

Mathematics, the Science of Reasoning

Mathematics, the Science of Reasoning

Julia Kameryn Williams
Cape Cod Community College

August 19, 2026

Website: juliakw.net/books/reasoning/

©2026–2026 Julia Kameryn Williams

This work is licensed under the Creative Commons Attribution-NonCommercial-ShareAlike 4.0 International License. To view a copy of this license, visit CreativeCommons.org

Contents

Preface: To the student	1
1 Logic	2
1.1 Arguments	2
1.2 The language of propositional logic	5
1.3 Truth tables.	10
1.4 Entailment	13
1.5 Other semantic notions.	15
1.6 Quantifiers	16
1.7 Models and the semantics of quantifier logic	20
2 Statistics	24
2.1 Visualizing data	24
2.2 Averages	24
2.3 Range and quantiles.	24
2.4 Box and whisker plots	24
2.5 The standard deviation.	25
2.6 Normal distributions.	25
2.7 Hypothesis testing	25
3 Modeling	26
3.1 Functions.	26
3.2 Linear models	26
3.3 Exponential functions	26
3.4 Exponential models	26
3.5 Logistic models	27

Back Matter

Preface: To the student

Chapter 1

Logic

Logic is the study of **deductive reasoning**, whereby you explore the consequences of the presumed certainty of your premises. Logic is applied in mathematics, computer science, law, philosophy, and other disciplines.

There are many different **logics**, formal systems of reasoning, which attempt to capture different domains and applications. We will survey a small corner of these many possibilities. To drop some jargon (keywords the interested reader can google to learn more), this chapter is mainly focused on **propositional logic**, with some minor dabbling in **predicate logic**. Our approach is **classical**, meaning we will assume the **law of excluded middle**. Our focus will primarily be on **semantics** and we will do little with **syntactic** concepts such as **proofs** or **deduction rules**. The reason for these choices is to enable a simpler introduction to the subject.

This chapter aims for the following goals for the reader.

- Be familiar with the language of logic and translate between ordinary English and the formal language.
- Understand the meaning of core concepts like implication and validity.
- Compute truth tables and use them to analyze logical properties.

1.1 Arguments

A core use of logic is assessing **arguments**. Informally, an argument is a reasoned explanation for why someone should believe a conclusion. (Think argument as in “legal argument”, not as in a shouting match.) We want to be able to understand what separates good arguments—those that are compelling and don’t have holes in their reasoning—from bad arguments. We will focus on the deductive side of things. Rhetoric, appeals to emotion, and so on are all important if you want to convince people to adopt your views. But they will not be what we talk about in this class.

With this use in mind, we will adopt the following formalization of what an argument is.

Definition 1.1.1 Arguments. An **argument** is a list of **premises** followed by a **conclusion**. A premise or conclusion is a **statement**, by which is meant a **truth-apt**, declarative statement. ♦

Not all English sentences are statements. Questions like “What time is it?” or commands like “Obey the speed limit!” are not truth-apt. It doesn’t make

sense to ask whether they are true or false. They aren't what we're interested in with logic.

Truth is a difficult and slippery notion. There are philosophers who spend their entire careers trying to understand this concept. Fortunately, we can do logic just fine with a common sensical notion of truth. We will assume it makes sense to talk about whether statements like "The sky is blue." or "Rhode Island has a higher average household income than Massachusetts." are true or false. We will also consider statements like "Frodo threw the One Ring into Mount Doom." or "Coffee tastes good." as truth-apt, even though the first is a statement about a fictional world and the second is an opinion. We are not interested in truth in an absolute sense but rather truth in the sense of X . Here, X could be many different things: true in the sense of physical reality, true in the sense of *Lord of the Rings*, true in the sense of Professor Williams's opinions, or so on. The tools of logic can be applied to all of these.

When speaking we will usually just say true or false and leave the "in the sense of X " part silent.

Example 1.1.2 Here are a few examples of arguments, each in the form of a short paragraph. We can use words like "therefore" or "thus" to mark which statement is the conclusion.

- If you want to pass math class then you should study. Therefore you should study.
- Gollum threw the One Ring into the fires of Mount Doom. Therefore Frodo did not throw the One Ring into Mount Doom.
- It is raining. Therefore I should bring an umbrella.

Do you think these are good, compelling arguments? Or are they unconvincing? ▲

Maybe you found each of these arguments convincing, in the sense that if the premises are true then the conclusion must also be true. But they are each missing something. Namely, each of them has a hidden premise which is assumed but not stated. For clarity's sake, we want to make all premises explicitly. Let's fix these arguments.

Example 1.1.3 Here are the arguments from before, but with the hidden premises added in.

- If you want to pass math class then you should study. *You want to pass math class.* Therefore you should study.
- Gollum threw the One Ring into the fires of Mount Doom. *If Gollum threw the One Ring then Frodo did not.* Therefore Frodo did not throw the One Ring into Mount Doom.
- It is raining. *If it is raining then I should bring an umbrella.* Therefore I should bring an umbrella.

