Soundness of Modal Tableaux

Lecture Notes for M1 — Philosophy of Science and Cognitive Science

Contents

Part I
Soundness

1 What Soundness Says

Theorem 1.1 (Soundness of the tableau method for K)**.** *If $\Gamma \vdash \varphi$ (i.e. every tableau for $\Gamma \cup \{\neg\varphi\}$ closes), then $\Gamma \models \varphi$ (i.e. φ is a semantic consequence of Γ).*

In other words: *the tableau method never proves something false.* If the method says "φ follows from Γ," then indeed every Kripke model that satisfies all of Γ also satisfies φ.

2 The Proof Strategy: Working with the Contrapositive

To prove soundness, it is more convenient to work with the **contrapositive**:

> If $\Gamma \vdash \varphi$, then $\Gamma \models \varphi$.
>
> **Contrapositive**: If $\Gamma \not\models \varphi$, then $\Gamma \nvdash \varphi$.

Let us unpack this.

$\Gamma \not\models \varphi$ means there exists a model in which all formulas in Γ are true but φ is false. In other words, $\Gamma \cup \{\neg\varphi\}$ is *satisfiable*.

$\Gamma \nvdash \varphi$ means the tableau built from $\Gamma \cup \{\neg\varphi\}$ does *not* close: it stays open.

So, to prove soundness, we need to show:

> **Key Claim**: If $\Gamma \cup \{\neg\varphi\}$ is satisfiable, then the tableau we build from
>
> $$\mathsf{T} : \gamma_1, \ldots, \mathsf{T} : \gamma_n, \mathsf{F} : \varphi$$
>
> (where $\Gamma = \{\gamma_1, \ldots, \gamma_n\}$) stays open when we apply the tableau rules.

If we prove this claim, we will have proved the soundness theorem for tableaux.

3 The Proof Method: Preserving Satisfiability

A tableau is built by repeatedly applying rules to extend branches. The core insight is:

> **Preservation Principle**: If a branch is satisfiable before a rule application, then at least one resulting branch is still satisfiable afterwards.

Why does this suffice? By induction on the number of rule applications: if the root is satisfiable, then at every stage of construction, *some* branch remains satisfiable. A satisfiable branch cannot be closed (it cannot contain both $i : \mathsf{T}\ p$ and $i : \mathsf{F}\ p$ for the same atom p at the same world i, since no model makes p both true and false). Therefore the tableau stays open.

There are two kinds of rules, and the Preservation Principle takes a different form for each:

Branching rules split a branch into two sub-branches. We need: if the branch is satisfiable, *at least one* of the sub-branches is satisfiable.

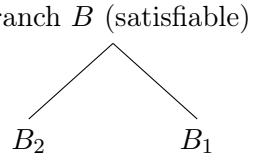

We need: at least one of B_1, B_2 is satisfiable.

Non-branching rules extend a branch by adding formulas (and possibly new world labels) without splitting. We need the stronger condition: the extended branch *itself* is satisfiable.

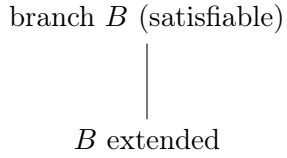

We need: B extended is satisfiable.

What we have already proved. In previous lectures, we proved the Preservation Principle for all **propositional rules**: $\mathsf{T}\wedge$, $\mathsf{F}\wedge$, $\mathsf{T}\vee$, $\mathsf{F}\vee$, $\mathsf{T}\to$, $\mathsf{F}\to$, $\mathsf{T}\neg$, $\mathsf{F}\neg$. Some of these are branching (e.g. $\mathsf{T}\vee$ splits into two branches); for each, we showed that satisfiability is preserved in the required sense.

The key observation.

Observation: All four modal rules ($\mathsf{T}\Box$, $\mathsf{F}\Box$, $\mathsf{T}\Diamond$, $\mathsf{F}\Diamond$) are **non-branching**.

This simplifies our task. To complete the soundness proof for K, we only need to show:

What remains: For each modal rule, if a branch is satisfiable and we apply the rule, the extended branch remains satisfiable.

4 Satisfiability of a Branch: A Precise Definition

Before proving anything, we need to be precise about what "satisfiable" means for a branch.

Definition 4.1 (Satisfiability of a branch)**.** A branch B of a tableau is **satisfiable** if there exist a Kripke model $\mathcal{M} = \langle W, R, I\rangle$ and an assignment σ mapping each world label i occurring on B to a world $\sigma(i) \in W$, such that:

1. For every $i : \mathsf{T}\,\psi$ on B: $\mathcal{M}, \sigma(i) \models \psi$.
2. For every $i : \mathsf{F}\,\psi$ on B: $\mathcal{M}, \sigma(i) \not\models \psi$.
3. For every accessibility statement iRj on B: $\sigma(i)\,R\,\sigma(j)$.

In words: there is a model that "makes good" on every claim recorded on the branch — every T-signed formula is true, every F-signed formula is false, and every recorded accessibility link actually holds in the model.

Remark 4.1. The initial branch of a tableau for $\Gamma \cup \{\neg\varphi\}$ consists of

$$1 : \mathsf{T}\,\gamma_1,\ \ldots,\ 1 : \mathsf{T}\,\gamma_n,\ 1 : \mathsf{F}\,\varphi.$$

This branch is satisfiable if and only if $\Gamma \cup \{\neg\varphi\}$ is satisfiable: just set $\sigma(1) = w$ where w is a world satisfying all of Γ and falsifying φ.

5 The Four Modal Rules

We recall the four modal rules for K. Each uses **prefixed signed formulas**: a world label i, a sign (T or F), and a formula.

Rule $\mathsf{T}\Box$ (universal, no new world).

If the branch contains $i : \mathsf{T}\,\Box\varphi$ and iRj, add $j : \mathsf{T}\,\varphi$.

$$i : \mathsf{T}\,\Box\varphi \quad \text{and} \quad iRj \quad \text{(on branch)}$$
$$\Big|$$
$$j : \mathsf{T}\,\varphi$$

Rule $\mathsf{F}\Box$ (existential, new world).
From $i : \mathsf{F}\,\Box\varphi$, introduce fresh j. Add iRj and $j : \mathsf{F}\,\varphi$.

$$i : \mathsf{F}\,\Box\varphi$$
$$\Big|$$
$$iRj \quad \text{(fresh } j\text{)}$$
$$\Big|$$
$$j : \mathsf{F}\,\varphi$$

Rule $\mathsf{T}\Diamond$ (existential, new world).
From $i : \mathsf{T}\,\Diamond\varphi$, introduce fresh j. Add iRj and $j : \mathsf{T}\,\varphi$.

$$i : \mathsf{T}\,\Diamond\varphi$$
$$\Big|$$
$$iRj \quad \text{(fresh } j\text{)}$$
$$\Big|$$
$$j : \mathsf{T}\,\varphi$$

Rule $\mathsf{F}\Diamond$ (universal, no new world).
If the branch contains $i : \mathsf{F}\,\Diamond\varphi$ and iRj, add $j : \mathsf{F}\,\varphi$.

$$i : \mathsf{F}\,\Diamond\varphi \quad \text{and} \quad iRj \quad \text{(on branch)}$$
$$\Big|$$
$$j : \mathsf{F}\,\varphi$$

Rule	**Type**	**New world?**	**Branching?**
$\mathsf{T}\Box$	Universal	No	No
$\mathsf{F}\Box$	Existential	Yes	No
$\mathsf{T}\Diamond$	Existential	Yes	No
$\mathsf{F}\Diamond$	Universal	No	No

6 The Four Soundness Lemmas

We now prove that each modal rule preserves branch satisfiability. The proofs come in two pairs, reflecting the two types of rules.

6.1 Universal rules: T□ and F◇

These rules apply to an *existing* accessibility link iRj already on the branch. They add a signed formula at j but do not create a new world. The model and the assignment σ do not need to be modified at all.

Lemma 6.1 (Soundness of T□)**.** *If a branch B is satisfiable and contains both $i : \mathsf{T}\,\Box\varphi$ and iRj, then $B' = B \cup \{j : \mathsf{T}\,\varphi\}$ is satisfiable.*

Proof. Since B is satisfiable, there exist $\mathcal{M} = \langle W, R, I\rangle$ and σ satisfying B.

