First-Order Logic by Raymond Smullyan

Notes by Aditya Dwarkesh

April 2021

1 Preliminaries

1.1 Foreword on trees

Definition 1.1 (Unordered trees)**.** An unordered tree τ is an ordered triplet $\{S, l, R\}$.

1. S is a set of elements called *points*

2. l is a function which assigns each point a positive integer $l(x)$ called the *level* of x

3. R is a relation defined on S such that we read xRy as "x is a *predecessor* of y" or "y is a *successor* of x", which satisfies the following conditions:

   (a) There is a unique point of level 1, called the *origin* of the tree.

   (b) Every point other than the origin has a unique predecessor.

   (c) For any points x, y, if y is a successor of x, then $l(y) = l(x) + 1$.

If x has no successor, is an *end point* of the tree; if one successor, a *simple point*, and if multiple, a *junction point.*
A *path* is a denumerable sequence of points wherein each term is the predecessor of the next. A *maximal path/branch* is one whose last term is an end point, or a path which is infinite.
For any point x, there exists a unique path P_x whose last term is x. If y lies on P_x, we say that y *dominates* x. Further, if $y \neq x$, we say that x *lies below* y. Two points are *comparable* if one dominates the other. y is *between* x and z if it is above one and below the other.

Definition 1.2 (Ordered trees)**.** An ordered tree is a quadruplet $\{S, l, R, \theta\}$ wherein the first three elements are as explicated above and θ is a function which assigns a sequence $\theta(z)$ to each junction point z which satisfies the following conditions:

1. It contains no repetitions

2. Its set of terms consists of all the successors of z

The *nth* successor of z refers to the *nth* term of $\theta(z)$. The successor of a simple point will be spoken of as a *sole* successor.
If each point has only finitely many successors, the tree is said to be *finitely generated.* If the tree has only finitely many points, it is said to be *finite*; otherwise, *infinite.*
Ordered trees in which each junction point has exactly 2 successors are called *dyadic trees.* The first successor

(in the sequence) is the *left successor*, and the second is the *right successor*.

Exercise. Let a be the junction point and x', y' be the successors by virtue of which x is to the left of y. Further, let a' and y'', z'' be the same corresponding to y being to the left of z. Either y' dominates y'' or vice-versa; suppose the former without loss of generality. This implies that y' dominates z, and the proof is concluded with the junction point a (had we taken y'' to dominate y', the relevant junction point would have been a').

1.2 Formulas of propositional logic

Definition 1.3 (Symbols)**.** The following symbols form the object language:

1. $\neg, \wedge, \vee, \implies$ (Logical connectives)
2. $\{p_1, p_2 ...\}$ (Propositional variables)

Definition 1.4 (Formula)**.** The set of formulas is the following:

1. Every propositional variable is a formula.
2. If X is a formula then so is the ordered pair $\langle \neg, X \rangle$, the negation of X.
3. If X, Y are formulas, then for each of the binary connectives, the ordered triplet $\langle X, b, Y \rangle$ is a formula.

Theorem 1.1 (**Metatheorem**)**.** *[Uniqueness of decomposition] For every formula X, one and only one of the following conditions hold:*

1. *X is a propositional variable.*
2. *There is a unique formula Y such that $X = \neg Y$.*
3. *There is a unique pair X_1, X_2 and a unique binary connective b such that $X = X_1 b X_2$.*

Proof. A formula is either a propositional variable, an ordered pair or an ordered triplet. These are three are mutually exclusive. Furthermore, an ordered pair (resp. triplet) uniquely determines its first and second (resp. first, second and third) elements. Hence, proved. $\square$

Definition 1.5 (Subformula)**.** The following characterizes the set of subformulas:

1. Propositional variables have no immediate subformulas.
2. $\langle \neg, X \rangle$ has X as an immediate subformula and no others.
3. $\langle X, b, Y \rangle$ has X, Y as immediate subformulas and no others.
4. If X is an immediate subformula of Y or identical to Y, then X is a subformula of Y.
5. If X is a subformula of Y and Y is a subformula of Z then X is a subformula of Z.

An *atomic* formula has no subformulas; otherwise, the formula is *compound*.

Definition 1.6 (Degree)**.** The following characterizes the degree of a formula:

1. A propositional variable is of degree 0.

2. If X is of degree n, then $\langle \neg, X \rangle$ is of degree $n+1$.
3. If X, Y are of degrees n_1, n_2 then $\langle X, b, Y \rangle$ is of degree $n_1 + n_2 + 1$.

Definition 1.7 (Formation tree)**.** A formation tree τ for a formula X is an ordered dyadic tree whose points are (occurrences of) formulas and whose origin is (an occurrence of) X such that

1. Each end point is (an occurrence of) a propositional variable
2. Each simple point is of the form $\langle \neg, Y \rangle$ and has (an occurrence of) Y as its sole successor
3. Each junction point is of the form $\langle X, b, Y \rangle$ and has (occurrences of) X, Y as respective left and right successors.

The principle of induction:
Let S be a set of formulas and P be a certain *property* of formulas. If

1. Every element of S of degree 0 has P
2. For every element X of S of positive degree, all the elements S of degree less than X have P

then X also has P.

