can be neither proved nor refuted from those axioms.

Remark 3.1. In algebra, many theories are deliberately incomplete: the aim is to characterise a *class* of structures abstractly, not to pin down a single structure.

3.2 Robinson's arithmetic Q

Robinson's arithmetic Q is formulated in the language $\{0, ', +, \cdot\}$ with the following axioms:

(Q1) $\forall x\, (0 \neq x')$

(Q2) $\forall x\, \forall y\, (x' = y' \to x = y)$

(Q3) $\forall x\, (x \neq 0 \to \exists y\, (x = y'))$

(Q4) $\forall x\, (x + 0 = x)$

(Q5) $\forall x\, \forall y\, (x + y' = (x + y)')$

(Q6) $\forall x\, (x \cdot 0 = 0)$

(Q7) $\forall x\, \forall y\, (x \cdot y' = (x \cdot y) + x)$

Q is not complete. For instance, the sentence $\forall x\, (0 + x = x)$ is independent of the axioms of Q: there exist non-standard models of Q in which the function interpreting '+' does not behave like standard addition, and which provide counter-examples to this sentence.

3.3 Peano arithmetic PA

Peano arithmetic PA extends Robinson's arithmetic by dropping axiom (Q3) and adding an axiom schema of induction. Its axioms are (Q1), (Q2), (Q4)–(Q7) together with every sentence of the form:

$$\big(\varphi(0) \ \wedge\ \forall x\, (\varphi(x) \to \varphi(x'))\big) \ \to\ \forall x\, \varphi(x),$$

where φ is any formula in which x occurs free, and $\varphi(0)$ is the result of replacing those free occurrences by 0.

Theorem 3.2 (Gödel, 1931)**.** PA *is not complete.*

This is far less obvious than the incompleteness of Q, and its proof is deep. If we accept that there is a standard, intended model for arithmetic—the structure $(\mathbb{N}, 0, ', +, \cdot)$—then Gödel's result tells us that $\mathrm{Th}(\mathbb{N}, 0, ', +, \cdot) \neq \mathrm{PA}$. There are arithmetical truths that do not belong to PA and that therefore cannot be derived from its axioms.

This raises a philosophical question: does Gödel's result show that mathematical thought is, in some sense, not mechanisable?

Remark 3.3. Mathematicians typically start with a structure and search for an axiomatisation Σ of its theory. They write down sentences capturing known properties of the structure; if they find a model $\mathcal{M}$ satisfying Σ together with a sentence φ such that $\mathcal{M} \not\models \varphi$, they add φ to Σ as a new axiom. When no further such sentence can be found, they have a plausible candidate for a complete axiomatisation. In the case of Peano arithmetic, Gödel showed that this process can never be completed. Counterintuitively, the theory of the real closed field *is* complete (Tarski), though the proof is involved.

4 Characterising the Size of Structures

Certain properties of structures can be expressed using only the purely logical vocabulary of FOL (i.e. without any non-logical symbols).

4.1 At least n elements

Consider the sentence

$$\exists x\, \exists y\, (x \neq y).$$

A structure $\mathcal{M}$ satisfies this sentence if and only if its domain contains at least two elements. More generally, for each natural number $n \geq 1$, define:

$$\varphi_n \;:\; \exists x_1\, \exists x_2 \cdots \exists x_n \bigwedge_{1 \leq i < j \leq n} (x_i \neq x_j).$$

A structure satisfies φ_n if and only if its domain has at least n elements.

4.2 At most n elements

Consider the sentence

$$\forall x\, \forall y\, \forall z\, (x = y \;\vee\; x = z \;\vee\; y = z).$$

This forces the domain to have at most two elements. In general, define:

$$\psi_n \;:\; \forall x_1\, \forall x_2 \cdots \forall x_{n+1} \bigvee_{1 \leq i < j \leq n+1} (x_i = x_j).$$

A structure satisfies ψ_n if and only if its domain has at most n elements.

4.3 Exactly n elements

Setting $\theta_n = \varphi_n \wedge \psi_n$, a structure satisfies θ_n if and only if its domain has *exactly* n elements.