From $i : \mathsf{T}\,\Box\varphi$: $\mathcal{M}, \sigma(i) \models \Box\varphi$.

From iRj: $\sigma(i)\,R\,\sigma(j)$.

By the semantics of □: if $\mathcal{M}, \sigma(i) \models \Box\varphi$, then for *every* w with $\sigma(i)\,R\,w$, $\mathcal{M}, w \models \varphi$. In particular, since $\sigma(i)\,R\,\sigma(j)$:

$$\mathcal{M}, \sigma(j) \models \varphi.$$

This is exactly the condition for $j : \mathsf{T}\,\varphi$. The same $\mathcal{M}$ and σ satisfy B'. □

Lemma 6.2 (Soundness of F◇)**.** *If a branch B is satisfiable and contains both $i : \mathsf{F}\,\Diamond\varphi$ and iRj, then $B' = B \cup \{j : \mathsf{F}\,\varphi\}$ is satisfiable.*

Proof. Since B is satisfiable, there exist $\mathcal{M}$ and σ satisfying B.

From $i : \mathsf{F}\,\Diamond\varphi$: $\mathcal{M}, \sigma(i) \not\models \Diamond\varphi$.

From iRj: $\sigma(i)\,R\,\sigma(j)$.

By the semantics of ◇: if $\mathcal{M}, \sigma(i) \not\models \Diamond\varphi$, then for *every* w with $\sigma(i)\,R\,w$, $\mathcal{M}, w \not\models \varphi$. In particular, since $\sigma(i)\,R\,\sigma(j)$:

$$\mathcal{M}, \sigma(j) \not\models \varphi.$$

This is exactly the condition for $j : \mathsf{F}\,\varphi$. The same $\mathcal{M}$ and σ satisfy B'. □

Remark 6.1. The two proofs are structurally identical. Both exploit the *universal* character of the rule: $\mathsf{T}\,\Box\varphi$ says "φ true at all accessible worlds," $\mathsf{F}\,\Diamond\varphi$ says "φ false at all accessible worlds." In both cases, instantiating at $\sigma(j)$ is immediate.

6.2 Existential rules: F□ and T◇

These rules create a *new* world label j together with a new accessibility link iRj. Here the model $\mathcal{M}$ stays the same, but the assignment σ must be *extended* to cover j. The freshness of j is what makes this possible.

Lemma 6.3 (Soundness of F□)**.** *If a branch B is satisfiable and contains $i : \mathsf{F}\,\Box\varphi$, then $B' = B \cup \{iRj,\ j : \mathsf{F}\,\varphi\}$ (with j fresh) is satisfiable.*

Proof. Since B is satisfiable, there exist $\mathcal{M} = \langle W, R, I\rangle$ and σ satisfying B.

From $i : \mathsf{F}\,\Box\varphi$: $\mathcal{M}, \sigma(i) \not\models \Box\varphi$.

By the semantics of □: there exists $w_0 \in W$ such that

$$\sigma(i)\,R\,w_0 \qquad \text{and} \qquad \mathcal{M}, w_0 \not\models \varphi.$$

Since j is **fresh** — it does not appear anywhere on B — we may extend σ by setting $\sigma(j) = w_0$. This does not affect the satisfaction of any formula already on B (those formulas involve only labels present before j was introduced).

Now check the new items:

- iRj: $\sigma(i)\,R\,\sigma(j) = \sigma(i)\,R\,w_0$. Holds by choice of w_0. ✓
- $j : \mathsf{F}\,\varphi$: $\mathcal{M}, \sigma(j) \not\models \varphi$, i.e. $\mathcal{M}, w_0 \not\models \varphi$. Holds by choice of w_0. ✓

The same model $\mathcal{M}$ and the extended σ satisfy B'. □

Lemma 6.4 (Soundness of $\mathsf{T}\Diamond$)**.** *If a branch B is satisfiable and contains $i : \mathsf{T}\,\Diamond\varphi$, then $B' = B \cup \{iRj,\ j : \mathsf{T}\,\varphi\}$ (with j fresh) is satisfiable.*

Proof. Since B is satisfiable, there exist $\mathcal{M}$ and σ satisfying B.

From $i : \mathsf{T}\,\Diamond\varphi$: $\mathcal{M}, \sigma(i) \models \Diamond\varphi$.

By the semantics of $\Diamond$: there exists $w_0 \in W$ such that

$$\sigma(i)\,R\,w_0 \qquad \text{and} \qquad \mathcal{M}, w_0 \models \varphi.$$

Set $\sigma(j) = w_0$ (this is allowed because j is fresh). Check:

- iRj: $\sigma(i)\,R\,w_0$. ✓
- $j : \mathsf{T}\,\varphi$: $\mathcal{M}, w_0 \models \varphi$. ✓

The same model $\mathcal{M}$ and the extended σ satisfy B'. □

Remark 6.2. Again, the two proofs are structurally identical. Both exploit the *existential* character of the semantics: $\mathsf{F}\,\Box\varphi$ says "there exists an accessible world where φ fails," $\mathsf{T}\,\Diamond\varphi$ says "there exists an accessible world where φ holds." In both cases, we use the guaranteed witness as the interpretation of the fresh label.

The role of freshness. Freshness is essential. If j already appeared on B, then $\sigma(j)$ would already be determined by the satisfying model, and we could not freely set $\sigma(j) = w_0$. The existential rules work precisely because they introduce a label that carries no prior commitments.

7 The Soundness Theorem: Full Proof

We now assemble the pieces.

Theorem 7.1 (Soundness of the tableau method for K)**.** *If $\Gamma \vdash_K \varphi$, then $\Gamma \models_K \varphi$.*

Proof. We prove the contrapositive: if $\Gamma \not\models_K \varphi$, then $\Gamma \nvdash_K \varphi$.

Step 1: The initial branch is satisfiable. Suppose $\Gamma \not\models_K \varphi$. Then $\Gamma \cup \{\neg\varphi\}$ is satisfiable: there exist a Kripke model $\mathcal{M}$ and a world w such that $\mathcal{M}, w \models \gamma$ for every $\gamma \in \Gamma$ and $\mathcal{M}, w \not\models \varphi$.

The initial branch of the tableau is:

$$1 : \mathsf{T}\,\gamma_1,\ \ldots,\ 1 : \mathsf{T}\,\gamma_n,\ 1 : \mathsf{F}\,\varphi.$$

Setting $\sigma(1) = w$, this branch is satisfiable (Definition 4.1).

Step 2: Every rule application preserves satisfiability. At each step of the tableau construction, we apply one rule. By the Preservation Principle:

- **Propositional rules** (proved in previous lectures): if the branch before the rule is satisfiable, at least one resulting branch is satisfiable.
- **Modal rules** (Lemmas 6.1–6.4): if the branch before the rule is satisfiable, the extended branch is satisfiable. (These are non-branching, so the condition is simpler.)

Step 3: Induction. By induction on the number of rule applications: at every stage of the tableau, at least one branch is satisfiable.

Step 4: A satisfiable branch cannot close. A branch closes only if it contains $i : \mathsf{T}\ p$ and $i : \mathsf{F}\ p$ for some atom p at some label i. But if the branch is satisfiable, the model $\mathcal{M}$ and assignment σ would require both $\mathcal{M}, \sigma(i) \models p$ and $\mathcal{M}, \sigma(i) \not\models p$, which is impossible.

Conclusion. Since at least one branch remains satisfiable at every stage, at least one branch remains open. The tableau does not close. Therefore $\Gamma \nvdash_K \varphi$. □

8 Structural Summary

The soundness proof has a clean modular structure:

1. **Reduce** soundness to its contrapositive: satisfiability of $\Gamma \cup \{\neg\varphi\}$ implies the tableau stays open.
2. **Decompose** the problem into individual rules: show that each rule preserves branch satisfiability.
3. **Classify** the rules:
   - Branching rules (propositional): at least one branch stays satisfiable. *(Already proved.)*
   - Non-branching rules (all modal rules): the extended branch stays satisfiable. *(Proved here.)*
4. **Assemble**: by induction, some branch remains satisfiable at every stage, so the tableau never fully closes.