Note: This is a meta-logical principle.

1.3 Boolean valuations and truth sets

Definition 1.8 (Boolean valuation)**.** A *valuation* is a function from a set of formulas to *truth-values* $F : S \to \{t, f\}$.
A valuation v of E is called a *Boolean* valuation if for every X, Y in E, the following conditions hold:

1. The formula $\langle \neg, X \rangle$ receives the value f if X receives the value t and t if X receives the value f.
2. The formula $\langle X, \wedge, Y \rangle$ receives the value t if both X, Y receive the value t; otherwise, it receives the value f.
3. The formula $\langle X, \vee, Y \rangle$ receives the value t if at least one of X, Y receive t; otherwise, it receives the value f.
4. The formula $\langle X, \implies, Y \rangle$ receives the value f if X, Y receive the values t, f respectively; otherwise, it receives the value t.

If S_1 is a subset of S_2 and v_1, v_2 are respective valuations then we say v_2 is an *extension* of v_1 if they both agree on the smaller set.
An *interpretation* of a set E means an assignment of truth values to all the variables which occur in any of the elements of W.

Theorem 1.2 (Metatheorem). *An interpretation on a set E can be extended to exactly one Boolean valuation of E.*

Proof. It is easily provable by induction that an interpretation on a set E can be extended to *at most* one Boolean valuation of E (assume there exist two and show that they must agree on all its elements).
To prove that it can be extended to *at least* one, consider an interpretation v_0 of a single formula X. It is easy to show using induction that there exists only one way of assigning truth values to all subformulas of X such that the atomic subformulas are assigned the same truth values as under v_0, and the compound subformulas have their truth values determined by the above rules. Since X is also a subformula of itself, we can define what it means for it to be true under an interpretation. Next, consider an interpretation for the set E. We now let v be the valuation which assigns to each element of E its truth value under the interpretation v_0. This yields the required valuation and completes the proof. □

Definition 1.9 (Tautology)**.** X is a *tautology* if and only if X is true in all Boolean valuations (or under every interpretation) of E.

Definition 1.10 (Satisfiability)**.** A formula X is called (truth-functionally) *satisfiable* if and only if X is true in at least one Boolean valuation. A set is satisfiable iff there exists at least one Boolean valuation in which every element of S is true.

Definition 1.11 (Truth-functional implication & equivalence)**.** *S truth-functionally implies X* if X is true in every Boolean valuation which satisfies S.
X, Y are *truth-funcionally equivalent* if they are true in the same Boolean valuations. (Note: This is true iff $X \leftrightarrow Y$ is a tautology.)

Definition 1.12 (Truth set)**.** A set S is called *saturated* or a *truth set* if it satisfies the following conditions:

1. $\langle \neg, X\rangle \in S \leftrightarrow \langle X\rangle \notin S$.
2. $\langle X \wedge Y\rangle \in S \leftrightarrow \langle X\rangle, \langle Y\rangle \in S$.
3. $\langle X \vee Y\rangle \in S \leftrightarrow \langle X\rangle$ or $\langle Y\rangle \in S$.
4. $\langle X \implies Y\rangle \in S \leftrightarrow \langle X\rangle \notin S$ or $\langle Y\rangle \in S$.

If S is the set of sentences true under an arbitrary valuation v, then (v is a Boolean valuation) $\leftrightarrow$ (S is saturated).
X is a tautology iff it is an element of the intersection of every truth set.
X is satisfiable iff it is an element of the union of every truth set.
S truth-functionally implies X iff X belongs to every truth set which includes S.

Exercise 1.

1. Suppose $\neg X_1 \to t$. If $X_2 \to t$, then $X_2 \implies X_1 \to f$, contrary to our assumption of it being a tautology. Therefore, $\neg X_2 \to t$. We may prove the converse similarly and conclude $\neg X_1 \leftrightarrow \neg X_2$.
2. Suppose $X_1 \wedge Y \to t$. Now, $\neg(X_2 \wedge Y) \to t \implies X_2 \to f \implies X_1 \to f$. But since X_1 was t, we conclude that $X_2 \wedge Y \to t$. The other way round may be proven similarly.
3. Suppose $X_1 \implies Y \to t$. If $X_2 \implies Y \to f$, then $\neg X_2 \vee Y \to f \implies \neg X_1 \vee Y \to f$ (from another part to the exercise) $\implies X_1 \implies Y \to f$. This gives us the contradiction required.
4. etc.

By induction from the above results, we can show that for any formula Z containing X_1, replacing any number of occurrences of it with X_2 yields an equivalent formula.

Exercise 2. $(X \implies f) \to t$ is equivalent to $\neg X \to t$ or $f \to t$. Since the latter can never be true, we have $\neg X \to t$. We may work our way back in a similar manner. All the others are similarly provable.
Using induction from the rest of the parts, we can show that a formula with propositional constants is either equivalent to some other one without the, or to t or f.