Having seen these examples let's make clear what we mean by a good argument. ▲

Definition 1.1.4 Validity and Soundness. An argument is **valid** if whenever all its premises are true its conclusion must also be true. An argument is **sound** if it is valid and all of its premises are true. ◆

Sound arguments are good arguments because they tell you their conclu-

sion is true. If you care about believing true things they are what you want. However it can be very difficult to determine whether an argument is sound, because that requires you to determine whether the premises are true.

Example 1.1.5 Consider the following argument, presented as a list of statements with the therefore symbol $\therefore$ marking the conclusion.

- If there are an even number of students enrolled at Cape Cod Community college, then there are not an odd number of students.
- There are an even number of students.
- $\therefore$ there are not an odd number of students.

Checking the soundness of this argument isn't too hard. One premise is a mathematical fact—a number cannot be both even and odd. The other can be figured out because the college keeps good records and has an exact count of the number of students currently enrolled. But you imagine similar arguments whose premises are very hard to check. What if instead of 4Cs we wanted to know about members of an organization from ancient Egypt? How could you possibly get an exact count for that? ▲

As such, validity will be the main thing we focus on for what makes an argument good. Soundness is even better, but it is someone else's problem to do the extra work.

Question 1.1.6 Is there a method that can be used to always determine whether an argument is valid? ▲

Here's a start at answering this question. Let's use the language that a **world** is a possible collection of facts, where we only care to specify facts relative to an argument. For example, consider the following two arguments.

- Lee Harvey Oswald didn't shoot JFK. So the CIA did it.
- Either Lee Harvey Oswald shot JFK, or else the CIA did it. Oswald didn't shoot JFK. Therefore the CIA did it.

If an argument is invalid, it means that there is a world (= collection of facts) where all the premises are true but the conclusion is false. For both worlds, the only relevant fact is who shot JFK. We can see that the first argument is invalid by imagining the world where JFK was shot by an alien. In that world, the one premise is true but the conclusion is false. You could specify other facts about the world—maybe the alien's favorite food was hamburgers—but those facts aren't relevant so we don't need to think about them.

On the other hand, an argument is valid if any world which makes the premises true also makes the conclusion true. For the second argument, the only way to make the first premise true is if either Oswald or the CIA shot JFK. To make the second premise true it must be that Oswald didn't do it. The only remaining option is that the CIA shot JFK. But then the conclusion must be true. So this argument is valid.

Note that there is an asymmetry here. Showing an argument is invalid only takes imagining up a single **counterexample** world. But showing an argument is valid requires you to think through the meaning of the argument. This is inconvenient. It'd be nice to have a single method we can use that always tells us the answer. As such, it would be pre-mature to say we have settled this question.

It will take us some time to build up the machinery to settle this question. But we can take it as a guiding star on developing some concepts in logic.

Exercises

1. Some arguments are presented in ordinary language. For each of them, identify the premises and the conclusion.
 - (a) Everyone who has a stake in how a state is run should be allowed to vote. Immigrants have a stake in how a state is run. Therefore immigrants should be allowed to vote.
 - (b) You should buy me lunch because it's my birthday tomorrow.
 - (c) All cats are carnivores. Martin is a cat, so he is a carnivore.
 - (d) If you study for math class you will pass. You passed, so you must have studied.
2. For each of the arguments from the previous problem, state whether you think it is valid. If it is invalid, is there a hidden premise you can add to make it valid? If so, add this premise. If not explain why you think it's invalid.
3. Each of the following arguments is invalid. Kindly demonstrate this by exhibiting a counterexample world.
 - (a) Chocolate ice cream is delicious. Therefore vanilla ice cream is mid.
 - (b) If I win the lottery then I will quit my job. I quit my job. Therefore I won the lottery.
 - (c) It is raining outside. Therefore I should bring my umbrella.
4. Explain why the following argument is valid, by means of analyzing worlds.
 - If I win the lottery then I will quit my job. I didn't quit my job. Therefore I didn't win the lottery.

1.2 The language of propositional logic

As a first task we will establish a formal language for talking about logic. The reason for doing this is to avoid the ambiguities inherent in natural language.

A **proposition** is a truth-apt statement—something that makes sense to say is true or false. We say a proposition is **atomic** if it can't be broken down any further into simpler propositions.

Example 1.2.1 “Grass tastes good” is an atomic proposition. If you try to break it down into smaller pieces you get partial sentences that don't make sense by themselves. On the other hand “grass tastes good and dirt tastes bad” is not atomic. you can break it down into simpler pieces, namely “grass tastes good” and “dirt tastes bad”. ▲

When we were analyzing worlds to check validity of an argument, you maybe noticed that the meaning of an atomic proposition didn't matter, merely its truth value. With that in mind, when doing logic we usually don't care about the content of an atomic proposition; it's irrelevant to our analysis. Accordingly, we will use **propositional variables** $A, B, C, \dots$, similar to how mathematicians use number variables to stand in for a known or unknown quantity.