The individual lemmas for the modal rules fall into two symmetric pairs:

Rule	**Semantic fact used**	**What happens to σ**
$\mathsf{T}\Box$	$\Box\varphi$ true $\Rightarrow$ φ true at all successors	σ unchanged
$\mathsf{F}\Diamond$	$\Diamond\varphi$ false $\Rightarrow$ φ false at all successors	σ unchanged
$\mathsf{F}\Box$	$\Box\varphi$ false $\Rightarrow$ $\exists$ successor where φ false	σ extended $(j \mapsto w_0)$
$\mathsf{T}\Diamond$	$\Diamond\varphi$ true $\Rightarrow$ $\exists$ successor where φ true	σ extended $(j \mapsto w_0)$

The universal rules (*top pair*) simply instantiate a universal quantifier at an existing world. The existential rules (*bottom pair*) use a semantic witness as the interpretation of a fresh label. In both cases, the model itself is never modified — only the assignment σ may grow.

Part II
Completeness

9 What Completeness Says

Theorem 9.1 (Completeness of the tableau method for K)**.** *If* $\Gamma \models \varphi$*, then* $\Gamma \vdash \varphi$*.*

In other words: *the tableau method is powerful enough to prove everything that is true.* If φ really does follow from Γ, then the tableau will eventually close.

10 Soundness vs. Completeness: A Contrast in Strategy

The soundness and completeness proofs are *not* symmetric. It is worth pausing to understand how they differ.

Soundness was atomistic. We examined each rule *individually* and showed it preserves satisfiability. The proof never needed to look at the tableau as a whole — it worked one rule at a time.

Completeness must be holistic. We must show that the system *as a whole* is complete. The contrapositive formulation makes the task concrete:

$$\Gamma \nvdash \varphi \implies \Gamma \not\models \varphi.$$

$\Gamma \nvdash \varphi$ means the tableau for $\Gamma \cup \{\neg\varphi\}$ stays open: there is a finished open branch. $\Gamma \not\models \varphi$ means $\Gamma \cup \{\neg\varphi\}$ is satisfiable.

So we need to prove:

> **Key Claim**: If a finished tableau has an open branch B, then the set of signed formulas on B is satisfiable — i.e., there exists a Kripke model making every T-signed formula true and every F-signed formula false.

The word "finished" is doing crucial work here. A branch is *finished* when every applicable rule has been applied: every $\mathsf{T}\Box$ has been instantiated at every accessible world on the branch, every $\mathsf{T}\Diamond$ has been given its witness, and so on. A half-built branch tells us nothing.

11 The Strategy: Build a Model from the Branch

The completeness proof proceeds in two steps:

1. **Construct** a Kripke model $\mathcal{M}$ directly from the information recorded on the open branch B.
2. **Show** (by induction on the complexity of formulas) that $\mathcal{M}$ satisfies every T-signed formula and falsifies every F-signed formula on B.

This is the same strategy used for first-order logic, where constants on the branch become elements of the domain. Here, *world labels* on the branch become the possible worlds of the model. The idea is simple and powerful: **read off the model from the branch**.

12 Building a Kripke Model from an Open Branch

Let B be a finished open branch of a tableau.

Definition 12.1 (Kripke model induced by a branch)**.** The **branch model** $\mathcal{M}_B = \langle W, R, I \rangle$ is defined as follows:

- **Worlds**: W is the set of all world labels appearing on B.
  (Every label i that occurs anywhere on B — whether as the prefix of a signed formula or in an accessibility statement — becomes a world.)
- **Accessibility**: $R = \{(i, j) : iRj \text{ appears on } B\}$.
  (We simply read off the accessibility links recorded on the branch.)
- **Valuation**: for each propositional variable p and each world $i \in W$:

$$I(p, i) = \begin{cases} T & \text{if } i : \mathsf{T}\, p \text{ appears on } B, \\ F & \text{otherwise.} \end{cases}$$

This definition deserves a few comments.

Remark 12.1 (Well-definedness). The valuation is well-defined because B is an *open* branch: it does not contain both $i : \mathsf{T}\, p$ and $i : \mathsf{F}\, p$ for the same atom p at the same world i (otherwise B would be closed). So there is no conflict in the definition of I.

Remark 12.2 (Where do the worlds come from?). Some of the world labels on B come from the initial setup (typically just label 1). Others were introduced during the tableau construction by the existential rules ($\mathsf{F}\Box$ and $\mathsf{T}\Diamond$), which create fresh labels. The set W includes *all* of them.

Remark 12.3 (No construction needed for R). In the term model for first-order logic, we had to define the interpretation of predicates by reading off atomic formulas. Here, we additionally read off the accessibility relation from the iRj statements on the branch. This is the only genuinely new ingredient compared to the propositional case.

Let us visualise the idea with a small example. Suppose a finished open branch contains:

$$1 : \mathsf{T}\, p, \quad 1 : \mathsf{T}\, \Diamond q, \quad 1R2, \quad 2 : \mathsf{T}\, q, \quad 1 : \mathsf{F}\, q.$$

The branch model has $W = \{1, 2\}$, $R = \{(1, 2)\}$, $I(p, 1) = T$, $I(q, 1) = F$, $I(q, 2) = T$:

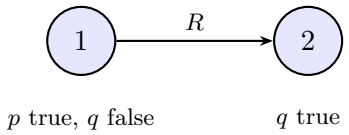

13 The Inductive Proof

We now prove that the branch model $\mathcal{M}_B$ satisfies every signed formula on the branch. The proof proceeds by induction on the complexity of formulas.

Theorem 13.1 (Branch satisfiability)**.** *Let B be a finished open branch, and let $\mathcal{M}_B$ be the branch model. For every formula φ appearing on B:*

(a) if $i : \mathsf{T}\, \varphi$ is on B, then $\mathcal{M}_B, i \models \varphi$;
(b) if $i : \mathsf{F}\, \varphi$ is on B, then $\mathcal{M}_B, i \not\models \varphi$.

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

Base case: φ is a propositional variable p.

- If $i : \mathsf{T}\, p$ is on B: by definition of I, $I(p, i) = T$, so $\mathcal{M}_B, i \models p$. ✓
- If $i : \mathsf{F}\, p$ is on B: since B is open, $i : \mathsf{T}\, p$ is *not* on B, so $I(p, i) = F$, and $\mathcal{M}_B, i \not\models p$. ✓

This is the step where the openness of the branch is used: no atom is both T-signed and F-signed at the same world.

Inductive step: propositional connectives. The cases for $\neg$, $\wedge$, $\vee$, and $\rightarrow$ are exactly as in the propositional completeness proof. We recall the pattern briefly for one representative case and omit the others.

Case $\varphi = \psi_1 \wedge \psi_2$.

- $i : \mathsf{T}\ \psi_1 \wedge \psi_2$ on B. The rule is non-branching: both $i : \mathsf{T}\ \psi_1$ and $i : \mathsf{T}\ \psi_2$ are on B. By the inductive hypothesis, $\mathcal{M}_B, i \models \psi_1$ and $\mathcal{M}_B, i \models \psi_2$. Therefore $\mathcal{M}_B, i \models \psi_1 \wedge \psi_2$. ✓
- $i : \mathsf{F}\ \psi_1 \wedge \psi_2$ on B. The rule is branching: on B, at least one of $i : \mathsf{F}\ \psi_1$ or $i : \mathsf{F}\ \psi_2$ appears (whichever sub-branch B follows). By the inductive hypothesis, $\mathcal{M}_B, i \not\models \psi_1$ or $\mathcal{M}_B, i \not\models \psi_2$. Either way, $\mathcal{M}_B, i \not\models \psi_1 \wedge \psi_2$. ✓

The other propositional cases ($\neg$, $\vee$, $\rightarrow$) follow the same pattern: non-branching rules give us both sub-formulas on the branch; branching rules give us at least one. In every case, the inductive hypothesis delivers the result. We omit the details since they were proved in the propositional completeness lecture.

We now come to the **modal cases**, which are the heart of this proof. There are four, organised in two pairs that mirror the soundness lemmas.

Case $\varphi = \Box\psi$ (true-signed): $i : \mathsf{T}\ \Box\psi$ on B. We need to show: $\mathcal{M}_B, i \models \Box\psi$, i.e. for every $j \in W$ with iRj, $\mathcal{M}_B, j \models \psi$.

