Proofs for First-Order Logic

Part 2: Natural Deduction

Advanced Logic Course Notes

Contents

1 Introduction

Natural deduction is a proof system that mirrors the way mathematicians actually reason. Unlike semantic methods (such as truth tables or semantic tableaux) that work by analyzing truth conditions, natural deduction is purely **syntactic**: we manipulate formulas according to fixed rules without ever mentioning interpretations or models.

The key idea is that each logical connective and quantifier has:

- An **introduction rule** that tells us how to *prove* a formula with that connective as its main operator.
- An **elimination rule** that tells us how to *use* a formula with that connective as its main operator.

This document covers the rules for quantifiers ($\forall$ and $\exists$), the rules for identity ($=$), important proof-theoretic concepts, and concludes with the soundness theorem for natural deduction.

2 Rules for Quantifiers

2.1 Existential Introduction ($\exists$I)

The existential introduction rule, also called **existential generalization**, allows us to infer an existentially quantified statement from a specific instance.

Definition 2.1 (Existential Introduction)**.** *If t is a closed term (a term with no free variables, typically an individual constant), then from a formula φt (which contains t), we may infer $\exists x \varphi x$:*

1	φt	
2	$\exists x \varphi x$	$\exists I$, 1

Intuition: If we know that a specific individual t has property φ, then we can conclude that *something* has property φ. For instance, if we know that Socrates is mortal (Ma where a denotes Socrates), we can conclude that something is mortal ($\exists x Mx$).

Example 2.2. *Let us prove that $Aa \to Bb \vdash \exists x \exists y (Ax \to By)$.*

1	$Aa \to Bb$	
2	$\exists y (Aa \to By)$	$\exists I$, 1
3	$\exists x \exists y (Ax \to By)$	$\exists I$, 2

Explanation:

- *Line 1: We start with the premise $Aa \to Bb$.*
- *Line 2: We apply $\exists I$, replacing b with the variable y. The formula $Aa \to Bb$ is an instance of $Aa \to By$ (with y instantiated to b), so we can existentially generalize over b.*
- *Line 3: We apply $\exists I$ again, this time replacing a with x. The formula $\exists y (Aa \to By)$ is an instance of $\exists y (Ax \to By)$ (with x instantiated to a).*

Remark 2.3. *When applying $\exists I$, we are free to choose which occurrences of the term to replace with the quantified variable. We don't have to replace all occurrences.*

2.2 Universal Elimination (∀E)

The universal elimination rule, also called **universal instantiation**, allows us to infer a specific instance from a universally quantified statement.

Definition 2.4 (Universal Elimination)**.** *If t is any closed term, then from $\forall x\varphi x$, we may infer φt:*

$$\begin{array}{r|ll} 1 & \forall x\varphi x & \\ \cline{2-2} 2 & \varphi t & \forall E,\ 1 \end{array}$$

Intuition: If something is true of *everything*, then it is true of any particular thing we care to name. If all humans are mortal, then Socrates (being a human) is mortal.

Example 2.5. *Let us prove that $\forall xPx \vdash \exists xPx$.*

$$\begin{array}{r|ll} 1 & \forall xPx & \\ \cline{2-2} 2 & Pa & \forall E,\ 1 \\ 3 & \exists xPx & \exists I,\ 2 \end{array}$$

Explanation:

- *Line 1: We assume the premise that everything has property P.*
- *Line 2: By ∀E, we instantiate to the specific individual a. Since everything has P, the individual a has P.*
- *Line 3: By ∃I, since a has P, something has P.*

Remark 2.6 (Non-empty domains)**.** *This proof shows that $\forall xPx \vdash \exists xPx$. This might seem surprising: how can we get existence from universality? The key is that in standard first-order logic, we assume the domain is **non-empty**. The proof works because we can always pick some individual a to instantiate to. In logics that allow empty domains (called "free logics"), this inference would not be valid.*

Example 2.7 (The Socrates Syllogism)**.** *Let us prove that the classical syllogism "Every human is mortal. Socrates is human. Therefore Socrates is mortal" is valid.*

Formalization: $\forall x(Hx \to Mx), Ha \vdash Ma$

$$\begin{array}{r|ll} 1 & \forall x(Hx \to Mx) & \\ 2 & Ha & \\ \cline{2-2} 3 & Ha \to Ma & \forall E,\ 1 \\ 4 & Ma & \Rightarrow E,\ 2,\ 3 \end{array}$$

Explanation:

- *Lines 1–2: We state our premises.*
- *Line 3: We apply ∀E to line 1, instantiating x with a (Socrates). This gives us "If Socrates is human, then Socrates is mortal."*
- *Line 4: We apply modus ponens (→E) to lines 2 and 3. Since Socrates is human (line 2) and if Socrates is human then Socrates is mortal (line 3), we conclude that Socrates is mortal.*

2.3 Universal Introduction (∀I)

The universal introduction rule, also called **universal generalization**, is more subtle. To prove that something holds for everything, we cannot check every instance (there are infinitely many). Instead, we prove it for an **arbitrary** individual.