4.4 Infinite domains

Now consider the set $\Delta = \{\theta_n \mid n \in \mathbb{N}\}$, which contains infinitely many sentences but is decidable (there is an effective procedure to check membership). A structure $\mathcal{M}$ satisfies Δ if and only if the domain of $\mathcal{M}$ is infinite.

Remark 4.1. While we can express "the domain is infinite" by an infinite (but decidable) set of sentences, the property "the domain is finite" *cannot* be expressed in first-order logic. This is a consequence of the compactness theorem, which we do not prove here.

5 Completeness of the Tableaux Method

We now turn to the completeness theorem for the tableaux proof system. Throughout this section, we restrict attention to first-order logic *without identity and without function symbols.* (The extension to identity and function symbols requires additional care in the construction of the term model, but the underlying strategy is the same.)

5.1 Soundness vs. completeness: a contrast in proof strategy

The soundness proof for tableaux is *atomistic*: we show that each individual rule preserves satisfiability. If a branch is satisfiable before a rule is applied, it remains satisfiable (i.e. open) after the rule is applied.

The completeness proof must be *holistic*: we must show that the derivation system *as a whole* is complete. In contrapositive form, completeness says:

$$\Gamma \nvdash \theta \implies \Gamma \not\models \theta,$$

which is equivalent to: if a tableau for $\Gamma, \neg\theta$ remains open, then $\Gamma, \neg\theta$ is satisfiable.

We therefore need to prove that whenever a finished tableau has an open branch, the set of formulas along that branch is satisfiable. By semantic monotonicity (if $\Gamma \subseteq \Delta$ and Δ is satisfiable, then Γ is satisfiable), the satisfiability of the branch entails the satisfiability of the original set $\Gamma \cup \{\neg\theta\}$.

5.2 Strategy

The argument proceeds in two steps:

(1) From the information on a finished open branch, **construct a model** $\mathcal{M}$.

(2) **Show** that $\mathcal{M}$ satisfies every sentence marked as true on the branch and falsifies every sentence marked as false.

5.3 Building a term model

The key idea is simple: treat the individual constants occurring on the branch as the elements of the domain, and define the interpretation so that the true atomic formulas come out true and the false atomic formulas come out false.

Two points deserve emphasis. First, the set of constants on a finished branch may be larger than the set of constants in the original formulas, because the existential-instantiation rule introduces *fresh* witness constants. Second, in the presence of quantifiers, a finished branch may be infinite (the universal-instantiation rule must be re-applied whenever a new constant appears). We therefore allow the domain of the term model to be countably infinite.

Definition 5.1 (Term model induced by a branch)**.** Let B be a finished open branch of a tableau, and let C_B be the set of *all* constants occurring on B (including any witness constants introduced by the existential-instantiation rule). The **term model induced by** B is the structure $\mathcal{M} = \langle D, \mathcal{I} \rangle$ defined as follows:

- The domain of discourse is $D = C_B$.
- Each constant $a \in C_B$ is interpreted as itself: $\mathcal{I}(a) = a$.
- The interpretation function $\mathcal{I}$ is defined on predicates via the atomic sentences on B: for any n-ary predicate R and constants $a_1, \ldots, a_n$,
  - if $T\colon Ra_1 \ldots a_n$ appears on B, then $\langle a_1, \ldots, a_n \rangle \in \mathcal{I}(R)$;
  - if $F\colon Ra_1 \ldots a_n$ appears on B, then $\langle a_1, \ldots, a_n \rangle \notin \mathcal{I}(R)$.

Remark 5.2. This definition is well-defined because B is an open branch: no atomic formula appears on B marked as both true and false (otherwise the branch would be closed). Note also that C_B is non-empty: the tableaux rules for quantifiers guarantee that at least one constant always occurs on a finished branch.

5.4 The inductive proof

Before proceeding, let us recall the four tableaux rules governing quantifiers. In what follows, a denotes any constant already occurring on the branch, and c denotes a *fresh* constant not previously occurring on the branch.