Let j be any world with iRj in $\mathcal{M}_B$. By definition of $\mathcal{M}_B$, this means iRj appears on B.

Now, B is a *finished* branch. The rule $\mathsf{T}\Box$ says: if $i : \mathsf{T}\ \Box\psi$ and iRj are both on the branch, add $j : \mathsf{T}\ \psi$. Since B is finished, this rule *has been applied*. Therefore $j : \mathsf{T}\ \psi$ is on B.

By the inductive hypothesis (since ψ is less complex than $\Box\psi$): $\mathcal{M}_B, j \models \psi$.

Since j was an arbitrary world accessible from i: $\mathcal{M}_B, i \models \Box\psi$. ✓

Remark 13.1. This is the step that relies most heavily on the branch being *finished*. If the $\mathsf{T}\Box$ rule had not yet been applied to some particular iRj on the branch, we would not know that $j : \mathsf{T}\ \psi$ is there, and the argument would fail.

Case $\varphi = \Box\psi$ (false-signed): $i : \mathsf{F}\ \Box\psi$ on B. We need to show: $\mathcal{M}_B, i \not\models \Box\psi$, i.e. there exists some j with iRj such that $\mathcal{M}_B, j \not\models \psi$.

Since B is finished, the rule $\mathsf{F}\Box$ has been applied to $i : \mathsf{F}\ \Box\psi$. This rule introduces a *fresh* label j and adds iRj and $j : \mathsf{F}\ \psi$ to the branch. So there exists a label j such that:

- iRj is on B, which means iRj in $\mathcal{M}_B$;
- $j : \mathsf{F}\ \psi$ is on B.

By the inductive hypothesis: $\mathcal{M}_B, j \not\models \psi$.

Since j is accessible from i and ψ fails at j: $\mathcal{M}_B, i \not\models \Box\psi$. ✓

Case $\varphi = \Diamond\psi$ (true-signed): $i : \mathsf{T}\ \Diamond\psi$ on B. We need to show: $\mathcal{M}_B, i \models \Diamond\psi$, i.e. there exists some j with iRj such that $\mathcal{M}_B, j \models \psi$.

Since B is finished, the rule $\mathsf{T}\Diamond$ has been applied. It introduced a fresh label j with iRj and $j : \mathsf{T}\ \psi$ on B. So:

- iRj in $\mathcal{M}_B$;
- $j : \mathsf{T}\ \psi$ is on B.

By the inductive hypothesis: $\mathcal{M}_B, j \models \psi$.

So there exists an accessible world satisfying ψ: $\mathcal{M}_B, i \models \Diamond\psi$. ✓

Case $\varphi = \Diamond\psi$ (false-signed): $i : \mathsf{F}\,\Diamond\psi$ on B. We need to show: $\mathcal{M}_B, i \not\models \Diamond\psi$, i.e. for every j with iRj, $\mathcal{M}_B, j \not\models \psi$.

Let j be any world with iRj in $\mathcal{M}_B$. Then iRj is on B.

Since B is finished, the rule $\mathsf{F}\Diamond$ has been applied: from $i : \mathsf{F}\,\Diamond\psi$ and iRj, we get $j : \mathsf{F}\,\psi$ on B.

By the inductive hypothesis: $\mathcal{M}_B, j \not\models \psi$.

Since j was arbitrary: $\mathcal{M}_B, i \not\models \Diamond\psi$. ✓

All cases have been verified. The induction is complete. □

14 The Pattern: Universal vs. Existential

Before stating the completeness theorem, let us pause to admire the structure of the four modal cases. They fall into two pairs that are worth comparing explicitly.

Rule	What it says about B	Why the case works
$\mathsf{T}\Box$	For *every* j with iRj on B: $j : \mathsf{T}\,\psi$ is on B.	These are all the accessible worlds in $\mathcal{M}_B$ (by construction of R). So ψ holds at all of them.
$\mathsf{F}\Diamond$	For *every* j with iRj on B: $j : \mathsf{F}\,\psi$ is on B.	Same reasoning: ψ fails at every accessible world.
$\mathsf{F}\Box$	There *exists* a j with iRj and $j : \mathsf{F}\,\psi$ on B.	This j is a world in $\mathcal{M}_B$ witnessing the failure of $\Box\psi$.
$\mathsf{T}\Diamond$	There *exists* a j with iRj and $j : \mathsf{T}\,\psi$ on B.	This j is a world in $\mathcal{M}_B$ witnessing the truth of $\Diamond\psi$.

The top pair (universal rules) works because the branch is *finished*: the rule has been applied to *every* iRj on the branch, and the worlds of $\mathcal{M}_B$ are precisely the labels on the branch — nothing more, nothing less.

The bottom pair (existential rules) works because the rule *created* the witness: the fresh label j is now a world in $\mathcal{M}_B$ with exactly the right properties.

This is the same universal/existential duality we saw in the soundness proof, now viewed from the other side:

> **Soundness**: the model was *given*, and we checked that each rule preserves it.
> **Completeness**: the model is *built from the branch*, and we check that it satisfies everything on the branch.

15 The Completeness Theorem: Full Proof

Theorem 15.1 (Completeness of the tableau method for K)**.** *If* $\Gamma \models_K \varphi$*, then* $\Gamma \vdash_K \varphi$.

Proof. We prove the contrapositive: if $\Gamma \nvdash_K \varphi$, then $\Gamma \not\models_K \varphi$.

Step 1. Suppose $\Gamma \nvdash_K \varphi$. Then the tableau for $\Gamma \cup \{\neg\varphi\}$ does not close: there exists a finished open branch B.

Step 2. Construct the branch model $\mathcal{M}_B = \langle W, R, I \rangle$ from B (Definition 12.1).

Step 3. By Theorem 13.1, $\mathcal{M}_B$ satisfies every T-signed formula on B and falsifies every F-signed formula on B.

Step 4. The initial signed formulas $1 : \mathsf{T}\ \gamma_1, \ldots, 1 : \mathsf{T}\ \gamma_n, 1 : \mathsf{F}\ \varphi$ are all on B (they are at the root of the tableau, which every branch passes through). Therefore:

- $\mathcal{M}_B, 1 \models \gamma_k$ for each $\gamma_k \in \Gamma$;
- $\mathcal{M}_B, 1 \not\models \varphi$.

So $\Gamma \cup \{\neg\varphi\}$ is satisfiable, which means $\Gamma \not\models_K \varphi$. □

Part III
Decidability

16 The Question

Soundness tells us: if the tableau closes, the answer is correct. Completeness tells us: if φ really follows from Γ, the tableau will close. But neither theorem addresses a practical question:

Does the tableau procedure always terminate?

If it does, then we have a **decision procedure**: an algorithm that, given any finite Γ and any formula φ, will always stop after finitely many steps and tell us whether $\Gamma \models \varphi$ or not.

Why termination is not obvious. The existential rules ($\mathsf{F}\Box$ and $\mathsf{T}\Diamond$) create *new* world labels. Each new label may trigger further applications of the universal rules ($\mathsf{T}\Box$ and $\mathsf{F}\Diamond$). Could this cascade indefinitely, producing an ever-growing chain of worlds?

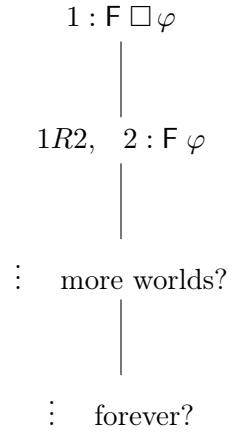

The answer is *no*, and the reason is the **subformula property**: tableau rules always decompose formulas into simpler ones. This bounds both the number of worlds and the information at each world.

Contrast with first-order logic. In the first-order tableau method, the rule for $\mathsf{T}\ \forall x\,\varphi(x)$ must be re-applied every time a new constant appears on the branch. New constants are introduced by the $\mathsf{T}\ \exists$ rule, and those new instances of φ may contain further existential quantifiers, generating yet more constants. This process may never terminate. Indeed, first-order logic is *undecidable*: no algorithm can decide validity in general.

In propositional modal logic, the analogue of "new constants" is "new world labels." But as we shall see, new worlds can only be created to a *bounded depth*, determined by the input formula. This is the fundamental difference that makes propositional modal logic decidable.