Definition 2.8 (Universal Introduction)**.** *From φa, we may infer $\forall x \varphi x$,* ***provided that*** *the constant a is* ***arbitrary*** *in the proof:*

$$
\begin{array}{r|ll}
1 & \varphi a & \\ \cline{2-2}
2 & \forall x \varphi x & \forall I,\ 1
\end{array}
$$

The constant a is ***arbitrary*** *if and only if it does not occur in:*

1. *the conclusion $\forall x \varphi x$,*
2. *any premise of the proof, nor*
3. *any undischarged assumption in the subproof leading to φa.*

Intuition: If we can prove something about an individual a without assuming anything special about a—that is, using only facts that would apply to any individual—then what we proved must be true of everything.

Practical test: Check whether you could replace every occurrence of a with any other constant (say, b or c) and the proof would still work. If yes, then a is arbitrary.

Example 2.9. *Let us prove that $\neg\exists x(Fx \wedge Gx) \vdash \forall x(Fx \to \neg Gx)$.*

This says: if nothing is both F and G, then everything that is F is not G.

$$
\begin{array}{l|l|l|ll}
1 & \multicolumn{3}{l}{\neg\exists x(Fx \wedge Gx)} & \\ \cline{2-4}
2 & & \multicolumn{2}{l}{Fa} & \\ \cline{3-3}
3 & & & Ga & \\ \cline{4-4}
4 & & & Fa \wedge Ga & \wedge I,\ 2,\ 3 \\
5 & & & \exists x(Fx \wedge Gx) & \exists I,\ 4 \\
6 & & & \bot & \neg I,\ 1,\ 5 \\
3-67 & & \multicolumn{2}{l}{\neg Ga} & \\
8 & \multicolumn{3}{l}{Fa \to \neg Ga} & \Rightarrow I,\ 2\text{–}7 \\
9 & \multicolumn{3}{l}{\forall x(Fx \to \neg Gx)} & \forall I,\ 8
\end{array}
$$

Explanation:

- *Line 1: Our premise states that nothing is both F and G.*
- *Lines 2–7: We open a subproof assuming Fa (to prove $Fa \to \neg Ga$).*
- *Lines 3–6: Within that, we open another subproof assuming Ga (to derive a contradiction, thereby proving $\neg Ga$).*
- *Line 4: From Fa and Ga, we derive $Fa \wedge Ga$.*
- *Line 5: By $\exists I$, we get $\exists x(Fx \wedge Gx)$.*
- *Line 6: This contradicts line 1, so we have $\bot$.*

- *Line 7: Discharging the assumption Ga, we conclude $\neg Ga$.*
- *Line 8: Discharging the assumption Fa, we conclude $Fa \to \neg Ga$.*
- *Line 9: Since a is arbitrary (it doesn't appear in line 1 or line 9), we generalize to $\forall x(Fx \to \neg Gx)$.*

***Why is** a **arbitrary?** The constant a does not occur in the premise (line 1) or in the conclusion (line 9). All assumptions involving a (lines 2 and 3) were discharged before we applied $\forall I$. We could replace a with any other constant throughout the proof, and it would still work.*

2.4 Existential Elimination (∃E)

The existential elimination rule is the trickiest of the quantifier rules. The core idea is that an existentially quantified sentence $\exists x \varphi x$ can be thought of as an infinite disjunction: $\varphi a \vee \varphi b \vee \varphi c \vee \cdots$ over all individuals.

To use such a disjunction, we need to show that no matter which disjunct is true, we can derive the same conclusion. This is analogous to proof by cases (∨E).

Definition 2.10 (Existential Elimination)**.** *If we have $\exists x \varphi x$, and we can derive ψ from the assumption φa (where a is arbitrary), then we may conclude ψ:*

1	$\exists x \varphi x$	
2	$\quad \varphi a$	
3	$\quad \psi$	
4	ψ	$\exists E$, 1, 2–3

*The constant a must be **arbitrary**: it must not occur in:*

1. *the conclusion ψ,*
2. *any premise,*
3. *$\exists x \varphi x$ itself, nor*
4. *any undischarged assumption other than φa in the subproof.*

Intuition: We know that *something* satisfies φ. We don't know which thing, so we give it a temporary name a (the "witness"). We then reason about this a, but we must not assume anything special about a—we can only use what we know: that it satisfies φ. If we can derive ψ without a appearing in ψ, then ψ must be true regardless of which specific individual the witness turned out to be.