When you are abbreviating a natural language proposition as a variable it's often helpful to pick a letter to remind you of the meaning. For example, we might use G for “grass tastes good” and D for “dirt tastes bad”. But the choice

of letter doesn't affect the meaning; we could've instead picked G for “good” and B for “bad” and it'd all be the same. If you're working abstractly and the variable doesn't stand for anything in particular then we'll usually pick letters near each other, just for convenience.

We want to combine simpler propositions to create more complex propositions. There's multiple ways to do this. For instance, “I'm rich and I'm a nice person” has a different meaning from “I'm rich or I'm a nice person” has a different meaning from “If I'm rich then I'm a nice person”. So we need multiple different ways to combine.

Logical **connectives** are those ways to combine. We write each of them with a single symbol. Compare to how mathematicians write symbols like $+$ or $\times$ to denote ways to combine numbers.

Definition 1.2.2 The main five logical connectives. There are five main connectives. Four of them are **binary**, meaning they connect two simpler propositions. One of them is **unary**, meaning it modifies a single proposition. We write the binary connectives in between the two.

Conjunction	We write $A \wedge B$ to mean “ A and B ”.
Disjunction	We write $A \vee B$ to mean “ A or B ”.
Negation	We write $\neg A$ to mean “not A ”.
Conditional	We write $A \Rightarrow B$ to mean “if A then B ”.
Biconditional	We write $A \iff B$ to mean “ A if and only if B ”.



One advantage of this compact language is that it reduces the phenomenon of having many ways to say the same thing. For instance, here's a few different ways to express $A \wedge B$.

- A and B .
- A but also B .
- Both of A and B .

There can be slight variation in the connotation of these different ways of phrasing things. For example, if I say “it is raining but it is warm” then the but indicates some sort of contrast or unexpectedness. This fine gradation of meaning isn't relevant to our logical work however.

Conditionals—if-then statements—are especially notorious for how many different ways there are to express the same idea. Here's a few different ways people say $A \Rightarrow B$.

- If A then B .
- A implies B .
- B because A .
- B only if A .
- A is a sufficient condition for B .
- B is a necessary condition for A .

These last two are especially beloved by philosophers. To explain them: if A is a sufficient condition for B , it means that whenever A is true then that is sufficient to conclude B must also be true. That is, if A then B . For the other

one, B being a necessary condition for A means if A is true then necessarily so is B . Again, you can phrase that as if A then B .

The phrase “if and only if” gets used a lot in mathematics, but it’s less common in ordinary speech. So let me make sure the meaning is clear. You can think of it as saying the implication goes both ways—hence why it’s called a *biconditional*. That is $A \iff B$ means “ A if and only if B ” means “if A then B and if B then A ”.

Checkpoint 1.2.3 Philosophers will sometimes say $A \iff B$ as “ A is a necessary and sufficient condition for B ”. Explain why that has the same meaning, based on the meaning of necessary and sufficient conditions.

It also should be made clear what or ($\vee$) means. In ordinary use there’s two thing it might mean. The first is saying one of the options, but not both. For example, if your waiter informs you that your burger comes with fries or a side salad, “both” is not an acceptable answer. Because we exclude the possibility of both, we call this a **exclusive or**. The other possible meaning is an **inclusive or**, allowing the possibility of both. For example, imagine your math professor promises to bring pizza if either the whole class passes the midterm or the average is at least 85%. If everyone passed and the average was at least 85% you would be upset if she tried to weasel out of the promise by saying “I said or, which means not both”. Rather, the clear meaning was that both happening counts as fulfilling the requirement.

In mathematics and in logic, it’s convenient to take the inclusive or as the standard. This is the meaning we most often want. For example, it’s a true mathematical statement that every number is either ≥ 0 or ≤ 0 . But this is only true because we allow the possibility of both, as happens with 0. With this in mind, $A \vee B$ will mean inclusive or— A or B , or possibly both.

It turns out that we can express exclusive or using the five connectives, so it’s no loss to not take it as a basic concept.

Example 1.2.4 A exclusive-or B can be understood as “either A and not B , or else B and not A ”. We can write this as a sentence in our formal language:

$$(A \wedge \neg B) \vee (\neg A \wedge B).$$

Thus we can express exclusive-or using the three connectives of conjunction (and/ $\wedge$), disjunction (or/ $\vee$), and negation (not/ $\neg$). ▲

Checkpoint 1.2.5 the $\vee$ in the formal sentence in the example is an inclusive or. Explain why it’s impossible for both **disjuncts**—the formulas on each side of the or—to be true.

Hint. Can you have both A and $\neg A$ true at once?

In general, we want to be able to combine multiple connectives to make more complicated sentences. Here’s some rules for how we do this.

- Negation ($\neg$) always happens first. If you write $\neg A \wedge B$ this means $(\neg A) \wedge B$ —the not happens before the and. If you want to do it in the other order you need to use parentheses; see the next rule.
- We use parentheses to denote priority—what comes first. This is similar to the use of parentheses for priority in arithmetic formulas. Just like $2 \cdot (3 + 1)$ means to add before you multiply, $A \vee (B \wedge C)$ means to and before you or.
- If we have three or more things connected by ands (or three or more things connected by ors), we don’t need to write parentheses to say which and (or which or) comes first. This is because the meaning is the same

regardless of order. Compare to how order doesn't matter for addition; you get the same sum no matter what order you do the additions in $1 + 2 + 3$.