17 The Subformula Property

Definition 17.1 (Subformulas)**.** The set $\mathrm{Sub}(\varphi)$ of **subformulas** of φ is defined recursively:

- $\mathrm{Sub}(p) = \{p\}$ for a propositional variable p.
- $\mathrm{Sub}(\neg\psi) = \{\neg\psi\} \cup \mathrm{Sub}(\psi)$.
- $\mathrm{Sub}(\psi_1 \star \psi_2) = \{\psi_1 \star \psi_2\} \cup \mathrm{Sub}(\psi_1) \cup \mathrm{Sub}(\psi_2)$, where $\star \in \{\wedge, \vee, \to\}$.
- $\mathrm{Sub}(\Box\,\psi) = \{\Box\,\psi\} \cup \mathrm{Sub}(\psi)$.

- $\mathrm{Sub}(\Diamond\,\psi) = \{\Diamond\,\psi\} \cup \mathrm{Sub}(\psi)$.

For a finite set $\Gamma = \{\gamma_1, \ldots, \gamma_n\}$, we write $\mathrm{Sub}(\Gamma) = \mathrm{Sub}(\gamma_1) \cup \cdots \cup \mathrm{Sub}(\gamma_n)$.

The crucial point: $\mathrm{Sub}(\varphi)$ is always *finite*, and its size is bounded by the length of φ.

Lemma 17.1 (Subformula property of tableaux)**.** *Let $\Sigma = \Gamma \cup \{\neg\varphi\}$ be the initial set of formulas. On every branch of the tableau for Σ, every formula appearing in a signed entry $i : \mathsf{T}\ \psi$ or $i : \mathsf{F}\ \psi$ is a subformula of some formula in Σ.*

In other words: no rule ever introduces a formula that is not already a subformula of the input.

Proof. By inspection of each rule.

Propositional rules. Each propositional rule decomposes a compound formula into its immediate sub-formulas. For example:

- $\mathsf{T}\ \psi_1 \wedge \psi_2$ produces $\mathsf{T}\ \psi_1$ and $\mathsf{T}\ \psi_2$: both are subformulas of $\psi_1 \wedge \psi_2$.
- $\mathsf{F}\ \psi_1 \vee \psi_2$ produces $\mathsf{F}\ \psi_1$ and $\mathsf{F}\ \psi_2$: same.
- $\mathsf{T}\ \neg\psi$ produces $\mathsf{F}\ \psi$: ψ is a subformula of $\neg\psi$.

All other propositional rules follow the same pattern.

Modal rules. Each modal rule strips away one layer of $\Box$ or $\Diamond$:

- $\mathsf{T}\Box$: from $i : \mathsf{T}\ \Box\,\psi$ and iRj, produce $j : \mathsf{T}\ \psi$. The formula ψ is a subformula of $\Box\,\psi$.
- $\mathsf{F}\Box$: from $i : \mathsf{F}\ \Box\,\psi$, produce $j : \mathsf{F}\ \psi$. Same.
- $\mathsf{T}\Diamond$: from $i : \mathsf{T}\ \Diamond\,\psi$, produce $j : \mathsf{T}\ \psi$. The formula ψ is a subformula of $\Diamond\,\psi$.
- $\mathsf{F}\Diamond$: from $i : \mathsf{F}\ \Diamond\,\psi$ and iRj, produce $j : \mathsf{F}\ \psi$. Same.

In every case, the output formula is a proper subformula of the input formula. No rule ever creates a formula more complex than the input. □

Remark 17.1. This is where propositional modal logic differs sharply from first-order logic. In FOL, the rule for $\mathsf{T}\ \forall x\,\varphi(x)$ produces $\mathsf{T}\ \varphi(a)$, which is an *instance* of φ, not a subformula. The formula $\varphi(a)$ may itself contain quantifiers that generate further instances with new constants. In modal logic, the rule for $\mathsf{T}\ \Box\,\psi$ produces $\mathsf{T}\ \psi$, which is a genuine *subformula* — strictly simpler. This is what prevents the infinite regress.

18 Finiteness of the Tableau

We now use the subformula property to establish two bounds, which together guarantee that every branch — and therefore the entire tableau — is finite.

Definition 18.1 (Modal depth)**.** The **modal depth** $\mathrm{md}(\varphi)$ of a formula φ is the maximum nesting depth of modal operators in φ:

- $\mathrm{md}(p) = 0$.
- $\mathrm{md}(\neg\psi) = \mathrm{md}(\psi)$.
- $\mathrm{md}(\psi_1 \star \psi_2) = \max(\mathrm{md}(\psi_1), \mathrm{md}(\psi_2))$, where $\star \in \{\wedge, \vee, \to\}$.
- $\mathrm{md}(\Box\,\psi) = 1 + \mathrm{md}(\psi)$.
- $\mathrm{md}(\Diamond\,\psi) = 1 + \mathrm{md}(\psi)$.

For a set Σ, $\mathrm{md}(\Sigma) = \max\{\mathrm{md}(\psi) : \psi \in \Sigma\}$.

For example, $\mathrm{md}(\Box(p \wedge \Diamond\,q)) = 1 + \max(0, 1 + 0) = 2$.

18.1 Bound 1: The number of worlds is bounded

Lemma 18.1 (Depth bound)**.** *On any branch of the tableau for Σ, if a world label j is at "distance" d from the root label 1 (i.e. it was reached through a chain $1Ri_1Ri_2R\cdots Rj$ of d accessibility links), then every formula at world j has modal depth at most* $\mathrm{md}(\Sigma) - d$.

Proof. By induction on d.

Base case ($d = 0$). At the root world 1, the formulas are the input formulas from Σ and their subformulas. These all have modal depth at most $\mathrm{md}(\Sigma)$.

Inductive step. Suppose at distance d, all formulas have modal depth at most $\mathrm{md}(\Sigma) - d$. A new world j at distance $d + 1$ is created by an existential rule ($\mathsf{F}\Box$ or $\mathsf{T}\Diamond$) applied to a modal formula at distance d. This modal formula has the form $\Box\psi$ or $\Diamond\psi$ with modal depth at least $1 + \mathrm{md}(\psi)$, and the rule produces ψ at the new world. So:

$$\mathrm{md}(\psi) \leq (\mathrm{md}(\Sigma) - d) - 1 = \mathrm{md}(\Sigma) - (d + 1).$$

Furthermore, the only formulas that can appear at world j (via universal rules $\mathsf{T}\Box$ or $\mathsf{F}\Diamond$ applied from ancestors) are obtained by stripping one modal operator from formulas at ancestor worlds, which also decreases modal depth by at least 1. □

Remark 18.1. The key consequence: when $d = \mathrm{md}(\Sigma)$, all formulas at that world have modal depth 0 — they are purely propositional, with no $\Box$ or $\Diamond$. No existential rule can fire (there is no modal formula to decompose), so **no new world can be created beyond depth** $\mathrm{md}(\Sigma)$.

This gives our first bound:

Bound 1: The chain of worlds on any branch has length at most $\mathrm{md}(\Sigma)$.

More precisely, the number of worlds on a branch is bounded. At each world, there can be at most as many existential rule applications as there are modal subformulas in $\mathrm{Sub}(\Sigma)$, and each application creates exactly one new world. So the total number of worlds on a branch is bounded (exponentially) by a function of $|\mathrm{Sub}(\Sigma)|$ and $\mathrm{md}(\Sigma)$.

18.2 Bound 2: The information at each world is bounded

Lemma 18.2 (Width bound)**.** *At any world i on a branch, the set of signed formulas $i : \mathsf{T}\ \psi$ and $i : \mathsf{F}\ \psi$ is drawn from* $\mathrm{Sub}(\Sigma)$. *In particular, there are at most $2 \times |\mathrm{Sub}(\Sigma)|$ signed formulas at each world.*

Proof. Immediate from the subformula property (Lemma 17.1): every formula at any world is a subformula of Σ, and each subformula can appear with sign T or F. □

This gives our second bound:

Bound 2: At each world, there are at most $2 \times |\mathrm{Sub}(\Sigma)|$ signed formulas.

Since the number of signed formulas at each world is bounded, each world reaches a point where no new propositional rule can fire (all compound formulas have been decomposed). Combined with Bound 1, this means every branch of the tableau is finite.