Example 2.11. *Let us prove that $\exists x(Fx \wedge Gx) \vdash \exists x Fx$.*

1	$\exists x(Fx \wedge Gx)$	
2	$\quad Fa \wedge Ga$	
3	$\quad Fa$	$\wedge E$, 2
4	$\quad \exists x Fx$	$\exists I$, 3
5	$\exists x Fx$	$\exists E$, 1, 2–4

Explanation:

- *Line 1: Our premise says something is both F and G.*

- *Line 2: We assume $Fa \wedge Ga$ for an arbitrary a. This is our "witness" assumption.*
- *Line 3: From the conjunction, we extract Fa.*
- *Line 4: By ∃I, since a is F, something is F.*
- *Line 5: We apply ∃E. Since we derived $\exists x Fx$ (which doesn't contain a) from the assumption that a witnesses the existential, we can conclude $\exists x Fx$ outside the subproof.*

Why is a arbitrary? *The constant a does not occur in the premise (line 1) or in the conclusion (line 5). It was introduced fresh in line 2 and only appears within the subproof.*

2.5 Why the Restrictions Matter: Invalid "Proofs"

The restrictions on ∀I and ∃E are essential for soundness. Without them, we could "prove" invalid arguments. Let's examine two instructive failures.

Example 2.12 (Invalid Argument 1)**.** *Consider: "Someone is walking. Socrates is whistling. Therefore, someone is both walking and whistling."*

Formalization: $\exists x Fx, Ga \nvdash \exists x(Fx \wedge Gx)$

*This argument is **invalid**—the person walking might not be Socrates! Here is a fallacious "proof":*

1	$\exists x Fx$	
2	Ga	
3	$\quad Fa$	
4	$\quad Fa \wedge Ga$	∧I, 3, 2
5	$\quad \exists x(Fx \wedge Gx)$	∃I, 4
6	$\exists x(Fx \wedge Gx)$	∃E, 1, 3–5

What went wrong? *At line 3, we used the constant a for our witness. But a already appears in premise 2! This violates the requirement that the witness constant be arbitrary. The constant a names a specific individual (Socrates, who whistles), so we cannot use it as an arbitrary witness for "someone walks."*

Test: *If we replaced a with a fresh constant b, the proof would break: we couldn't derive $Fb \wedge Gb$ because we only have Ga, not Gb.*

Example 2.13 (Invalid Argument 2)**.** *Consider: "Someone is walking. Everyone who walks also whistles. Therefore, Socrates whistles."*

Formalization: $\exists x Fx, \forall x(Fx \to Gx) \nvdash Ga$

*This argument is **invalid**—the person walking might not be Socrates! Here is a fallacious "proof":*

1	$\exists x Fx$	
2	$\forall x(Fx \to Gx)$	
3	$\quad Fa$	
4	$\quad Fa \to Ga$	∀E, 2
5	$\quad Ga$	⇒E, 3, 4
6	Ga	∃E, 1, 3–5

What went wrong? *At line 6, we try to apply* $\exists E$*. But the conclusion* Ga *contains the constant* a*! This violates the requirement that* a *not appear in the conclusion. We were reasoning about an arbitrary walker and concluded something specifically about* a *(Socrates).*

Moral: *When using* $\exists E$*, the witness constant must be completely "forgotten" by the time we exit the subproof. It cannot appear in the final conclusion.*

3 Rules for Identity

First-order logic with identity includes a special binary predicate $=$ with its own rules.

3.1 Identity Introduction (=I)

Definition 3.1 (Identity Introduction)**.** *For any closed term* t*, we may introduce* $t = t$ *at any point in a proof:*

$$\begin{array}{r|ll} 1 & t = t & \quad =I \end{array}$$

Intuition: Everything is identical to itself. This is the reflexivity of identity and requires no premises or assumptions.

3.2 Identity Elimination (=E)

Definition 3.2 (Identity Elimination—Leibniz's Law)**.** *If* $t_1 = t_2$ *and* φt_1*, then we may infer* φt_2 *(and vice versa):*

$$\begin{array}{r|ll} 1 & t_1 = t_2 & \\ 2 & \underline{\varphi t_1} & \\ 3 & \varphi t_2 & \quad =E,\ 1,\ 2 \end{array} \qquad \textit{or} \qquad \begin{array}{r|ll} 1 & t_1 = t_2 & \\ 2 & \underline{\varphi t_2} & \\ 3 & \varphi t_1 & \quad =E,\ 1,\ 2 \end{array}$$

Intuition: If two terms denote the same individual, then whatever is true of one is true of the other. This is Leibniz's principle of the *indiscernibility of identicals*.

Remark 3.3 (Intensional Contexts)**.** *Leibniz's Law applies to the extensional language of first-order logic. It does not always work in natural language, especially with propositional attitude verbs like "believes," "knows," or "wants."*

For example: Maya believes that Eric Rohmer directed Ma nuit chez Maud*. Eric Rohmer = Maurice Schérer (his real name). But it does not follow that Maya believes that Maurice Schérer directed* Ma nuit chez Maud*—she might not know they are the same person!*

Such contexts are called ***intensional*** *or* ***opaque****, and they require more sophisticated logical frameworks to handle properly.*

Exercise 3.1. *Prove the following theorems of first-order logic with identity:*

1. $\vdash \forall x(x = x)$ *(Reflexivity)*
2. $\vdash \forall x \forall y(x = y \to y = x)$ *(Symmetry)*
3. $\vdash \forall x \forall y \forall z((x = y \land y = z) \to x = z)$ *(Transitivity)*

Hint: *Prove these with arbitrary constants first, then generalize step by step.*

4 Proof-Theoretic Concepts

4.1 Theorems and Derivability

Definition 4.1 (Theorem). *A sentence φ is a* ***theorem****, written $\vdash \varphi$, if and only if there is a derivation of φ from an empty set of premises in which all assumptions are discharged.*

A theorem is a logical truth that can be proven without any premises—it follows from the rules of logic alone.