When translating between the formal language and ordinary language, you should give a **translation key** specifying the meaning of the atomic variables.

Example 1.2.6 We want to translate “If you get a college degree you will be happy and wealthy” into the logical language. Our first step is to identify the pieces of the sentence we can't break down any further. You might identify two pieces making up the two parts of an if-then statement, namely

- You get a college degree.
- You will be happy and wealthy.

But that second piece can be broken down even further:

- You will be happy
- You will be wealthy.

That is, to break a sentence down into pieces you sometimes have to not follow exactly how it's worded. In this example, you have to slightly rephrase one of the propositions connected by the “and” to be a complete sentence on its own.

In all, we have three atomic propositions, which we need to give names to.

C = You get a college degree.

H = You will be happy.

W = You will be wealthy.

With this translation key we can translate the sentence into the formal language. Looking at the connecting words, we notice an “if-then” and a “and”. The “and” connects H and W , so that part of the sentence becomes $H \wedge W$. And the if-then is if C then H and W . So we see the whole thing comes together as

$$C \Rightarrow (H \wedge W).$$

We put the parentheses to make clear that the and is inside the then part of the if-then. ▲

Example 1.2.7 Let's use the same translation key as the previous example.

C = You get a college degree.

H = You will be happy.

W = You will be wealthy.

We use this key to translate the following into ordinary language

$$(C \wedge W) \vee (H \wedge \neg W).$$

If you just replace each connective with the corresponding English word you get something like:

- You get a college degree and you will be wealthy or you will be happy and not you will be wealthy.

This isn't how normal people talk, so let's pretty it up.

- Either you get a college degree and will be wealthy or else you will be happy but not wealthy.

With ordinary language you don't have parentheses to mark priority so you sometimes need to phrase your sentence to make it clear where things break. ▲

Checkpoint 1.2.8 Write a translation key for and translate into formal logic the following sentences.

- Richard's novel is a commercial success but has no artistic merit.
- Cass respects Richard if and only if his novel has artistic merit.

Solution.

$$C \wedge \neg A$$

$$R \iff A$$

Using the translation key:

C = Richard's novel is a commercial success.

A = Richard's novel has artistic merit.

R = Cass respects Richard.

Exercises

1. Translate the formal sentences into ordinary language, using the following translation key for the meaning of atomic variables:

A = Antinous wants to marry Penelope.

T = Telemachus wants to kill Antinous.

P = Penelope worries Antinous will kill Telemachus.

- (a) $T \Rightarrow P$
- (b) $P \wedge \neg A$
- (c) $P \Rightarrow (T \vee \neg A)$
2. Translate the following sentences of ordinary language into the formal logical language. Provide a translation key giving the meaning of your atomic variables.
 - (a) If the Trojans let the horse into their city, then the Greeks will win the Trojan war.
 - (b) Either the Trojans will let the horse into their city or else they won't.
 - (c) If the Trojans let the horse into their city and the Greeks win the Trojan war then Odysseus will feel guilty.
3. If you forget to put in parentheses, the meaning of a sentence can be ambiguous. Consider the formal sentence

$$A \wedge B \Rightarrow C.$$

There's two places to put parentheses, either around $A \wedge B$ or around $B \Rightarrow C$. For each of the two possibilities, translate it into an ordinary language sentence. As part of this, create a translation key giving the meaning of the atomic variables (your choice what they mean!). Explain why the two ordinary language sentences have different meanings.

4. Translate the following pairs of sentences into ordinary language, providing a translation key for the variables. Do your translations convince you that the paired statements have equivalent meaning? Explain

$$\begin{array}{lll}
 A \wedge B & \text{and} & B \wedge A \\
 A \vee B & \text{and} & B \vee A \\
 (A \wedge B) \wedge C & \text{and} & A \wedge (B \wedge C) \\
 (A \vee B) \vee C & \text{and} & A \vee (B \vee C)
 \end{array}$$

5. Translate the following into ordinary language, providing a translation key for the variable. Do you think any of them have equivalent meaning? If so, which? Explain.

$$\begin{array}{c}
 A \\
 \neg A \\
 \neg\neg A
 \end{array}$$

1.3 Truth tables

Now that we know how to write logically, let's talk about the meaning—as logicians would say, the **semantics**—of the symbols. We will do this through the device of a **truth table**, invented by the Austrian philosophy Ludwig Wittgenstein. The idea is, a truth table lays out how a symbol behaves for any possible combination of inputs. This exhaustive listing of all possibilities constitutes a complete description of its meaning.

This is a **truth-functional** explanation of the meaning of the connectives. That is, if we want to know whether a statement is true or false the meaning of the input doesn't matter, only their truth values.