Signed formula	**Rule**
$T\colon \forall x\,\varphi(x)$	Add $T\colon \varphi(a)$ for every constant a on the branch.
$F\colon \forall x\,\varphi(x)$	Add $F\colon \varphi(c)$ for a fresh constant c.
$T\colon \exists x\,\varphi(x)$	Add $T\colon \varphi(c)$ for a fresh constant c.
$F\colon \exists x\,\varphi(x)$	Add $F\colon \varphi(a)$ for every constant a on the branch.

The rules involving fresh constants ($F:\forall$ and $T:\exists$) are applied once, introducing a single witness. The rules involving every constant ($T\colon\forall$ and $F\colon\exists$) must be re-applied whenever a new constant appears on the branch; a branch is *finished* only when these rules have been applied to all available constants.

We prove by induction on the complexity of formulas that the term model $\mathcal{M}$ induced by a finished open branch B satisfies all the formulas on B in the appropriate sense: true-signed formulas are satisfied, and false-signed formulas are not satisfied.

Theorem 5.3 (Branch satisfiability)**.** *Let B be a finished open branch of a tableau, and let $\mathcal{M}$ be the term model induced by B. For every formula φ appearing on B:*
(a) if $T\colon\varphi$ is on B, then $\mathcal{M}\models\varphi$;
(b) if $F\colon\varphi$ is on B, then $\mathcal{M}\not\models\varphi$.

Proof. The proof is by induction on the complexity of φ.

Base case. Suppose φ is an atomic formula $Ra_1\dots a_n$.

- If $T\colon Ra_1\dots a_n$ is on B, then by definition of $\mathcal{M}$ we have $\langle a_1,\dots,a_n\rangle\in\mathcal{I}(R)$, so $\mathcal{M}\models Ra_1\dots a_n$.
- If $F\colon Ra_1\dots a_n$ is on B, then $\langle a_1,\dots,a_n\rangle\notin\mathcal{I}(R)$, so $\mathcal{M}\not\models Ra_1\dots a_n$.

Inductive hypothesis. Suppose the result holds for all formulas of complexity at most n.

Inductive step. Let θ be a formula of complexity $n+1$. Then θ has one of the following forms: $\neg\varphi$, $\varphi\wedge\psi$, $\varphi\vee\psi$, $\varphi\to\psi$, $\forall x\,\varphi(x)$, or $\exists x\,\varphi(x)$. We treat each case in turn.

Case 1: $\theta=\neg\varphi$.

- Suppose $T\colon\neg\varphi$ is on B. Since B is finished, the negation rule has been applied, so $F\colon\varphi$ is also on B. The formula φ has complexity n, so by the inductive hypothesis $\mathcal{M}\not\models\varphi$. Therefore $\mathcal{M}\models\neg\varphi$, i.e. $\mathcal{M}\models\theta$.
- Suppose $F\colon\neg\varphi$ is on B. Since B is finished, $T\colon\varphi$ is on B. By the inductive hypothesis $\mathcal{M}\models\varphi$, so $\mathcal{M}\not\models\neg\varphi$, i.e. $\mathcal{M}\not\models\theta$.

Case 2: $\theta=\varphi\wedge\psi$.

- Suppose $T\colon\varphi\wedge\psi$ is on B. Since B is finished and the conjunction rule for true-signed conjunctions is non-branching, both $T\colon\varphi$ and $T\colon\psi$ are on B. By the inductive hypothesis, $\mathcal{M}\models\varphi$ and $\mathcal{M}\models\psi$. Therefore $\mathcal{M}\models\varphi\wedge\psi$.

- Suppose $F\colon \varphi \wedge \psi$ is on B. Since B is finished and the rule for false-signed conjunctions is a branching rule, at least one of $F\colon \varphi$ or $F\colon \psi$ appears on B (we follow whichever branch B continues along). If $F\colon \varphi$ is on B, then by the inductive hypothesis $\mathcal{M} \not\models \varphi$, so $\mathcal{M} \not\models \varphi \wedge \psi$. If $F\colon \psi$ is on B, then $\mathcal{M} \not\models \psi$, and again $\mathcal{M} \not\models \varphi \wedge \psi$.

Case 3: $\theta = \varphi \vee \psi$.

This case is dual to the conjunction case.