Exercise 3. We prove this by induction.
This is true for every formula of degree 0 by definition.
Let $X = Y_1 \vee Y_2$, where Y_1 and Y_2 are formulae in disjunctive normal form. Then it follows directly that so is X.
Let $X = Y \wedge Y_2$, where Y is a propositional variable and Y_2 is a compound formula in disjunctive normal form. An application of de Morgan's law yields the required form for X. If we now let Y be compound as well, then iterated application of the law yields the desired form.
Let $X = \neg Y$, where Y is a formula in disjunctive normal form. $\neg Y = \neg C_1 \wedge ... \neg C_n$. $C_1 = \neg p_1 \vee ... \neg p_n$. This completes the transformation.

Exercise 4.

1. $X \wedge Y \leftrightarrow \neg(\neg X \vee \neg Y)$.
2. $X \implies Y \leftrightarrow \neg X \vee Y$. Use the fact that $p \vee q \leftrightarrow \neg(\neg p \wedge \neg q)$.
3. Done above.
4. From 1 and 2, $X \wedge Y \leftrightarrow \neg(X \implies \neg Y$.
5. Replace X with $\neg X$ in 2.

Exercise 5.

1. $X|Y \leftrightarrow \neg X \vee \neg Y$.
2. $X \wedge Y \leftrightarrow \neg(X|Y)$; $X \implies Y \leftrightarrow X|\neg Y$.
3. $X \downarrow Y \leftrightarrow \neg X \wedge \neg Y$.
   For $|, \downarrow$, the remaining can be done using the above results.

Extra problem: Prove that $\{|\}, \{\downarrow\}$ are the only singleton sets which are functionally complete.

2 Analytic tableaux

2.1 The method of tableaux

Definition 2.1 (Signed formulas)**.** Under any interpretation, a *signed formula* TX is called *true* if X is true and false otherwise. FX is true if X is false and false otherwise.

A *conjugate* of a signed formula is the result of changing T to F (or vice-versa).

Notation.
α stands for any formula of the following type:

1. $T(X \wedge Y)$
2. $F(X \vee Y)$
3. $F(X \implies Y)$
4. $T\neg X$
5. $F\neg X$

These are formulas of the *conjunctive* type, with *direct* consequences.
β stands for any formula of the following type:

1. $F(X \wedge Y)$
2. $T(X \vee Y)$
3. $T(X \implies Y)$

These are formulas of the *disjunctive* type, which *branch*.

Exercise. Show that if the following conditions hold, S is a truth set:

1. If $\alpha \in S$, so are its direct consequences α_1, α_2
2. If $\beta \in S$, so is at least one of its branches β_1, β_2.
3. Exactly one of $X, \neg X$ belongs to S.

Note: Sets satisfying 1 and 2 are called *downward closed*, while ones satisfying their converses are called *upwards closed*.
In the following, read $\neg$ as conjugation.
Suppose we have $\alpha_1, \alpha_2, \neg\alpha \in S$. $\neg\alpha$ is a formula of type β, and by 2 we infer that at least one of $\neg\alpha_1, \neg\alpha_2$ must be in S. Here we have a contradiction.
The case for the disjunction may be proven similarly ($\neg\beta$ becomes a formula of type α).
Note: Unlike the principle of induction, contradiction isn't being used as a meta-logical principle here.

Definition 2.2 (Analytic tableaux)**.** An *analytic tableau* for X is an ordered dyadic tree, whose points are (occurences of) formulas, which is constructed as follows:
We start by placing X at the origin. Now suppose τ is a tableau for X which has already been constructed; let Y be an end point. Then we may extend τ by either of the following two operations:

1. If some α occurs on the path P_Y, then we may adjoin either α_1 or α_2 as the sole successor of Y.
2. If some β occurs on the path P_Y, then we may simultaneously adjoin β_1 as the left successor of Y and β_2 as the right successor of Y.

Given two trees of the above kind τ_1 and τ_2, we say the latter is a *direct extension* of the former if it can be obtained by only one application of 1 or 2.
τ is a tableau for X if and only if there exists a finite sequence $\{\tau_1, ...\tau_n\} = \tau$ such that τ_1 is a one-point tree whose origin is X and such that for each $i < n, \tau_{i+1}$ is a direct extension of τ_i.

A branch θ of a tableau is *closed* if it contains some signed formula and its conjugate. τ is called closed if every branch of τ is closed.
A branch θ of a tableau is *complete* if for every α occurring in it, both α_1, α_2 occur and for every β, one of β_1, β_2 occur. τ is called complete if every branch of τ is complete.
By a *proof* of X is meant a closed tableau for FX.

Exercise. Prove that $(p \vee (q \wedge r)) \implies ((p \vee q) \wedge (p \vee r))$ is a tautology.

1. $F[(p \vee (q \wedge r)) \implies ((p \vee q) \wedge (p \vee r))]$
2. $T[(p \vee (q \wedge r))]$ (1)
3. $F[((p \vee q) \wedge (p \vee r))]$ (1)
4. $F[(p \vee q)]$ (3)
5. $F[(p \vee r)]$ (3)
6. $F[q]$ (4)
7. $F[p]$ (4,5)
8. $F[r]$ (5)
    - $T[p]$ (1)

    Closed.

    - $T[(q \wedge r)]$ (1)
    - $T[q]$

    Closed.