Let's see truth tables for the five connectives. We begin with conjunction (and) and disjunction (or). These each take two inputs, on the left and right. Each input has two possible truth values, namely **true** (T) or **false** (F). So in total we need $2 \cdot 2 = 4$ rows to represent all possibilities.

Remember, we're looking at propositions, statements which are unambiguously true or false. So those are the only possible truth values. You might say, what about statements are only kinda true? For instance, someone might be kinda tall so "John is tall" isn't totally true but nor is it totally false.

You can do logic allowing for intermediate truth values. But things get more complicated. We focus on the two-valued logic with only true and false to make the analysis less difficult. I hope you forgive the trade-off.

We organize the table so that the leftmost columns give the values of the inputs, while on the right the values of the output are given.

Table 1.3.1 Truth table for conjunction

A	B	$A \wedge B$
T	T	T
T	F	F
F	T	F
F	F	F

What this table says is, A and B is true just in case both A and B are true. We take the common sense definition of "and" and write it in tabular form.

Table 1.3.2 Truth table for disjunction

A	B	$A \vee B$
T	T	T
T	F	T
F	T	T
F	F	F

Here, $A \vee B$ is true when at least one of the two is true. Recall that or is inclusive-or, allowing the possibility of both disjuncts being true.

For $\neg$ (not) there is only one input, written on the right. Accordingly we only need two rows for the table.

Table 1.3.3 Truth table for negation

A	$\neg A$
T	F
F	T

Negation acts as a toggle flipping the value of its input to the opposite.

For the biconditional ($\iff$), what it is saying is that the two sides are equivalent. From that you can deduce the rule for filling out the table: the output should be true just when the two inputs are the same.

Table 1.3.4 Truth table for biconditional

A	B	$A \iff B$
T	T	T
T	F	F
F	T	F
F	F	T

The conditional was saved for last because it's the most subtle and controversial. An issue is, the ordinary use of if-then statements aren't truth-functional. Instead, they are used to imply some kind of connection between the statements. Consider a statement like "if it is sunny out you should wear sunscreen". There is a causal connection; the weather presents a minor health risk and so there is prudential reason to take steps to minimize that risk. In contrast imagine someone saying "if it is sunny out then Mars has two moons". Probably your reaction is to be confused. What is the connection between those two things?

But if the meaning of $A \Rightarrow B$ can only depend upon their truth values, then we can't require a connection. That would require information beyond whether they are true or false. Subject to that restriction, the best we can do is the following.

Table 1.3.5 Truth table for conditional

A	B	$A \Rightarrow B$
T	T	T
T	F	F
F	T	T
F	F	T

A second, more substantive use of truth tables is to analyze more complicated logical sentences. We can use the truth tables defining the connectives as rules for how to determine the behavior of more complicated sentences, and thereby understand what they are saying.

When you are doing this, you want to think of how the logical sentence is built up. For example, suppose you are trying to fill out the truth table for

$$A \wedge (B \vee \neg A).$$

The last connective to be evaluated is the $\wedge$. You can evaluate it what it does by using the rule given in the truth table for conjunction (and). But to do this you need to know the inputs on both sides. You know the input on the left, since it's just A . But you have to do some work to figure out the input on the right. How do you do that? The right conjunct is an or statement. You can evaluate that using the truth table for $\vee$, but first you need to know the pattern for its two inputs. The left disjunct B you already know, but to figure out the right disjunct $\neg A$ you need to apply the rule for $\neg$. So building up this truth table is a multi-step process. The following table represents the process, where going from left to right gives you the information needed for the next step.

Table 1.3.6 Truth table for $A \wedge (B \vee \neg A)$

A	B	$\neg A$	$B \vee \neg A$	$A \wedge (B \vee \neg A)$
T	T	F	T	T
T	F	F	F	F
F	T	T	T	F
F	F	T	T	F

To explain the first row, as a representative example: to fill out the first column, you look at the A column and apply the rule for $\neg$. Namely, you flip T to become F . Then to fill out the $B \vee \neg A$ columns you use the rule for $\vee$, looking at the B and $\neg A$ columns for input. In this case, you have an instance of $T \vee F$ which evaluates to T . Finally, to fill out the $A \wedge (B \vee \neg A)$ column you use the rule for $\wedge$, looking at the A and $B \vee \neg A$ columns for inputs. Namely, you have an instance of $T \wedge T$ which evaluates to T . This finishes the first row, and the other three are done similarly.

When filling these out, you may prefer to save on space by not writing every subformula separately and just filling in the value below it in the one big formula. That is fine, but you should make it clear what the final output is.

Exercises

1. Fill out a truth table for $A \Rightarrow (A \vee B)$.
2. Fill out truth tables for $\neg A \vee B$ and $A \Rightarrow B$. What do you observe about the two truth tables? Take a stab at explaining the meaning of what you see.
3. Fill out truth tables for $(A \Rightarrow B) \wedge (B \Rightarrow A)$ and $A \iff B$. What do you observe about the two truth tables? Take a stab at explaining the meaning of what you see.
4. Fill out a truth table for $(A \wedge \neg B) \vee (\neg A \wedge B)$. Explain why the pattern on the truth table is the pattern for exclusive-or.
5. How many rows do you need in a truth table for a proposition with three variables? Fill out a truth table for $A \wedge (B \vee \neg C)$.
6. Fill out a truth table for $A \Rightarrow (B \Rightarrow C)$.