Definition 4.2 (Derivability). *A sentence φ is* ***derivable from*** *a set of sentences Γ, written $\Gamma \vdash \varphi$, if and only if there is a derivation with conclusion φ in which every undischarged assumption is an element of Γ.*

We write $\Gamma \nvdash \varphi$ to mean that φ is not derivable from Γ.

Definition 4.3 (Consistency). *A set of sentences Γ is* ***inconsistent*** *if and only if $\Gamma \vdash \bot$ (a contradiction can be derived from Γ).*

A set is ***consistent*** *if and only if it is not inconsistent, i.e., $\Gamma \nvdash \bot$.*

4.2 Properties of Derivability

Proposition 4.4 (Reflexivity). *If $\varphi \in \Gamma$, then $\Gamma \vdash \varphi$.*

Proof. Consider the one-line "proof" that simply states φ as an assumption. Since $\varphi \in \Gamma$, this undischarged assumption is in Γ, making it a valid derivation of φ from Γ. □

Proposition 4.5 (Monotonicity). *If $\Gamma \subseteq \Delta$ and $\Gamma \vdash \varphi$, then $\Delta \vdash \varphi$.*

Proof. Any derivation of φ from Γ has all its undischarged assumptions in Γ. Since $\Gamma \subseteq \Delta$, these assumptions are also in Δ. Thus the same derivation serves as a derivation of φ from Δ. □

Proposition 4.6 (Transitivity / Cut). *If $\Gamma \vdash \varphi$ and $\Delta \cup \{\varphi\} \vdash \psi$, then $\Gamma \cup \Delta \vdash \psi$.*

Proof. We construct a derivation as follows:

1. From $\Delta \cup \{\varphi\} \vdash \psi$, we can treat φ as an assumption and discharge it using $\to$I to get $\Delta \vdash \varphi \to \psi$.
2. We have $\Gamma \vdash \varphi$.
3. By modus ponens ($\to$E) on these two results, we get $\Gamma \cup \Delta \vdash \psi$.

All undischarged assumptions come from either Γ or Δ, hence from $\Gamma \cup \Delta$. □

4.3 Characterizations of Inconsistency

Proposition 4.7. *The following are equivalent:*

1. *Γ is inconsistent (i.e., $\Gamma \vdash \bot$).*
2. *For any sentence φ: $\Gamma \vdash \varphi$.*
3. *For some sentence φ: $\Gamma \vdash \varphi$ and $\Gamma \vdash \neg\varphi$.*

Proof. (1) $\Rightarrow$ (2): From $\bot$, anything follows by the $\bot$-elimination rule (ex falso quodlibet).

(2) $\Rightarrow$ (3): Trivial—pick any φ and its negation.

(3) $\Rightarrow$ (1): From φ and $\neg\varphi$, we can derive $\bot$ by $\neg$-elimination. □

Proposition 4.8 (Derivability and Inconsistency). *$\Gamma \cup \{\neg\varphi\}$ is inconsistent if and only if $\Gamma \vdash \varphi$.*

Proof. ($\Rightarrow$) Assume $\Gamma \cup \{\neg\varphi\}$ is inconsistent. Then:

1. $\Gamma \cup \{\neg\varphi\} \vdash \bot$
2. By $\to$I (discharging $\neg\varphi$): $\Gamma \vdash \neg\varphi \to \bot$, i.e., $\Gamma \vdash \neg\neg\varphi$
3. By double negation elimination: $\Gamma \vdash \varphi$

($\Leftarrow$) Assume $\Gamma \vdash \varphi$. Then in $\Gamma \cup \{\neg\varphi\}$:

1. $\neg\varphi$ is derivable (by reflexivity, since $\neg\varphi \in \Gamma \cup \{\neg\varphi\}$)
2. φ is derivable (by monotonicity, since $\Gamma \subseteq \Gamma \cup \{\neg\varphi\}$)
3. Therefore $\bot$ is derivable (by $\neg$E)

□

4.4 Compactness

Theorem 4.9 (Syntactic Compactness). *If $\Gamma \vdash \varphi$, then there is a finite subset $\Gamma_0 \subseteq \Gamma$ such that $\Gamma_0 \vdash \varphi$.*