- Suppose $T\colon \varphi \vee \psi$ is on B. The rule for true-signed disjunctions is a branching rule, so at least one of $T\colon \varphi$ or $T\colon \psi$ appears on B. In either case, the inductive hypothesis gives $\mathcal{M} \models \varphi$ or $\mathcal{M} \models \psi$, whence $\mathcal{M} \models \varphi \vee \psi$.
- Suppose $F\colon \varphi \vee \psi$ is on B. The rule is non-branching: both $F\colon \varphi$ and $F\colon \psi$ are on B. By the inductive hypothesis, $\mathcal{M} \not\models \varphi$ and $\mathcal{M} \not\models \psi$, so $\mathcal{M} \not\models \varphi \vee \psi$.

Case 4: $\theta = \varphi \to \psi$.

- Suppose $T\colon \varphi \to \psi$ is on B. The rule for true-signed conditionals is a branching rule: on B, either $F\colon \varphi$ or $T\colon \psi$ appears. If $F\colon \varphi$ is on B, then $\mathcal{M} \not\models \varphi$ and so $\mathcal{M} \models \varphi \to \psi$. If $T\colon \psi$ is on B, then $\mathcal{M} \models \psi$ and so $\mathcal{M} \models \varphi \to \psi$.
- Suppose $F\colon \varphi \to \psi$ is on B. The rule is non-branching: both $T\colon \varphi$ and $F\colon \psi$ are on B. By the inductive hypothesis, $\mathcal{M} \models \varphi$ and $\mathcal{M} \not\models \psi$, so $\mathcal{M} \not\models \varphi \to \psi$.

Case 5: $\theta = \forall x\, \varphi(x)$.

Recall that the tableaux rule for a true-signed universal formula is a non-branching rule that may be applied repeatedly: for each constant a occurring on the branch, the rule adds $T\colon \varphi(a)$ to the branch. In a finished tableau, this rule has been applied for *every* constant on the branch.

- Suppose $T\colon \forall x\, \varphi(x)$ is on B. Since B is finished, for every constant $a \in C_B$ the signed formula $T\colon \varphi(a)$ also appears on B. The formula $\varphi(a)$ has complexity n, so by the inductive hypothesis $\mathcal{M} \models \varphi(a)$ for each $a \in C_B$. Now the domain of $\mathcal{M}$ is precisely C_B, and each constant is interpreted as itself. So $\varphi(x)$ is satisfied by every element of the domain, which means $\mathcal{M} \models \forall x\, \varphi(x)$.
- Suppose $F\colon \forall x\, \varphi(x)$ is on B. The tableaux rule for a false-signed universal formula produces $F\colon \varphi(c)$, where c is a *fresh* constant not previously occurring on the branch. (This is the same rule as for true-signed existentials; falsifying a universal amounts to witnessing an existential.) Since B is finished, $F\colon \varphi(c)$ appears on B for some $c \in C_B$. By the inductive hypothesis, $\mathcal{M} \not\models \varphi(c)$. Since c is in the domain and $\mathcal{I}(c) = c$, there exists an element of the domain that does not satisfy φ. Therefore $\mathcal{M} \not\models \forall x\, \varphi(x)$.

Case 6: $\theta = \exists x\, \varphi(x)$.

This case is dual to the universal case. The tableaux rule for a true-signed existential introduces a fresh witness constant; the rule for a false-signed existential requires instantiation for every constant on the branch.

- Suppose $T\colon \exists x\, \varphi(x)$ is on B. The tableaux rule introduces a fresh constant c not previously occurring on the branch and adds $T\colon \varphi(c)$ to B. Since B is finished, $T\colon \varphi(c)$ appears on B for some $c \in C_B$. By the inductive hypothesis, $\mathcal{M} \models \varphi(c)$. Since c is in the domain and $\mathcal{I}(c) = c$, there exists an element of the domain satisfying φ. Therefore $\mathcal{M} \models \exists x\, \varphi(x)$.
- Suppose $F\colon \exists x\, \varphi(x)$ is on B. Since B is finished, the rule for false-signed existentials has been applied for every constant $a \in C_B$: the signed formula $F\colon \varphi(a)$ appears on B for each such a. By the inductive hypothesis, $\mathcal{M} \not\models \varphi(a)$ for every $a \in C_B$. Since the domain of $\mathcal{M}$ is C_B, no element of the domain satisfies φ. Therefore $\mathcal{M} \not\models \exists x\, \varphi(x)$.