1.4 Entailment

With the device of truth tables in hand we are finally ready to give a method to answer our guiding question and give a method to check the validity of an argument.

Say that a sentence C is **entailed** by premises $P_1, P_2, \dots$ if whenever each premise is true then C is also true. Using this language, an argument is valid when its conclusion is entailed by its premises.

Because a truth table gives an exhaustive listing of the behavior of sentences under all possible inputs, we can use them to operationalize this definition. Namely, each row of a truth table corresponds to a possible world—a possible state of facts about the inputs. A sentence C is entailed by premises if any row where each premise comes out to true has C also true. This gives us a procedure to check whether an argument is valid.

Definition 1.4.1 Procedure for checking validity of an argument. You can check validity of an argument by the following procedure.

1. If not already written in formal logic, translate each proposition in the argument into the formal language.
2. Construct truth tables for all propositions in the argument. It's convenient to put them all in one table so you can easily compare them.
3. If the conclusion is true in every row where all premises are true, then the argument is valid. Otherwise, if the argument has a **counterexample** (a row where all premises are true but the conclusion is false), then the argument is invalid.

◆

This procedure, except for the translation step, is completely mechanical. You don't need to know the meaning of anything involved, you just need to follow rules to calculate an object and then look at it to extract an answer. We were able to take our question of how do you check whether an argument is valid and turn it into a kind of algebra problem.

Example 1.4.2 A fallacious argument pattern is **affirming the consequent**. Namely, it's an argument which takes the form:

- B
- $A \Rightarrow B$
- $\therefore A$

We can use the truth table procedure to show why it's an invalid argument.

Table 1.4.3 Truth table for affirming the consequent

A	B	B	$A \Rightarrow B$	A
T	T	T	T	T
T	F	F	F	T
F	T	T	T	F
F	F	F	T	F

Looking at this, we see that there are two rows where both premises are true, namely the first and third rows. In the first the conclusion is also true, but in the third the conclusion is false. We have a counterexample, so the argument is invalid. ▲

Example 1.4.4 A similar argument pattern is called **modus ponens**. It's an argument which takes the form:

- A
- $A \Rightarrow B$
- $\therefore B$

We can use the truth table procedure to show why it's a valid argument pattern.

Table 1.4.5 Truth table for *modus ponens*

A	B	A	$A \Rightarrow B$	B
T	T	T	T	T
T	F	T	F	F
F	T	F	T	T
F	F	F	T	F

The only row where both premises are true is the first row. In this row, the conclusion is also true. So the argument is valid. $\blacktriangle$

Using this same method you can check the validity (or invalidity) of various arguments. Here are a few, which you check yourself in the exercises.

Valid arguments

- *Modus tollens*: $\neg B, A \Rightarrow B, \therefore A$
- Disjunctive syllogism: $A \vee B, \neg B, \therefore A$
- Hypothetical syllogism: $A \Rightarrow B, B \Rightarrow C, \therefore A \Rightarrow C$

Invalid arguments

- Denying the antecedent: $\neg A, A \Rightarrow B, \therefore \neg B$
- Affirming a disjunct: $A \vee B, A, \therefore \neg B$
- *Non sequitur*: $A, \therefore B$

Exercises

1. For each of the valid argument patterns *modus tollens*, disjunctive syllogism, and hypothetical syllogism: (1) create a truth table to confirm its validity; (2) give an ordinary language instance of the argument pattern; (3) give an intuitive explanation for why that argument is valid.
2. For each of the invalid argument patterns denying the antecedent, affirming a disjunct, *non sequitur*: (1) create a truth table to confirm its invalidity; (2) give an ordinary language instance of the argument pattern; (3) explain a counterexample which shows the invalidity of the argument.
3. For each of the following arguments use truth tables to determine whether it is valid. Give an intuitive argument for its (in)validity.

- (a) $A, B, \therefore A \wedge B$
- (b) $A, A \vee B, \therefore A \wedge B$
- (c) $A, \therefore A \vee B$
- (d) $\neg A, \therefore A \Rightarrow B$
- (e) $\neg A, \therefore A \iff B$

1.5 Other semantic notions

Truth tables allowed us a process for checking whether is an argument is valid. They can be used for other notions. In this section we do this.

Definition 1.5.1 Consistency. A collection of propositions are called **jointly consistent** if it's possible for them all to be true at once. In terms of truth tables, they are jointly consistent if there's a row where they all are true. Intuitively, propositions are jointly consistent if there's some world where they're all true at the same time.

If we're talking about just one sentence we say it's **consistent** if there's a row of its truth table where it's true. ♦

Example 1.5.2 Consider the proposition $A \wedge \neg A$. We can see it is not consistent by checking its truth table.