18.3 Putting the bounds together

Theorem 18.3 (Termination)**.** *For any finite Σ, every tableau for Σ is finite.*

Proof. A branch of the tableau is a sequence of rule applications. By Bound 1, the number of worlds on a branch is bounded. By Bound 2, the number of signed formulas at each world is bounded. Each rule application either adds a signed formula at an existing world or creates a new world with a signed formula. Since both the number of worlds and the formulas per world are bounded, the number of rule applications on any branch is bounded. Therefore every branch is finite.

A tableau is a tree with finitely many branches (each branching rule creates at most two sub-branches, and branching can happen at most finitely many times on a finite branch). A finite tree with finite branches is finite. □

19 The Decision Procedure

We now have all three ingredients. Combining soundness, completeness, and termination gives us a decision procedure for propositional modal logic K.

Theorem 19.1 (Decidability of K)**.** *There exists an algorithm that, given a finite set of formulas Γ and a formula φ, decides whether $\Gamma \models_K \varphi$.*

Proof. The algorithm is simply: *build the tableau for $\Gamma \cup \{\neg\varphi\}$.*

Step 1: Termination. By Theorem 18.3, the tableau construction terminates after finitely many steps. The algorithm always halts.

Step 2: Correctness. When the algorithm halts, the tableau is either closed or has an open branch.

- If the tableau is **closed**: every branch contains a contradiction $i : \mathsf{T}\ p$ and $i : \mathsf{F}\ p$. By soundness (Theorem 1.1), $\Gamma \models_K \varphi$. The algorithm answers **yes**.
- If the tableau has an **open branch**: by completeness (Theorem 9.1), $\Gamma \not\models_K \varphi$. Moreover, the branch model $\mathcal{M}_B$ from Definition 12.1 is an explicit *counter-model*: a finite Kripke model satisfying Γ but falsifying φ. The algorithm answers **no** and produces the counter-model.

□

Remark 19.1 (The finite model property). A by-product of the completeness and termination proofs is the **finite model property** for K: if φ is satisfiable, it is satisfiable in a *finite* Kripke model. The branch model $\mathcal{M}_B$ is always finite (its worlds are the finitely many labels on a finite branch). This means:

If φ has a model at all, it has a *finite* model.

This is a strong property that not all logics enjoy. It is one of the reasons propositional modal logic is so well-behaved computationally.

20 Summary: The Three Pillars

The three results proved in these notes work together as follows:

Theorem	**What it guarantees**
Soundness (Part I)	If the tableau closes, the conclusion is correct. *Method*: show each rule preserves satisfiability.
Completeness (Part II)	If the conclusion is correct, the tableau closes. *Method*: build a model from an open branch.
Termination (Part III)	The tableau always halts. *Method*: bound worlds by modal depth, formulas by subformula closure.
Together: the tableau method is a *decision procedure* for K.	

Part IV
Completeness of Labelled Natural Deduction

21 Why a New Proof?

In Part II, we proved the completeness of the *tableau* method. That proof was easy: an open branch of a finished tableau already *is* a model. We just read it off.

The labelled natural deduction system is a different proof system. To show it is complete, we need a different argument. The difficulty is that ND proofs do not produce a model as a by-product. We must **construct** a model from scratch, using only the abstract fact that a set of formulas is consistent (i.e., does not derive $\bot$).

The technique we use is due to Leon Henkin (1949) and was adapted to modal logic by Arnould Bayart (1958), Dana Scott, and David Kaplan. The key idea: treat *maximal consistent sets of formulas* as possible worlds.

22 Derivability in K

Throughout Part IV, we write $\Gamma \vdash_K \varphi$ to mean: there exists a labelled ND derivation of $w : \varphi$ from the set of labelled premises $\{w : \gamma \mid \gamma \in \Gamma\}$, all at the same initial world w. We write $\vdash_K \varphi$ when φ is derivable from zero premises (a *theorem* of K).

All the formulas $\varphi, \psi, \ldots$ in Γ are *unlabelled*: they are ordinary modal formulas. The world labels only appear inside derivations.

Definition 22.1 (Consistency)**.** A set Γ of formulas is **consistent** (in K) if $\Gamma \nvdash_K \bot$. It is **inconsistent** if $\Gamma \vdash_K \bot$.

23 Maximal Consistent Sets

Definition 23.1 (Maximal consistent set)**.** A set ω of formulas is **maximal consistent** if:

1. ω is consistent: $\omega \nvdash_K \bot$.
2. ω is maximal: for every formula φ, either $\varphi \in \omega$ or $\neg\varphi \in \omega$.

Think of a maximal consistent set as a complete description of a possible world: it answers every yes/no question (condition 2), and its answers are never self-contradictory (condition 1).

Lemma 23.1 (Properties of maximal consistent sets)**.** *Let ω be a maximal consistent set. Then:*

(i) $\varphi \in \omega$ *if and only if* $\omega \vdash_K \varphi$.
(ii) $\varphi \notin \omega$ *if and only if* $\neg\varphi \in \omega$.
(iii) $\varphi \wedge \psi \in \omega$ *if and only if* $\varphi \in \omega$ *and* $\psi \in \omega$.
(iv) $\varphi \vee \psi \in \omega$ *if and only if* $\varphi \in \omega$ *or* $\psi \in \omega$.
(v) $\varphi \to \psi \in \omega$ *if and only if* $\varphi \notin \omega$ *or* $\psi \in \omega$.
(vi) $\neg\varphi \in \omega$ *if and only if* $\varphi \notin \omega$.

Proof. These all follow from maximality and consistency. We show (i) and (iii); the others are similar.

Property (i). ($\Rightarrow$): If $\varphi \in \omega$, then $\omega \vdash_K \varphi$ trivially (φ is a premise).

($\Leftarrow$): If $\omega \vdash_K \varphi$ but $\varphi \notin \omega$, then by maximality $\neg\varphi \in \omega$. But then $\omega \vdash_K \varphi$ and $\omega \vdash_K \neg\varphi$, giving $\omega \vdash_K \bot$. This contradicts consistency.

Property (iii). ($\Rightarrow$): If $\varphi \wedge \psi \in \omega$, then $\omega \vdash_K \varphi$ and $\omega \vdash_K \psi$ (by $\wedge$E). By (i), $\varphi \in \omega$ and $\psi \in \omega$.

($\Leftarrow$): If $\varphi \in \omega$ and $\psi \in \omega$, then $\omega \vdash_K \varphi \wedge \psi$ (by $\wedge$I). By (i), $\varphi \wedge \psi \in \omega$. □

24 Lindenbaum's Lemma

Lindenbaum's Lemma guarantees that every consistent set can be extended to a maximal consistent set. This is the tool that lets us pass from "Γ does not derive a contradiction" to "there is a complete, coherent description of a possible world extending Γ."

Lemma 24.1 (Lindenbaum)**.** *If Γ is consistent, then there exists a maximal consistent set $\omega \supseteq \Gamma$.*

Proof. Since our language is countable (finitely many propositional variables, built up with $\neg$, $\wedge$, $\vee$, $\rightarrow$, $\Box$, $\Diamond$), we can enumerate all formulas:

$$\varphi_1, \varphi_2, \varphi_3, \dots$$

We build an ascending chain of consistent sets $\Gamma = \Gamma_0 \subseteq \Gamma_1 \subseteq \Gamma_2 \subseteq \cdots$ by deciding one formula at each step:

$$\Gamma_{n+1} = \begin{cases} \Gamma_n \cup \{\varphi_{n+1}\} & \text{if this set is consistent,} \\ \Gamma_n \cup \{\neg\varphi_{n+1}\} & \text{otherwise.} \end{cases}$$

Finally, set $\omega = \bigcup_{n=0}^{\infty} \Gamma_n$.