Proof. If $\Gamma \vdash \varphi$, then there exists a derivation from Γ to φ. A derivation is a finite sequence of lines. Each undischarged assumption appears on some line, so the set of undischarged assumptions is finite. Let Γ_0 be this finite set. Then $\Gamma_0 \subseteq \Gamma$ and $\Gamma_0 \vdash \varphi$. □

Corollary 4.10. *If no finite subset $\Gamma_0 \subseteq \Gamma$ is such that $\Gamma_0 \vdash \varphi$, then $\Gamma \nvdash \varphi$.*

Corollary 4.11. *If no finite subset of Γ is inconsistent, then Γ is consistent.*

Proof. Contrapositive of the compactness theorem with $\varphi = \bot$. □

5 The Soundness Theorem for Natural Deduction

The soundness theorem establishes that our proof system is trustworthy: if we can derive something, then it is semantically valid.

Theorem 5.1 (Soundness). *If $\Gamma \vdash \varphi$, then $\Gamma \vDash \varphi$.*

In words: if there is a syntactic derivation of φ from Γ, then φ is a semantic consequence of Γ (every model that satisfies all formulas in Γ also satisfies φ).

5.1 Philosophical Significance

The soundness theorem has profound philosophical implications:

- **Formal computations preserve truth.** When we manipulate symbols according to syntactic rules, we never go from true premises to a false conclusion.
- **Syntax mirrors semantics.** The rules of natural deduction, designed to capture "natural" reasoning patterns, turn out to respect the truth conditions of the logical connectives.
- **Mechanization of reasoning.** Since syntactic derivation is mechanical (we just follow rules), soundness suggests that rational thought can, in principle, be simulated by computations. This is connected to Jerry Fodor's "Language of Thought" hypothesis: perhaps human cognition operates as a semantic engine driven by syntax.

5.2 Proof Strategy

We prove soundness by **strong induction on the length of proofs**.

Notation: For a proof, let Γ_i denote the set of undischarged assumptions at line i, and let φ_i denote the formula on line i.

Goal: Show that for every line i of any proof, $\Gamma_i \vDash \varphi_i$.

Method:

1. **Base case:** Show that for one-line proofs, $\Gamma_1 \vDash \varphi_1$.

2. **Inductive step:** Assume that for all lines $k < n$, we have $\Gamma_k \vDash \varphi_k$. Show that $\Gamma_n \vDash \varphi_n$, by considering which rule was used to derive line n.

We need to verify soundness for each rule. There are 15 rules without identity (or 18 with identity).

5.3 Preliminary Lemmas

Lemma 5.2 (Semantic Monotonicity)**.** *If $\Gamma_i \subseteq \Gamma_j$ and $\Gamma_i \vDash \varphi$, then $\Gamma_j \vDash \varphi$.*

Proof. Suppose $\Gamma_j \nvDash \varphi$. Then there exists a model $\mathcal{M}$ such that $\mathcal{M} \vDash \gamma$ for all $\gamma \in \Gamma_j$, but $\mathcal{M} \nvDash \varphi$. Since $\Gamma_i \subseteq \Gamma_j$, we also have $\mathcal{M} \vDash \gamma$ for all $\gamma \in \Gamma_i$. But then $\Gamma_i \nvDash \varphi$, contradicting our assumption. Hence $\Gamma_j \vDash \varphi$. □

Lemma 5.3 (Deduction Lemma—Semantic Version)**.** *If $\Gamma \cup \{\varphi\} \vDash \psi$, then $\Gamma \vDash \varphi \to \psi$.*

Proof. Suppose $\Gamma \nvDash \varphi \to \psi$. Then there exists a model $\mathcal{M}$ such that $\mathcal{M} \vDash \gamma$ for all $\gamma \in \Gamma$, but $\mathcal{M} \nvDash \varphi \to \psi$. By the semantics of $\to$: $\mathcal{M} \vDash \varphi$ and $\mathcal{M} \nvDash \psi$. But then $\mathcal{M}$ satisfies everything in $\Gamma \cup \{\varphi\}$ yet $\mathcal{M} \nvDash \psi$, contradicting $\Gamma \cup \{\varphi\} \vDash \psi$. □

5.4 Base Case: One-Line Proofs

In a one-line proof, the only line is an undischarged assumption. So φ_1 is itself an assumption, and $\Gamma_1 = \{\varphi_1\}$.

We need to show: $\{\varphi\} \vDash \varphi$.

This is trivially true: any model that satisfies φ (obviously) satisfies φ.

Conclusion: $\Gamma_1 \vdash \varphi_1 \Rightarrow \Gamma_1 \vDash \varphi_1$. ✓

5.5 Inductive Step

Inductive Hypothesis: For all lines $k < n$, $\Gamma_k \vdash \varphi_k \Rightarrow \Gamma_k \vDash \varphi_k$.

We must show $\Gamma_n \vDash \varphi_n$ by considering each possible rule that could produce line n.