This concludes the proof.

2.2 Consistency and completeness of the system

Definition 2.3. A system is consistent if no formula and its negation are both provable in it.

Theorem 2.1. *Any formula provable by the tableau method is a tautology.*

Proof. Call a tableau true under v_0 iff at least one branch is true under v_0; a branch true if every term of it is true under v_0.
Let τ_2 be an immediate extension of τ_1. If τ_1 is true, so is τ_2; for either the extension was independent of the true branch of τ_1, or it was by an operation on that branch. If the former, we are done. If the latter, if the extension was by operation 1, then τ_2 has as branches (θ_1, α_1) and (θ_2, α_2). But since α was true, both α_1 and α_2 are, and so both the branches are true, and so τ_2 is true. The case may be proven similarly for operation 2, and we are done.
Now, using induction, we may say that for any tableau, if the origin is true under an interpretation, then so is the tableau. But since the origin of a closed tableau cannot be true under any interpretation, it is the negation of a tautology; and so, every formula provable by the tableau method must be a tautology. □

Since no formula and its negation can both be tautologies, it follows that the tableau method is consistent.

Definition 2.4 (Completeness)**.** A system is complete if every tautology is provable in it.

Definition 2.5 (Hintikka set)**.** We call S a *Hintikka set* if it satisfies the following three conditions:

1. No signed variable and its conjugate both are in S
2. If $\alpha \in S$, then α_1 and $\alpha_2 \in S$
3. If $\beta \in S$, then β_1 or $\beta_2 \in S$.

We shall also call such sets *downward saturated.* If the set of terms in a denumerable sequence is a Hintikka set, we shall call it a *Hintikka sequence.*
The set of terms of a complete open branch θ of τ is a Hintikka set.

Theorem 2.2 (Hintikka's lemma)**.** *Every Hintikka set is satisfiable.*

Proof. Let S be a set of signed formulae. Assign an interpretation to the set in the following manner: if $Fp \in S$, give p the truth value false; otherwise, give p the truth value true. (Since no Tp and Fp both can occur in S, this is possible.)
It follows immediately that every *signed variable* is true under this interpretation. Now, consider an element X of degree greater than 0 (making it either α or β), and suppose every element of a lower degree is true.
If it is α, then both α_1, α_2 must also be in S, since it is a Hintikka set. But by the induction hypothesis, both α_1 and α_2 are true. Thus, α must be true. We may prove the same for β similarly.
Thus, we have found an interpretation in which every element of S is true, proving that S is satisfiable. □

It follows as a corollary that every finite Hintikka set is satisfiable, and so, that any complete open branch of any tableau is (simultaneously) satisfiable.

Theorem 2.3. *For any tautology X, every completed tableau starting with FX must close.*

Proof. From theorem 2.2, if τ is open, FX is satisfiable and so X cannot be a tautology. Therefore, for a given tautology X, every completed tableau must close. Since for any formula X there exists a completed tableau with X at the origin, this holds for any tautology. □

Theorem 2.4. *A finite set S is unsatisfiable iff there exists a closed tableau for S.*

Proof. Recall that a set is satisfiable iff there exists at least one Boolean valuation in which every element is true. But the definition of a Boolean valuation tells us that if every element is true, then so is $X_1 \wedge ... X_n$. However, the tableau drawn corresponds to just this formula, and since it is closed, we infer that it is not true. Therefore, the set is unsatisfiable if there exists a closed tableau for it. The converse may be proven similarly. □

Definition 2.6. A tableau is *atomically* closed if every branch contains some *atomic* element and its conjugate.

Theorem 2.5. *If S is unsatisfiable, then there exists an atomically closed tableau for S.*

Proof. Suppose that a completed tableau for S has an atomically open branch. Since this branch is a Hintikka set, it follows that it is satisfiable, and thus, that so is S. The theorem follows from this. □

A corollary to this is that if there exists a closed tableau for S, then there exists an atomically closed tableau for S.

Exercise 1.

1. This is true trivially for degree 0. Next, suppose X is a β. Then, at least one of β_1, β_2 is in the tableau. But since the conjugate of X is also there, the conjugates of the above two are also there. From the induction hypothesis, there is an atomically closed tableau for a set which has both β_1 and its conjugate (or 2). The other case may be dealt with similarly. Hence, proved.

2. Let X and its conjugate close an arbitrary branch of the tableau. From 1, we know that the branch can be closed atomically. Hence, proved.

Exercise 2. We have to show that X is truth functionally equivalent to $C_1 \vee ...C_n$. We do this by first showing that X is truth functionally equivalent to $B_1 \vee ...B_n$, and then that B_i is to C_i.
The first part true because an interpretation under which any one branch is true is immediately an interpretation under which X is true, and vice-versa.
For an open branch B_i to be true under a valuation, all of its elements have to be true. This is equivalent to the truth of C_i.
Hence, proved.