Claim 1: Each Γ_n is consistent. By induction on n. $\Gamma_0 = \Gamma$ is consistent by assumption. For the inductive step: if Γ_n is consistent and $\Gamma_n \cup \{\varphi_{n+1}\}$ is consistent, we take the first case and Γ_{n+1} is consistent. If $\Gamma_n \cup \{\varphi_{n+1}\}$ is inconsistent, we take the second case: $\Gamma_{n+1} = \Gamma_n \cup \{\neg\varphi_{n+1}\}$. Is this consistent? Suppose not: $\Gamma_n \cup \{\neg\varphi_{n+1}\} \vdash_K \bot$. Then $\Gamma_n \vdash_K \neg\neg\varphi_{n+1}$, hence $\Gamma_n \vdash_K \varphi_{n+1}$ (by double negation elimination). But then $\Gamma_n \cup \{\varphi_{n+1}\}$ is a trivial extension that adds a consequence of Γ_n, so it should be consistent — contradicting our assumption that it was inconsistent. So Γ_{n+1} is consistent.

Claim 2: ω is consistent. If $\omega \vdash_K \bot$, then by the compactness of derivations (a derivation uses only finitely many premises), there is a finite subset $\Delta \subseteq \omega$ with $\Delta \vdash_K \bot$. Since Δ is finite and $\omega = \bigcup \Gamma_n$, there exists N such that $\Delta \subseteq \Gamma_N$. But then $\Gamma_N \vdash_K \bot$, contradicting Claim 1.

Claim 3: ω is maximal. For every formula φ_{n+1}: at step $n+1$, we added either φ_{n+1} or $\neg\varphi_{n+1}$ to $\Gamma_{n+1} \subseteq \omega$. So either $\varphi_{n+1} \in \omega$ or $\neg\varphi_{n+1} \in \omega$. □

Remark 24.1. The construction is simple: we go through all formulas one by one and decide each one, always choosing in a way that preserves consistency. The result is a maximal consistent set that extends Γ. This is the modal version of the same argument used for classical propositional and first-order logic.

25 The Canonical Model

We are now ready to define the key construction.

Definition 25.1 (Canonical model for K)**.** The **canonical model** is the Kripke model $\mathcal{M}^c = \langle W^c, R^c, I^c \rangle$ defined as follows:

- **Worlds**: $W^c = \{\omega \mid \omega \text{ is a maximal consistent set of formulas}\}$.
  Each "world" is an entire maximally consistent theory.

- **Accessibility**: for $\omega, \omega' \in W^c$,

$$\omega \, R^c \, \omega' \quad \overset{\text{def}}{\iff} \quad \text{for every formula } \varphi,\ \Box\varphi \in \omega \implies \varphi \in \omega'.$$

  ω' is accessible from ω iff everything necessary in ω is true in ω'.
- **Valuation**: for each propositional variable p and each $\omega \in W^c$:

$$I^c(p, \omega) = T \quad \iff \quad p \in \omega.$$

Remark 25.1. The accessibility relation R^c captures the semantic meaning of $\Box$: saying $\Box\varphi \in \omega$ should mean φ is true at every accessible world. So we *define* accessibility to make this condition hold by fiat for the ($\Rightarrow$) direction.

The hard part will be the ($\Leftarrow$) direction: showing that if φ is true at every accessible world, then $\Box\varphi \in \omega$. This requires the Existence Lemma.

26 Auxiliary Facts

Before proving the Existence Lemma, we need four derivable facts about $\Box$ and $\Diamond$ in K. We prove them in labelled ND.

26.1 Fact 1: $\Box\varphi \wedge \Box\psi \vdash_K \Box(\varphi \wedge \psi)$

This was proved as an example in the previous lecture notes. For reference:

1	$w : \Box\varphi \wedge \Box\psi$	P
2	$w : \Box\varphi$	$\wedge$E, 1
3	$w : \Box\psi$	$\wedge$E, 1
4	$\quad wRv$	fresh v
5	$\quad v : \varphi$	$\Box$E, 2, 4
6	$\quad v : \psi$	$\Box$E, 3, 4
7	$\quad v : \varphi \wedge \psi$	$\wedge$I, 5, 6
8	$w : \Box(\varphi \wedge \psi)$	$\Box$I, 4—7

26.2 Fact 2: $\Box\alpha_1 \wedge \cdots \wedge \Box\alpha_n \vdash_K \Box(\alpha_1 \wedge \cdots \wedge \alpha_n)$

Generalisation of Fact 1 by repeated application. If we have n boxed formulas, we can collect them under a single box.

26.3 Fact 3: $\Box\varphi \wedge \Diamond\psi \vdash_K \Diamond(\varphi \wedge \psi)$

This fact says: if φ is necessary and ψ is possible, then $\varphi \wedge \psi$ is possible. The proof combines $\Box$E with the witness provided by $\Diamond$E:

1	$w : \Box\varphi \wedge \Diamond\psi$	P
2	$w : \Box\varphi$	$\wedge$E, 1
3	$w : \Diamond\psi$	$\wedge$E, 1
4	$\quad wRv, \;\; v : \psi$	fresh v
5	$\quad v : \varphi$	$\Box$E, 2, 4
6	$\quad v : \varphi \wedge \psi$	$\wedge$I, 5, 4
7	$\quad w : \Diamond(\varphi \wedge \psi)$	$\Diamond$I, 4, 6
8	$w : \Diamond(\varphi \wedge \psi)$	$\Diamond$E, 3, 4—7

26.4 Fact 4: If $\Diamond\varphi$ is consistent, then φ is consistent

Equivalently (by contrapositive): if φ is inconsistent ($\vdash_K \neg\varphi$), then $\Diamond\varphi$ is inconsistent.

Proof. Suppose $\vdash_K \neg\varphi$, i.e., $\neg\varphi$ is a theorem. By the necessitation rule: $\vdash_K \Box\neg\varphi$. Now we show $\Diamond\varphi \vdash_K \bot$:

1	$w : \Diamond\varphi$	P
2	$w : \Box\neg\varphi$	Nec. (from $\vdash_K \neg\varphi$)
3	$\quad wRv, \;\; v : \varphi$	fresh v
4	$\quad v : \neg\varphi$	$\Box$E, 2, 3
5	$\quad v : \bot$	$\neg$E, 3, 4
6	$w : \bot$	$\Diamond$E, 1, 3—5

So $\Diamond\varphi \vdash_K \bot$: $\Diamond\varphi$ is inconsistent. □

27 The Existence Lemma

This is the technical heart of the completeness proof. It guarantees that the canonical model has "enough" accessible worlds.

Lemma 27.1 (Existence Lemma)**.** *Let ω be a maximal consistent set. If $\Diamond\varphi \in \omega$, then there exists a maximal consistent set ω' such that $\omega \, R^c \, \omega'$ and $\varphi \in \omega'$.*

In other words: if $\Diamond\varphi$ belongs to a world ω, we can find an accessible world ω' where φ is true.

The strategy. Define the set

$$\Delta_\omega = \{\psi : \Box\psi \in \omega\}.$$

This is the set of all formulas that *must* hold at any world accessible from ω (by the definition of R^c). We will show that $\Delta_\omega \cup \{\varphi\}$ is consistent, then apply Lindenbaum's Lemma to extend it to a maximal consistent set ω'.

Proof. **Step 1: $\Delta_\omega \cup \{\varphi\}$ is consistent.**

Suppose for contradiction that $\Delta_\omega \cup \{\varphi\}$ is inconsistent. Then there exist finitely many formulas $\alpha_1, \ldots, \alpha_n \in \Delta_\omega$ such that:

$$\{\alpha_1, \ldots, \alpha_n, \varphi\} \vdash_K \bot.$$

(We use finitely many because a derivation has finitely many premises.)

This means $\alpha_1 \wedge \cdots \wedge \alpha_n \wedge \varphi$ is inconsistent, so $\vdash_K \neg(\alpha_1 \wedge \cdots \wedge \alpha_n \wedge \varphi)$. By Fact 4:

$$\Diamond(\alpha_1 \wedge \cdots \wedge \alpha_n \wedge \varphi) \text{ is inconsistent.}$$

Now, since $\alpha_1, \ldots, \alpha_n \in \Delta_\omega$, we have $\Box\alpha_1, \ldots, \Box\alpha_n \in \omega$. By Fact 2:

$$\omega \vdash_K \Box(\alpha_1 \wedge \cdots \wedge \alpha_n).$$

Since $\Diamond\varphi \in \omega$, by Fact 3:

$$\omega \vdash_K \Diamond((\alpha_1 \wedge \cdots \wedge \alpha_n) \wedge \varphi).$$

But we just showed that $\Diamond(\alpha_1 \wedge \cdots \wedge \alpha_n \wedge \varphi)$ is inconsistent, meaning $\vdash_K \neg\Diamond(\alpha_1 \wedge \cdots \wedge \alpha_n \wedge \varphi)$.