Table 1.5.3 Truth table

A	$A \wedge \neg A$
T	F
F	F

No row of the truth table has this proposition true, so it is not consistent.

On the other hand, consider the proposition $A \wedge \neg B$. We can see it's consistent by checking its truth table.

Table 1.5.4 Truth table

A	B	$A \wedge \neg B$
T	T	F
T	F	T
F	T	F
F	F	F

There is at least row where this proposition is true, namely the second row, so it is consistent. ▲

Definition 1.5.5 Tautology and contradiction. A proposition is a **tautology** if it always true—every row of its truth table is true. And a proposition is a **contradiction** if it is always false—every row of its truth table is false. ♦

The first truth table from the previous example shows that $A \wedge \neg A$ is a contradiction. Indeed, note that being a contradiction is the same as not being consistent. By a similar process you can check that $A \vee \neg A$ is a tautology. Note that every tautology is consistent, but not vice versa.

An important use of truth tables is confirming that two propositions are equivalent.

Definition 1.5.6 Equivalence. Two propositions are **logically equivalent** when one is true if and only if the other is true. In terms of truth tables, two propositions are logically equivalent if their truth tables have the same pattern of trues and falses. ♦

Using truth tables you can check, as you are asked to do in the exercises, that the following pairs of propositions are equivalent.

Associativity of $\wedge$ $A \wedge (B \wedge C)$ and $(A \wedge B) \wedge C$.

Associativity of $\vee$ $A \vee (B \vee C)$ and $(A \vee B) \vee C$.

Commutativity of $\wedge$ $A \wedge B$ and $B \wedge A$.
Commutativity of $\vee$ $A \vee B$ and $B \vee A$.

Together, these equivalences say that for conjunction and disjunction order doesn't matter. If you have three or more inputs the order you do the ands/ors won't affect the outcome, and if you swap which sides of the and/or the inputs are on you get the same outcome.

Another useful equivalence is that a conditional $A \Rightarrow B$ is logically equivalent to its **contrapositive** $\neg B \Rightarrow \neg A$. This can be checked using truth tables. On the other hand, $A \Rightarrow B$ is not equivalent to its **converse** $B \Rightarrow A$.

Exercises

- For each of the following propositions use truth tables to classify it as a contradiction, a tautology, or consistent.
 - $A \Rightarrow \neg A$
 - $A \Rightarrow A$
 - $B \Rightarrow (A \Rightarrow B)$
 - $A \Rightarrow (A \wedge B)$
 - $\neg((A \Rightarrow B) \vee (B \Rightarrow A))$
- Are $A \Rightarrow B$ and $\neg A \vee B$ logically equivalent? Check using truth tables. Give an intuitive explanation for your answer.
- Construct truth tables to demonstrate the equivalences for the associativity and commutativity of conjunction and disjunction.
- Construct truth tables to demonstrate that an implication is equivalent to its contrapositive but not its converse.
- Explain why two propositions P and Q are logically equivalent just in case $P \iff Q$ is a tautology.

1.6 Quantifiers

Previously we've looked at **propositional logic**. This logic has limited expressiveness, and there's many natural arguments it can't talk about. Consider these two examples.

Example 1.6.1 This is a classical argument going back to Aristotle.

- All humans are mortal. Socrates is a human. Therefore, Socrates is mortal.

If you try to cast this into propositional logic, you'd have to treat each sentence as its own atomic proposition. There's no way to break them down further using the concepts of propositional logic. Thus the argument becomes

$$A, B, \therefore C.$$

Doing this erases the logical inference being made. That is, we don't want to treat these statements as atomic, we want to treat them as saying something about connections between properties of things.

This kind of reasoning shows up all the time in mathematics.

- Every positive number has a square root. 2 is positive. Therefore, 2 has a square root.

Again, casting this into propositional logic would destroy all meaning. These aren't unconnected atomic statements, they are statements about how properties of numbers relate. ▲

We need to introduce some new ideas to be able to make logical sense of these arguments. This will give us the system known as **predicate logic**. The first new idea is we want to have names for individuals, not for atomic propositions. A name might be a **variable** like x or y standard for an unspecified individual. We use lowercase letters for variables for individuals. Or a name might pick out a specified individual, such as Socrates or the number 2.

With propositional logic, our logical arena was just about sentences. Is this sentence true or false? That was all we cared about. In this new arena we have a **domain** of individuals we are talking about. Your choice of domain will be determined by what you want to look at. If you want to argue that Socrates is mortal, it makes sense to take humanity as your domain. If you want to talk about square roots, it makes sense to take the real numbers (the numbers which measure continuous quantities) as your domain.

In this context our atomic statements are statements about individuals, which leads to our next new idea. We need **predicate** symbols to stand for statements that individuals have such-and-such property. We use capital letters for predicates to distinguish them from names for individuals. For example, we might write $M(x)$ to mean “ x is mortal”, so that $M(\text{Socrates})$ is true. Or we might write $R(x)$ to mean “ x has a square root”. Then $R(2)$ is true.