Exercise 3. We can reproduce the original 8 rules by replacing the other propositional connectives with Sheffer's symbol. For the joint denial connective, the rules will be (presented roughly)

1. $TX \downarrow Y \implies \neg X \wedge \neg Y$
2. $FX \downarrow Y \implies X \vee Y$

3 Compactness

3.1 Analytic proofs of the compactness theorem

Definition 3.1 (Consistency)**.** A (denumerable) set is consistent if every finite subset of it is satisfiable. If a set is not consistent, it is *inconsistent.*

A *finitely generated tree* is a tree in which each point has only finitely many successors.

Lemma 3.1 (König)**.** *Every finitely generated tree τ with infinitely many points must contain at least one infinite branch.*

Proof. Call a point *good* if it has infinitely many descendants and *bad* if finitely many. Since all points are dominated by the origin, the origin is good.
A good point must have at least one good successor. Thus the origin a_0 has at least one good successor a_1, which in turn has at least one good successor a_2...this generates an infinite branch $a_0...a_n...$ (using the axiom of choice, if the tree be unordered). □

Theorem 3.2 (Compactness)**.** *If S is consistent, then S is satisfiable.*

Proof. Let S be a consistent set arranged in a denumerable sequence $X_1, X_2...X_n, ...$. Two methods are given.

1. Since S is consistent, a complete tableau for X_1 cannot close. Attach X_2 to every open branch and complete the tableau again. Since S is consistent, a complete tableau for X_1, X_2 cannot close and so we will have at least one open branch again.
   This process is iterated ad infinitum, and since at no stage can the tableau close, we have an infinite tree. By König's lemma, this must have at least one infinite branch. The branch must be open, contain all the X_i, and be a Hintikka set. By Hintikka's lemma, it is satisfiable. It follows from this that S is satisfiable.

2. Note that consistency satisfies the following conditions:

   - No set containing a propositional variable and its negation is consistent.
     This is immediate.
   - If $\{S, \alpha\}$ is consistent, so is $\{S, \alpha_1, \alpha_2\}$.
     Suppose $\{S, \alpha_1, \alpha_2\}$ is inconsistent. Then, $\{S_1, \alpha_1, \alpha_2\}$ is unsatisfiable, where S_1 is a subset of S and disjoint from α_1, α_2. This means that $\{S_1, \alpha\}$ is unsatisfiable, and therefore, that $\{S, \alpha\}$ is inconsistent. The result follows.
   - If $\{S, \beta\}$ is consistent, then so is at least one of $\{S, \beta_1\}$ and $\{S, \beta_2\}$.
     This can be proven in a similar manner as the above.

   Now, we wish to construct a Hittika sequence whose terms include each X_i. Let us take X_1 as the first term.
   Suppose now that at the nth stage we have a finite sequence θ_n as $X_1, Y_2...Y_{n+i}, i \geq 0$, such that $\{S, Y_2...Y_{n+i}\}$ is consistent. The next extension is done as follows:

   - If Y_n is of the type α, then θ_{n+1} is $X_1, Y_2...Y_{n+i}, \alpha_1, \alpha_2, X_{n+1}$. From the second condition on consistency, $\{S, \theta_{n+1}\}$ is also consistent.
   - If Y_n is of the type β, then from the third condition on consistency, either $\{S, \theta_n, \beta_1, X_{n+1}\}$ or $\{S, \theta_n, \beta_2, X_{n+1}\}$ is consistent. We choose θ_{n+1} accordingly.

- If Y_n is a propositional variable, then we merely adjoin X_{n+1} to θ_n.

This gives us the required Hintikka sequence, and as with the first proof, satisfiability follows.

The statement is trivial for finite sets. □

Exercise. If there are k propositional variables in S, there are at most 2^k possible interpretations. Let X_1 be the element of S true under v_1, and so on until X_n is the one true under v_{2^k}. It is clear that their disjunction is satisfied (true) under all possible Boolean valuations, and it is thus a tautology.

3.2 Maximal consistency: Lindenbaum's Construction

A *proper extension* of a set is a superset which contains at least one element not in it.

Definition 3.2. A set of formulas is called *maximally consistent* if it is consistent and if no proper extension of it is consistent.

Lemma 3.3. *Any maximally consistent set is a truth set.*

Proof. First, we prove that if S is consistent, then for any formula X at least one of $\{S, X\}$ and $\{S, \neg X\}$ is consistent. For suppose both to be inconsistent; then there must be finite subsets S_1, S_2 of S such that $\{S_1, X\}$ and $\{S_2, \neg X\}$ are unsatisfiable. Let $S_3 = S_1 \cup S_2$. Then both $\{S_3, X\}$ and $\{S_3, \neg X\}$ are unsatisfiable; and so S_3 is unsatisfiable, contradicting the consistency of S.
It follows directly from this that if M is maximally consistent, then for any X either $X \in M$ or $\neg X \in M$.
Now, let M be a maximally consistent set. This means that for any formula X, at least one of $X, \neg X$ lies outside M and also that for any formula X, at least one of $X, \neg X$ lies inside M. This satisfies the first condition for truth sets. Next, suppose $\alpha \in M$. Then, $\neg\alpha_1 \notin M$, since $\{\alpha, \neg\alpha_1\}$ is not satisfiable; and so, $\alpha_1 \in M$. Similarly, $\alpha_2 \in M$. Conversely, let $\alpha_1 \in M, \alpha_2 \in M$. Since $\{\neg\alpha, \alpha_1, \alpha_2\}$ is not satisfiable, we have $\alpha \in M$. This completes the proof. □

A property P of a set is said to be of *finite character* if for any set S, it has the property P iff all finite subsets of S have the property P.
Note that consistency is of finite character.