5.5.1 Repetition (Reiteration)

If line n is obtained by repeating a formula φ_k from line $k < n$:

By the inductive hypothesis, $\Gamma_k \vDash \varphi_k$. Since we're just repeating the formula, $\varphi_n = \varphi_k$ and $\Gamma_k \subseteq \Gamma_n$. By semantic monotonicity, $\Gamma_n \vDash \varphi_n$. ✓

5.5.2 Conjunction Introduction ($\wedge$I)

Line n has the form $\varphi_k \wedge \varphi_l$, derived from lines k and l (both $< n$).

By the inductive hypothesis: $\Gamma_k \vDash \varphi_k$ and $\Gamma_l \vDash \varphi_l$.

Since $\Gamma_k \subseteq \Gamma_n$ and $\Gamma_l \subseteq \Gamma_n$, by monotonicity: $\Gamma_n \vDash \varphi_k$ and $\Gamma_n \vDash \varphi_l$.

By the semantics of $\wedge$: A model satisfies $\varphi_k \wedge \varphi_l$ iff it satisfies both φ_k and φ_l.

Therefore $\Gamma_n \vDash \varphi_k \wedge \varphi_l$. ✓

5.5.3 Conjunction Elimination (∧E)

Line n has the form φ (either the left or right conjunct), derived from line k where $\varphi_k = \psi \wedge \varphi$ or $\varphi_k = \varphi \wedge \psi$.

By the inductive hypothesis: $\Gamma_k \vDash \varphi_k$.

By monotonicity (since $\Gamma_k \subseteq \Gamma_n$): $\Gamma_n \vDash \varphi_k$.

By the semantics of $\wedge$: if a model satisfies $\psi \wedge \varphi$, it satisfies φ.

Therefore $\Gamma_n \vDash \varphi$. ✓

5.5.4 Implication Introduction (→I)

Line n has the form $\varphi_l \to \varphi_k$, where lines l through k form a subproof:

- Line l is the assumption φ_l.
- Line k is the conclusion φ_k of the subproof.
- Line n discharges the assumption.

We have $\Gamma_l \subseteq \Gamma_k \subseteq \Gamma_n \cup \{\varphi_l\}$ (the undischarged assumptions in the subproof are those at line n plus the assumption φ_l).

By the inductive hypothesis: $\Gamma_k \vDash \varphi_k$.

By monotonicity: $\Gamma_n \cup \{\varphi_l\} \vDash \varphi_k$.

By Lemma 2 (Deduction Lemma): $\Gamma_n \vDash \varphi_l \to \varphi_k$. ✓

5.5.5 Implication Elimination / Modus Ponens (→E)

Line n has the form ψ, derived from lines k and l where $\varphi_k = \varphi$ and $\varphi_l = \varphi \to \psi$.

By the inductive hypothesis: $\Gamma_k \vDash \varphi$ and $\Gamma_l \vDash \varphi \to \psi$.

By monotonicity: $\Gamma_n \vDash \varphi$ and $\Gamma_n \vDash \varphi \to \psi$.

By the semantics of $\to$: if a model satisfies both φ and $\varphi \to \psi$, it satisfies ψ.

Therefore $\Gamma_n \vDash \psi$. ✓

5.5.6 Disjunction Introduction (∨I)

Line n has the form $\varphi \vee \psi$, derived from line k where $\varphi_k = \varphi$ (or $\varphi_k = \psi$).

By the inductive hypothesis: $\Gamma_k \vDash \varphi$.

By monotonicity: $\Gamma_n \vDash \varphi$.

By the semantics of $\vee$: if a model satisfies φ, it satisfies $\varphi \vee \psi$.

Therefore $\Gamma_n \vDash \varphi \vee \psi$. ✓

5.5.7 Disjunction Elimination (∨E)

Line n has the form χ, derived from:

- Line m: $\varphi \vee \psi$
- Subproof from j to k: assuming φ, derive χ
- Subproof from j' to k': assuming ψ, derive χ

By the inductive hypothesis: $\Gamma_m \vDash \varphi \vee \psi$, $\Gamma_k \vDash \chi$, and $\Gamma_{k'} \vDash \chi$.

We have $\Gamma_k \subseteq \Gamma_n \cup \{\varphi\}$ and $\Gamma_{k'} \subseteq \Gamma_n \cup \{\psi\}$.

By monotonicity: $\Gamma_n \cup \{\varphi\} \vDash \chi$ and $\Gamma_n \cup \{\psi\} \vDash \chi$.

By the semantics of $\vee$: any model satisfying $\varphi \vee \psi$ satisfies at least one of φ or ψ.

Combined with monotonicity for $\Gamma_m \subseteq \Gamma_n$: $\Gamma_n \vDash \chi$. ✓

5.5.8 Negation Introduction ($\neg$I)

Line n has the form $\neg\varphi$, derived from a subproof where assuming φ leads to $\bot$.

By the inductive hypothesis and the structure similar to $\to$I: $\Gamma_n \cup \{\varphi\} \vDash \bot$.