So ω derives both $\Diamond(\alpha_1 \wedge \cdots \wedge \alpha_n \wedge \varphi)$ and its negation. Hence $\omega \vdash_K \bot$, contradicting the consistency of ω.

Therefore $\Delta_\omega \cup \{\varphi\}$ is consistent.

Step 2: Apply Lindenbaum's Lemma. By Lemma 24.1, extend $\Delta_\omega \cup \{\varphi\}$ to a maximal consistent set ω'.

Step 3: Check the two required properties.

$\varphi \in \omega'$: Since $\varphi \in \Delta_\omega \cup \{\varphi\} \subseteq \omega'$. ✓

$\omega\, R^c\, \omega'$: We need to show that for every ψ, if $\Box\psi \in \omega$ then $\psi \in \omega'$. But if $\Box\psi \in \omega$, then $\psi \in \Delta_\omega \subseteq \omega'$. ✓ □

Remark 27.1. The Existence Lemma is where the four auxiliary facts come together:

- Fact 2 collects the boxed formulas under one box.
- Fact 3 combines the single box with the diamond.
- Fact 4 transfers the contradiction from φ to $\Diamond\varphi$.

Without any one of these, the argument would not go through.

28 The Truth Theorem

We can now prove the central result: in the canonical model, a formula belongs to a world if and only if it is true at that world.

Theorem 28.1 (Truth Theorem)**.** *For every maximal consistent set $\omega \in W^c$ and every formula φ:*

$$\mathcal{M}^c, \omega \models \varphi \quad \Longleftrightarrow \quad \varphi \in \omega.$$

Proof. By induction on the complexity of φ.

Base case: $\varphi = p$ (propositional variable). $\mathcal{M}^c, \omega \models p$ iff $I^c(p, \omega) = T$ iff $p \in \omega$ (by definition of I^c). ✓

Propositional cases. These follow from the properties of maximal consistent sets (Lemma 23.1). For instance:

Case $\varphi = \psi_1 \wedge \psi_2$: $\mathcal{M}^c, \omega \models \psi_1 \wedge \psi_2$ iff $\mathcal{M}^c, \omega \models \psi_1$ and $\mathcal{M}^c, \omega \models \psi_2$ iff (by induction) $\psi_1 \in \omega$ and $\psi_2 \in \omega$ iff (by Lemma 23.1(iii)) $\psi_1 \wedge \psi_2 \in \omega$. ✓

The cases for $\neg$, $\vee$, $\rightarrow$ are analogous, using properties (ii), (iv), (v) respectively.

Modal case: $\varphi = \Box\psi$**.** This is the heart of the proof. We must show both directions.

Direction ($\Leftarrow$)*: if* $\Box\psi \in \omega$*, then* $\mathcal{M}^c, \omega \models \Box\psi$*.*
We need: for every $\omega' \in W^c$ with $\omega \, R^c \, \omega'$, $\mathcal{M}^c, \omega' \models \psi$.
Let ω' be such that $\omega \, R^c \, \omega'$. By definition of R^c: since $\Box\psi \in \omega$, we get $\psi \in \omega'$. By the inductive hypothesis: $\mathcal{M}^c, \omega' \models \psi$. ✓

Direction ($\Rightarrow$)*: if* $\mathcal{M}^c, \omega \models \Box\psi$*, then* $\Box\psi \in \omega$*.*
This is the hard direction. We prove it by *reductio ad absurdum*.
Suppose $\mathcal{M}^c, \omega \models \Box\psi$ but $\Box\psi \notin \omega$. Since ω is maximal consistent, $\neg\Box\psi \in \omega$, and therefore $\Diamond\neg\psi \in \omega$ (since $\neg\Box\psi$ and $\Diamond\neg\psi$ are provably equivalent in K).
By the **Existence Lemma** (Lemma 27.1): there exists a maximal consistent set ω' with $\omega \, R^c \, \omega'$ and $\neg\psi \in \omega'$.
By the inductive hypothesis: $\neg\psi \in \omega'$ implies $\mathcal{M}^c, \omega' \not\models \psi$.
But $\omega \, R^c \, \omega'$ and $\mathcal{M}^c, \omega \models \Box\psi$ imply $\mathcal{M}^c, \omega' \models \psi$. Contradiction. ✓

Modal case: $\varphi = \Diamond\psi$**.** *Direction* ($\Rightarrow$)*: if* $\mathcal{M}^c, \omega \models \Diamond\psi$*, then* $\Diamond\psi \in \omega$*.*
$\mathcal{M}^c, \omega \models \Diamond\psi$ means there exists ω' with $\omega \, R^c \, \omega'$ and $\mathcal{M}^c, \omega' \models \psi$. By induction, $\psi \in \omega'$.
Suppose $\Diamond\psi \notin \omega$. Then $\neg\Diamond\psi \in \omega$, so $\Box\neg\psi \in \omega$. By definition of R^c: $\neg\psi \in \omega'$. But $\psi \in \omega'$ and $\neg\psi \in \omega'$ contradicts the consistency of ω'.

Direction ($\Leftarrow$)*: if* $\Diamond\psi \in \omega$*, then* $\mathcal{M}^c, \omega \models \Diamond\psi$*.*
By the Existence Lemma: there exists ω' with $\omega \, R^c \, \omega'$ and $\psi \in \omega'$. By induction: $\mathcal{M}^c, \omega' \models \psi$. Since ω' is accessible from ω: $\mathcal{M}^c, \omega \models \Diamond\psi$. ✓

All cases are verified. The induction is complete. □

29 The Completeness Theorem

Theorem 29.1 (Completeness of labelled ND for K)**.** *If* $\Gamma \models_K \varphi$*, then* $\Gamma \vdash_K \varphi$*.*

Proof. Contrapositive: suppose $\Gamma \nvdash_K \varphi$. We show $\Gamma \not\models_K \varphi$.

Step 1: $\Gamma \cup \{\neg\varphi\}$ **is consistent.** If $\Gamma \cup \{\neg\varphi\} \vdash_K \bot$, then $\Gamma \vdash_K \neg\neg\varphi$, hence $\Gamma \vdash_K \varphi$ (by double negation elimination). This contradicts $\Gamma \nvdash_K \varphi$.

Step 2: Extend to a maximal consistent set. By Lindenbaum's Lemma (Lemma 24.1), extend $\Gamma \cup \{\neg\varphi\}$ to a maximal consistent set ω.

Step 3: Build the canonical model. The canonical model $\mathcal{M}^c = \langle W^c, R^c, I^c \rangle$ (Definition 25.1) is a Kripke model, and $\omega \in W^c$.

Step 4: Apply the Truth Theorem. By the Truth Theorem (Theorem 28.1):
- For each $\gamma \in \Gamma$: $\gamma \in \omega$ (since $\Gamma \subseteq \omega$), so $\mathcal{M}^c, \omega \models \gamma$.
- $\neg\varphi \in \omega$ (since $\neg\varphi \in \Gamma \cup \{\neg\varphi\} \subseteq \omega$), so $\mathcal{M}^c, \omega \models \neg\varphi$, i.e., $\mathcal{M}^c, \omega \not\models \varphi$.

Therefore $\Gamma \cup \{\neg\varphi\}$ is satisfiable, so $\Gamma \not\models_K \varphi$. □

30 Comparison: Tableau vs. Canonical Model Completeness

	Tableau (Part II)	**Labelled ND** (Part IV)
Worlds	Labels on the open branch.	Maximal consistent sets.
Accessibility	iRj statements on the branch.	Defined via $\Box$: $\omega R\omega'$ iff $\Box\psi \in \omega \Rightarrow \psi \in \omega'$.
Valuation	T-signed atoms on the branch.	Membership: $p \in \omega$.
Key tool	The branch is already finished — just read off the model.	Lindenbaum + Existence Lemma: construct the model from scratch.
Difficulty	Low.	High (Existence Lemma uses Facts 1–4).
Advantage	Algorithmic: directly produces a finite counter-model.	General: works for any proof system, not just tableaux.