Theorem 3.4 (Tukey's lemma for the denumerable case)**.** *For any denumerable universe U and any property P of subsets of U of finite character any set S (of elements of U) having property P can be extended to a maximal subset of U having property P.*

Proof. Arrange the elements of U in some denumerable sequence $Y_1, Y_2 ... Y_n, ...$ and define a denumerable sequence of sets in the following manner:
We set $S_0 = S$. From S_n, we define S_{n+1} as follows: $S_{n+1} = S_n \cup \{Y_{n+1}\}$ if the right hand side has property P; otherwise, $S_{n+1} = S_n$.
It is immediate that $S_0 \subseteq S_1 ... S_n \subseteq S_{n+1} ...$, and that each S_i has property P. We claim that $M = S_0 \cup S_1 ... S_n \cup S_{n+1} ...$ is a maximal set having property P.
Let K be any finite subset of M. It follows that K must be a subset of S_i for some i. Since S_i has P and P is of finite character, K has P. Since this is true of any finite subset K, we conclude also that M has P.
Now, take any Y_i such that $M \cup \{Y_i\}$ has P. Since P is of finite character, so does $S_i \cup \{Y_i\}$. Then, $Y_i \in S_{i+1}$, and so $Y_i \in M$. This concludes the proof. □

The general version of the lemma (for any arbitrary universe U) is equivalent to the axiom of choice.

Corollary 3.4.1 (Lindenbaum's theorem)**.** *Every consistent set can be extended to a maximally consistent set.*

Note that the above construction isn't *analytic* in the sense of cut-free; the elements of the set constructed aren't limited to subformulas (or negations of subformulas) of the original set.

Definition 3.3 (Completeness of a set)**.** S is said to be complete if every formula or its negation is in S.

Exercise. Let M be a consistent complete set. Let us try to extend it by adding a formula X. By completeness, M has either X or $\neg X$; by consistency, it has only one of them. If it has X, then M isn't extended at all; if it has $\neg X$, then its extension is no longer consistent. Hence, proved.

3.3 An Analytic Modification of Lindenbaum's Proof

Call Y a *direct descendant* of X if either X is some α and Y is α_1 or α_2, or correspondingly for β.
Call Y a *descendant* of X if there exists a finite sequence beginning with X and ending with Y such that each term of the sequence (other than the first) is a direct descendant of the preceding term.
Let S^0 be the set of all descendants of elements of S.

Theorem 3.5. *Every maximally consistent subset of S^0 is a Hintikka set.*

Proof. Let M be a maximally consistent subset of S^0.
Since it is consistent, it contains no variable and its negation. This satisfies the first condition for Hintikka sets. Next, suppose $\alpha \in M$. This means that $\alpha_1, \alpha_2 \in S$. Then, since M is consistent, so is $\{M, \alpha\}$, and so is $\{M, \alpha_1\}$. By maximality, we have $\alpha_1 \in M$. The proof for α_2 and the third condition for Hintikka sets can be done similarly. □

3.4 The Compactness theorem for Deducibility

Definition 3.4 (Deducibility)**.** X is *deducible* from a set S if there are finitely many $X_1, X_2...X_n \in S$ such that $(X_1 \wedge X_2... \wedge X_n) \implies X$ is a tautology.

Theorem 3.6 (Compactness: second form)**.** *If X is true in all Boolean valuations which satisfy S, then X is deducible from S.*

Proof. By hypothesis, $\{S, \neg X\}$ is unsatisfiable. By the compactness theorem, some finite subset of it (which must include $\neg X$, since S is satisfiable) must be unsatisfiable. Therefore, $\{X_1, X_2...X_n, \neg X\}$ is unsatisfiable. It follows that $(X_1 \wedge X_2... \wedge X_n) \implies X$ is a tautology.
Hence, proved. □

A set is called *deductively closed* if every formula deducible from S lies in S.

Exercise. Prove that a consistent deductively closed set is the intersection of all its complete consistent extensions (Tarski's theorem).
Proof. We have to prove that $S = S_1 \cap S_2...$, where S is a deductively closed set and $S_1, S_2...$ are its complete consistent extensions.
First, suppose $X \in S$. We wish to show that $X \in$ each of $S_1, S_2....$ This follows because $\neg X \notin S$ and each of them are extensions, so $X \in S \implies X \in S_n$. Therefore, $S \subseteq S_! \cap S_2...$
Next, suppose $X \in$ each of $S_1, S_2....$ If X is not deducible from S, then both $S \cup \{X\}$ and $S \cup \{\neg X\}$ are possible as consistent extensions of S, and so this would mean $X \notin S_1 \cap S_2....$ Thus, X must be deducible from S.
Hence, proved.