But no model can satisfy $\bot$. So there is no model satisfying $\Gamma_n \cup \{\varphi\}$.

Therefore, every model satisfying Γ_n must make φ false, i.e., $\Gamma_n \vDash \neg\varphi$. ✓

5.5.9 Negation Elimination ($\neg$E)

Line n has $\bot$, derived from lines k and l where $\varphi_k = \varphi$ and $\varphi_l = \neg\varphi$.

By the inductive hypothesis: $\Gamma_k \vDash \varphi$ and $\Gamma_l \vDash \neg\varphi$.

By monotonicity: $\Gamma_n \vDash \varphi$ and $\Gamma_n \vDash \neg\varphi$.

By the semantics of $\neg$: no model can satisfy both φ and $\neg\varphi$.

Therefore, there is no model satisfying Γ_n, which means $\Gamma_n \vDash \bot$ vacuously. ✓

5.5.10 Falsum Elimination ($\bot$E) / Ex Falso Quodlibet

Line n has any formula φ, derived from line k where $\varphi_k = \bot$.

By the inductive hypothesis: $\Gamma_k \vDash \bot$.

By monotonicity: $\Gamma_n \vDash \bot$.

Since no model satisfies $\bot$, there is no model satisfying Γ_n.

Therefore $\Gamma_n \vDash \varphi$ holds vacuously (there are no counterexamples). ✓

5.5.11 Double Negation Elimination (DNE)

Line n has φ, derived from line k where $\varphi_k = \neg\neg\varphi$.

By the inductive hypothesis: $\Gamma_k \vDash \neg\neg\varphi$.

By monotonicity: $\Gamma_n \vDash \neg\neg\varphi$.

By the semantics of $\neg$: $\mathcal{M} \vDash \neg\neg\varphi$ iff $\mathcal{M} \nvDash \neg\varphi$ iff $\mathcal{M} \vDash \varphi$.

Therefore $\Gamma_n \vDash \varphi$. ✓

5.5.12 Universal Elimination ($\forall$E)

Line n has φt, derived from line k where $\varphi_k = \forall x \varphi x$.

By the inductive hypothesis: $\Gamma_k \vDash \forall x \varphi x$.

By monotonicity: $\Gamma_n \vDash \forall x \varphi x$.

By the semantics of $\forall$: if $\mathcal{M} \vDash \forall x \varphi x$, then for every individual d in the domain, $\mathcal{M} \vDash \varphi x$ under an assignment where $x \mapsto d$. In particular, taking $d = I(t)$ (the denotation of t), we get $\mathcal{M} \vDash \varphi t$.

Therefore $\Gamma_n \vDash \varphi t$. ✓

5.5.13 Universal Introduction ($\forall$I)

Line n has $\forall x \varphi x$, derived from line l where $\varphi_l = \varphi a$ and a is arbitrary.

By the inductive hypothesis: $\Gamma_l \vDash \varphi a$.

Since $\Gamma_l \subseteq \Gamma_n$ and a does not occur in Γ_n: $\Gamma_n \vDash \varphi a$.

Key argument: Let $\mathcal{M}$ be any model satisfying Γ_n. We want to show $\mathcal{M} \vDash \forall x \varphi x$.

Let α be any individual in the domain of $\mathcal{M}$.

Consider a model $\mathcal{M}'$ that is identical to $\mathcal{M}$ except that $I(a) = \alpha$.

Since a does not occur in any formula in Γ_n, we have $\mathcal{M}' \vDash \Gamma_n$ (changing the interpretation of a doesn't affect the truth of formulas not containing a).

Since $\Gamma_n \vDash \varphi a$, we have $\mathcal{M}' \vDash \varphi a$.

This means that under the assignment $\sigma[x \mapsto \alpha]$, $\mathcal{M}' \vDash \varphi x$, i.e., $\mathcal{M}' \vDash_{\sigma[x \mapsto \alpha]} \varphi x$.

Since a does not occur in $\forall x \varphi x$, and $\mathcal{M}'$ agrees with $\mathcal{M}$ on everything except a: $\mathcal{M} \vDash_{\sigma[x \mapsto \alpha]} \varphi x$.

Since α was arbitrary: $\mathcal{M} \vDash \forall x \varphi x$.

Since $\mathcal{M}$ was any model satisfying Γ_n: $\Gamma_n \vDash \forall x \varphi x$. ✓

5.5.14 Existential Introduction (∃I)

Line n has $\exists x \varphi x$, derived from line k where $\varphi_k = \varphi t$ for some closed term t.

By the inductive hypothesis: $\Gamma_k \vDash \varphi t$.

By monotonicity: $\Gamma_n \vDash \varphi t$.

By the semantics of $\exists$: if $\mathcal{M} \vDash \varphi t$, then there exists an individual (namely, $I(t)$) satisfying φx, so $\mathcal{M} \vDash \exists x \varphi x$.