In every case the inductive step goes through, completing the proof. □

Remark 5.4. The asymmetry between the universal and existential rules is worth noting. For the universal quantifier, the *true*-signed rule requires instantiation with every constant (and so may need to be applied repeatedly as new constants appear), while the *false*-signed rule introduces a single fresh witness. For the existential quantifier it is the reverse: the true-signed

rule introduces a witness, and the false-signed rule requires instantiation with every constant. This duality mirrors the semantic equivalence $\neg\forall x\,\varphi(x) \equiv \exists x\,\neg\varphi(x)$.

5.5 The completeness theorem

From Theorem 5.3, we can derive the completeness of the tableaux system.

Theorem 5.5 (Completeness of tableaux)**.** *For any set of sentences Γ and any sentence θ:*

$$\Gamma \models \theta \implies \Gamma \vdash \theta.$$

Proof. We prove the contrapositive: $\Gamma \nvdash \theta \Rightarrow \Gamma \not\models \theta$.

Suppose $\Gamma \nvdash \theta$. Then the tableau for $\Gamma \cup \{\neg\theta\}$ does not close: there exists a finished open branch B. By Theorem 5.3, the set of signed formulas on B is satisfiable: there is a term model $\mathcal{M}$ that satisfies every true-signed formula on B and falsifies every false-signed formula. Since $\Gamma \cup \{\neg\theta\}$ is a subset of the formulas on B, by semantic monotonicity $\mathcal{M}$ satisfies $\Gamma \cup \{\neg\theta\}$. In particular, $\mathcal{M} \models \Gamma$ and $\mathcal{M} \models \neg\theta$, which means $\mathcal{M} \not\models \theta$. Therefore $\Gamma \not\models \theta$. □

6 The Difficult Lemma

We conclude by establishing the relationship between the completeness theorem and a result that is often called the *difficult lemma*.

Lemma 6.1 (Difficult lemma)**.** *For any set of sentences Γ: if Γ is consistent, then Γ is satisfiable.*

This lemma is logically equivalent to the completeness theorem.

Proposition 6.2. *The completeness theorem (Theorem 5.5) and the difficult lemma (Lemma 6.1) are equivalent.*

Proof. We establish both directions.

Completeness ⇒ Difficult lemma. Suppose completeness holds: $\Gamma \models \theta \Rightarrow \Gamma \vdash \theta$. Substituting $\bot$ for θ, we obtain $\Gamma \models \bot \Rightarrow \Gamma \vdash \bot$. Taking the contrapositive: $\Gamma \nvdash \bot \Rightarrow \Gamma \not\models \bot$. That is, if Γ is consistent, then Γ is satisfiable.

Difficult lemma ⇒ Completeness. Suppose the difficult lemma holds: if Γ is consistent, then Γ is satisfiable. Taking the contrapositive: if Γ is not satisfiable, then Γ is inconsistent.

Now suppose $\Gamma \models \theta$. Then $\Gamma \cup \{\neg\theta\} \models \bot$, which means $\Gamma \cup \{\neg\theta\}$ is not satisfiable. By the contrapositive of the difficult lemma, $\Gamma \cup \{\neg\theta\}$ is inconsistent: $\Gamma \cup \{\neg\theta\} \vdash \bot$. By the rule of negation introduction and double-negation elimination, we conclude $\Gamma \vdash \theta$. □

Remark 6.3. The difficult lemma is so named because proving it directly (without appealing to the completeness theorem) requires constructing a model from purely syntactic information—a non-trivial task. The term-model construction we used in Section 5.3 for the tableaux system is one instance of precisely this kind of argument: from the syntactic data on an open branch, we built a structure satisfying the relevant formulas. In the general case (full first-order logic with identity, function symbols, and quantifiers), the construction is considerably more involved but follows the same guiding idea.