4 First-Order Logic: Preliminaries

4.1 Formulas of Quantification Theory

Definition 4.1 (Symbols)**.** The following symbols form the object language:

1. Logical connectives (as defined in 1.2)
2. Quantifiers (universal, $\forall$ reading "for all"; existential, $\exists$ reading "there exists")
3. Individual variables (denumerably many)
4. Individual parameters (denumerably many)
5. $n - ary$ predicates (for each positive integer n)

Definition 4.2 (Atomic formula)**.** An atomic formula is an $(n+1)-$tuple $Pc_1 ... c_n$ where P is a predicate of degree n and $c_1, ... c_n$ are individual symbols.

Definition 4.3 (Formula)**.** A is a formula iff there is a finite sequence of expressions which terminates with A such that each term is either an atomic formula or is the negation, conjunction, disjunction or conditional of earlier term(s), or is the existential or universal quantification of an earlier term (with respect to some variable x).

A *pure formula* is a formula without any individual parameters.

Definition 4.4 (Degree)**.** The following characterizes the degree of a formula.

1. Every atomic formula is of degree 0.
2. $d(\neg A) = d(A) + 1$
3. $d(A \wedge B) = d(A \vee B) = d(A \implies B) = d(A) + d(B) + 1$
4. $d((\forall x)A) = d((\exists x)A) = d(A) + 1$

Definition 4.5 (Substitution)**.** The substituted formula A^x_a is defined from A by the following induction scheme:

1. If A is atomic, then A^x_a is the result of substituting a for every occurence of x in A.
2. $[AcB]^x_a = A^x_a c B^x_a$, where c is any one of the binary connectives
3. $[\neg A]^x_a = \neg[A^x_a]$
4. $[(Cx)A]^x_a = (Cx)A$, $[(Cx)A]^y_a = (Cx)[A^y_a]$, where C is any one of the quantifiers.

A *closed formula* or a *sentence* is a formula such that $A = A^x_a$ for every variable x and parameter a.

The occurrence of a variable in a formula is *bound* if it is within the scope some occurrence of $\forall x$ or $\exists x$, or is immediately preceeded by a quantifier. Otherwise, we say that the variable is *free*.

Definition 4.6 (Subformula)**.** The following characterizes the set of subformulas (in the sense of first-order logic):

1. $\langle \neg, X \rangle$ has X as an immediate subformula and no others.

2. $\langle X, b, Y \rangle$ has X, Y as immediate subformulas and no others (for a binary connective b).

3. For any parameter a, variable x and formula A, A^x_a is an immediate subformula of $(\forall x)A$ and $(\exists x)A$.

4. If X is an immediate subformula of Y or identical to Y, then X is a subformula of Y.

5. If X is a subformula of Y and Y is a subformula of Z then X is a subformula of Z.

Definition 4.7. A *formation tree* (in the sense of first-order logic) is a tree in which each end point is atomic, and every other point satisfies one of the following conditions:

1. It is of the form AbB and has A, B for its first and second successors, and has no other successors

2. It is of the form $\neg A$ and has A for its sole successor

3. It is of the form $(Cx)A$ for some quantifier C and has $A^x_{a_1}, A^x_{a_2}...$ as its successors.

Note that, since we have denumerably many parameters, formation trees for quantification theory are not finitely generated.

4.2 First-order Valuations and Models

Definition 4.8 (First order valuation)**.** Let U be any non-empty set, which we shall call a *universe of individuals.* A U-formula is like a formula, except that it has elements from U instead of individual parameters. Let E^U be the set of all closed U-formulas, and let v be an assignment of truth values to all its elements.
If, for every $A \in E^U$ and every variable x, the following hold:

1. v is a Boolean valuation of E^U

2. $(\forall x)A$ is true iff for every $k \in U, A^x_k$ is true under v

3. $(\exists x)A$ is true iff for at least one element $k \in U, A^x_k$ is true under v.

then v is a first-order valuation with respect to the universe U.

Definition 4.9 (First-order truth set)**.** A subset S of E^U is a first-order truth set with respect to the universe U if it satisfies all the conditions of a propositional truth set, as well as the following:

1. $(\forall x)A$ belongs to S iff for every $k \in U, A^x_k$ belongs to S

2. $(\exists x)A$ belongs to S iff for at least one element $k \in U, A^x_k$ belongs to S.

It follows that S is a first-order truth set iff the characteristic function of S is a first-order valuation.

An *atomic valuation* is an assignment of truth values to all the atomic elements of E^U.

Theorem 4.1 (Metatheorem). *Any atomic valuation v_0 of E^U can be extended to exactly one first-order valuation v of E^U.*

Proof. A *valuation* tree for A is a formation tree along with a truth value at every points, being determined by the truth values of its successors. For any atomic valuation, one can show by induction that there exists only one valuation tree for A, and A thus receives a unique value. □

Note: Since a point can have infinitely many successors in general, formalizing the induction argument above will require a stronger language than first-order logic.

Definition 4.10 (First-order interpretation)**.** An interpretation is a function which maps an n-place predicate P to an n-place relation P^* between elements of U.

An atomic U-sentence $Pa_1..$ is *true under* I if the elements $a_1...$ stand in the relation P^*. Thus, an interpretation I induces a unique atomic valuation v_0.