Therefore $\Gamma_n \vDash \exists x \varphi x$. ✓

5.5.15 Existential Elimination (∃E)

Line n has ψ, derived from:

- Line m: $\exists x \varphi x$
- Subproof from j to k: assuming φa (with a arbitrary), derive ψ

By the inductive hypothesis: $\Gamma_m \vDash \exists x \varphi x$ and $\Gamma_k \vDash \psi$.

We have $\Gamma_k \subseteq \Gamma_n \cup \{\varphi a\}$, so by monotonicity: $\Gamma_n \cup \{\varphi a\} \vDash \psi$.

Key argument: Let $\mathcal{M}$ be any model satisfying Γ_n. We want to show $\mathcal{M} \vDash \psi$.

By monotonicity from $\Gamma_m \subseteq \Gamma_n$: $\mathcal{M} \vDash \exists x \varphi x$.

So there exists an individual α in the domain such that $\mathcal{M} \vDash_{\sigma[x \mapsto \alpha]} \varphi x$.

Consider a model $\mathcal{M}'$ identical to $\mathcal{M}$ except $I(a) = \alpha$.

Then $\mathcal{M}' \vDash \varphi a$.

Since a does not occur in Γ_n: $\mathcal{M}' \vDash \Gamma_n$.

So $\mathcal{M}'$ satisfies $\Gamma_n \cup \{\varphi a\}$.

Since $\Gamma_n \cup \{\varphi a\} \vDash \psi$: $\mathcal{M}' \vDash \psi$.

Since a does not occur in ψ, and $\mathcal{M}'$ agrees with $\mathcal{M}$ on everything except a: $\mathcal{M} \vDash \psi$.

Therefore $\Gamma_n \vDash \psi$. ✓

5.5.16 Identity Introduction (=I)

Line n has $t = t$ for some closed term t.

By the semantics of $=$: in any model, $I(t) = I(t)$ (reflexivity of identity in the metatheory).

Therefore any model satisfies $t = t$, regardless of what it satisfies otherwise.

Hence $\Gamma_n \vDash t = t$. ✓

5.5.17 Identity Elimination (=E)

Line n has φt_2 (or φt_1), derived from lines k and l where $\varphi_k = t_1 = t_2$ and $\varphi_l = \varphi t_1$ (or φt_2).

By the inductive hypothesis: $\Gamma_k \vDash t_1 = t_2$ and $\Gamma_l \vDash \varphi t_1$.

By monotonicity: $\Gamma_n \vDash t_1 = t_2$ and $\Gamma_n \vDash \varphi t_1$.

By the semantics of $=$: if $\mathcal{M} \vDash t_1 = t_2$, then $I(t_1) = I(t_2)$.

By Leibniz's Law in the semantics: if $\mathcal{M} \vDash \varphi t_1$ and $I(t_1) = I(t_2)$, then $\mathcal{M} \vDash \varphi t_2$.

Therefore $\Gamma_n \vDash \varphi t_2$. ✓

5.6 Conclusion of the Soundness Proof

We have verified that for every rule of natural deduction, if the premises of the rule yield semantic consequences (by the inductive hypothesis), then the conclusion also yields a semantic consequence.

By strong induction on the length of proofs, every line n of any proof satisfies: if $\Gamma_n \vdash \varphi_n$, then $\Gamma_n \vDash \varphi_n$.

In particular, for the last line of a proof (which is the derived conclusion), we have:

$$\Gamma \vdash \varphi \quad \Rightarrow \quad \Gamma \vDash \varphi$$

This completes the proof of the soundness theorem. $\square$

6 Exercises

Exercise 6.1. *Prove the following using natural deduction:*

1. $\forall x(Fx \to Gx), Fa \vdash \exists x(\neg Gx \to Hx)$
2. $\forall x(Fx \to Gx), \forall x Fx \vdash \forall x Gx$
3. $\exists x(Px \land Qx) \vdash \exists x Px$
4. $\exists x(Px \lor Qx) \vdash \exists x Px \lor \exists x Qx$
5. $\forall x(Bx \to Ax), \exists x(Bx \land Cx) \vdash \exists x(Ax \land Cx)$
6. $\neg \exists x Px \vdash \forall x \neg Px$
7. $\neg \forall x Px \vdash \exists x \neg Px$ *(Hint: use proof by contradiction)*

Exercise 6.2. *Explain why each of the following arguments is **invalid**, and identify what restriction would be violated if we attempted a proof:*

1. $\exists x Fx \vdash Fa$
2. $\exists x Fx, \exists x Gx \vdash \exists x(Fx \land Gx)$
3. $Pa \vdash \forall x Px$

Exercise 6.3. *Prove the identity theorems:*

1. $\vdash \forall x(x = x)$
2. $\vdash \forall x \forall y(x = y \to y = x)$
3. $\vdash \forall x \forall y \forall z((x = y \land y = z) \to x = z)$

Exercise 6.4. *Complete the soundness proof by providing full details for:*

1. *Disjunction elimination ($\lor$E)*
2. *Existential elimination ($\exists$E)*