Definition 4.11 (Models)**.** An interpretation in which every element of a given set is true is called a *model* for the set.

A pure formula A is *valid* if for every universe U, A is true under every first-order valuation of E^U.
A pure formula is *satisfiable* if for at least one universe U there is at least one first-order valuation of E^U under which A is true.
A sentence with parameters $a_1...a_n$ is satisfiable in a universe if there exists at least one interpretation in which there exists at least one $n-$tuple of elements in U such that the formula produced by substituting them for the parameters is true under I. Validity is defined similarly, but with the condition being for every $n-$tuple under every interpretation.

Exercise 1. Show that the validity or satisfiability of a formula in a universe U depends only on the cardinality of U.
Answer 1.

Exercise 2. Show that if a formula is satisfiable/valid in U, it is satisfiable/valid in any superset/subset (resp.) of U.
Answer 2.

Answer 3. $\forall x \neg R(x,x) \wedge \forall x \forall y \forall z (R(x,y) \wedge R(y,z) \implies R(x,z)) \wedge \forall x \exists y R(x,y)$.

Answer 4. If $\neg A$ is not satisfiable, there is no interpretation under which $\neg A$ is true, i.e., $\neg A$ is false under every interpretation. Since every first-order valuation also satisfies the basic properties of a Boolean valuation, this means that A is true under every interpretation, i.e., A is valid. The converse and the second half can be proven similarly.

Exercise 5.

1. Show that $A(a_1...a_n)$ is valid iff $(\forall x_1)...(\forall x_n)A(x_1...x_n)$ is valid.
2. Show that $A(a_1...a_n)$ is satisfiable iff $(\exists x_1)...(\exists x_n)A(x_1...x_n)$ is valid.

4.3 Boolean valuations vs. First-order valuations

Call a sentence a *Boolean atom* iff it is either an atomic sentence or of the form $(Cx)A$, where C is a quantifier. Consider the universe V whose elements are the individual parameters. The following statements hold:

1. Every first-order valuation of E^V is a Boolean valuation but not vice-versa.
2. Any assignment of truth values to all the Boolean atoms of E^V can be extended to exactly one Boolean valuation of E^V.
3. If S is an infinite subset of E^V such that every finite subset of S is *truth-functionally satisfiable* (true in at least one Boolean valuation of E^V), then S is truth-functionally satisfiable. (In the proof for this, the tableau will stop at $(Cx)A$ as an atom, rather than continuing to split off.)

Finally, to reiterate: A sentence is *valid* if it is true under all first-order valuations; a sentence is a *tautology* if it is true under all Boolean valuations.

Answer 1. Since the left-hand side and the right-hand side of the implication are two different atoms, we can assign f to the RHS and t to the LHS as a well-defined valuation.

Answer 2. $\forall x(P(x)) \wedge \forall x(\neg P(x))$

Answer 3.

5 First-order analytic tableaux

5.1 Extension of our unified notation

Formulae of the type α and β are defined as they were in propositional logic.
γ stands for any formula of the following type:

- T $(\forall x)A$
- F $(\exists x)A$

δ stands for any formula of the following type:

- F $(\forall x)A$
- T $(\exists x)A$

Four laws concerning first-order satisfiability:

1. If S is satisfiable and $\alpha \in S$, then $\{S, \alpha_1, \alpha_2\}$ is satisfiable. The proof is trivial.
2. If S is satisfiable and $\beta \in S$, then at least one of $\{S, \beta_1\}$ and $\{S, \beta_2\}$ is satisfiable. The proof is trivial.
3. If S is satisfiable and $\gamma \in S$, then for every parameter $a, \{S, \gamma(a)\}$ is satisfiable. The proof is similar to the fourth one.
4. If S is satisfiable and $\delta \in S$, then if a is any parameter which does not occur in $S, \{S, \delta(a)\}$ is satisfiable.
   Proof: By hypothesis, there exists an interpretation I of all predicates of S in some universe U and a mapping φ of all parameters of S into elements of U such that for every $A \in S$, the $U-$sentence A^{φ} is true under I.
   In particular, this means that δ^{φ} is true under I; therefore, there must be at least on element $k \in U$ such that $\delta^{\varphi}(k)$ is true under I.
   We want the parameter a to map to k; therefore, we set a to be any parameter not in S and extend φ by $\varphi^*(a) = k$. $\delta(a)^{\varphi^*}$ is true under I, and the proof is complete.

5.2 Analytic tableaux for quantification theory

In addition to the rules of the analytic tableaux for propositional logic, the following are added:

1. From any formula of the form γ, we infer $\gamma(a)$.
2. From any formula of the form δ, we infer $\delta(a)$ (with the proviso that a is new).

Theorem 5.1. *Every sentence provable by the tableau method for first-order logic is valid.*

Proof. Any immediate extension of a tableau which is satisfiable is again satisfiable.
If the origin is satisfiable, then by induction, at least one branch of the tableau is satisfiable and hence open.
Therefore, if a tableau closes, then the origin must be unsatisfiable. □

Since no sentence and its negation can both be valid, it follows that the tableau method is consistent for first order logic.

5.3 The Completeness Theorem