Those two examples are what we call **unary** predicates, meaning they are about a single individual. Predicates can also be **binary**, about two individuals, or **ternary** (about three individuals), or even more. For example, if you want to talk about a property like “ x is a square root of y ” this isn't about just one individual but about how two relate. So you'd need a binary predicate like $R(x, y)$ and then $R(2, 4)$ and $R(-2, 4)$ would both be true. Or consider the relation “ x is the mother of y ”. This isn't about just x or just y , but about both of them. You could write it as $M(x, y)$.

When using predicate symbols that stand for something you should include a translation key that explains their meanings.

These two concepts are enough to make sense of sentences like “Socrates is mortal”. But we need one more thing to make sense of “all humans are mortal”. This isn't a statement about a specified individual, but rather a general state about all individuals. When making general statements about the entire domain there's multiple things you might want to say. Two of them stand out as particularly useful.

The first is saying “all individuals have property P ”. This is a general statement about the domain as a whole, which doesn't reduce to a statement about a single individual. The second is saying “some individuals, at least one, have property P ”. Again, this is a statement about the domain as a whole. We introduce **quantifiers** to allow us to make these general statements.

The **universal quantifier** is how we say “all individuals have P ”. We write

$$\forall x P(x)$$

to mean “every individual x has property P ”. The variable x is the **bound variable** of the quantifier. Think of it as a pronoun; we introduce an ad-hoc way to refer to the unstated individuals we are looking at.

Example 1.6.2 We could express “all men are mortal” as

$$\forall x M(x)$$

where our domain is humanity and $M(x)$ means “ x is mortal”. ▲

The **existential quantifier** is how we say “some individuals have P ”. We write

$$\exists x P(x)$$

to mean “there exists at least one individual x which has property P ”. Again, the bound variable x functions as a way to refer to the unstated individuals we are looking at.

Example 1.6.3 You could formalize “2 has a square root” as a single predicate. But it’s more natural to use the existential quantifier. After all, this statement is saying there is some number which squared gives 2. With that in mind, let’s write $S(x, y)$ to mean $x^2 = y$. Then we can express “2 has a square root” as

$$\exists x S(x, 2)$$

where our domain is the real numbers.

Let’s consider another statement, namely “every positive number has a square root”. To express this we need a new predicate $P(x)$ meaning “ x is positive”. This is a statement we need two quantifiers to express. Namely, we can do it as

$$\forall y P(y) \Rightarrow \exists x S(x, y).$$

You can cast this into English as: “for every number y if y is positive then there is a number x which is the square root of y .”

This second example illustrates why we want quantifiers to have bound variables attached. When you are making complicated statements about multiple objects it’s convenient to have different names so you can distinguish which you are talking about. ▲

Warning! Two different variables can refer to the same individual. Take the square root example where $S(x, y)$ means $x^2 = y$. For most numbers, their square root must be different. But there’s a couple numbers which are their own square root, namely 0 and 1. If we plug in 1 for both x and y then $S(x, y)$ —that is, $1^2 = 1$ —is a true statement.

We can combine quantifiers with the connectives from propositional logic. These connectives can occur inside or outside of quantifiers. Carefully note that their exact position affects the meaning of a sentence. With that in mind, it’s never a bad idea to put in parentheses to make clear the order of things.

Exercises

1. Use the following translation key.

$R(x) = x$ is a rational number

$0 =$ the number 0

$\pi =$ the circle constant

Use this translation key with the real numbers as domain to translate the following sentences into formal logic.

- (a) 0 is rational but π is irrational. (Note: irrational means not rational.)
 - (b) There is a number which is irrational.
 - (c) There is a number which is irrational and there is a number which is rational.
 - (d) Every number is either rational or irrational.
2. If you've ever ridden a subway you may have heard the announcer say "all doors will not open" at a stop. What they actually mean is "not all doors will open". Translate both sentences into formal logic. As part of this, state what your domain is and give your translation key. Explain why the two sentences have different meaning.
 3. Use the following translation key.

$P(x) = x$ is prime

$M(x, y) = x$ is a multiple of y

$E(x, y) = x$ and y are equal

$2 =$ the number two

$9 =$ the number nine

Use the positive whole numbers as your domain. Translate the following sentences of formal logic into ordinary language. As much as possible, try to write ordinary sentences that don't have variables in them.

- (a) $P(2) \wedge \neg P(9)$
 - (b) $\neg \forall x P(x)$
 - (c) $\neg M(2, 9) \wedge \neg M(9, 2)$
 - (d) $\forall x \exists y M(y, x)$
 - (e) $\forall x \forall y [(M(x, y) \wedge M(y, x)) \Rightarrow E(x, y)]$
4. The order of quantifiers affects the meaning of a sentence. If you swap a universal and existential quantifier you get a different meaning. To illustrate, let $M(x, y)$ mean " x is the mother of y ". Translate the following two sentences into ordinary language and explain why they mean different things. Please in your translations don't refer to variables; that's not how ordinary people talk.

$\forall y \exists x M(x, y)$

$\exists x \forall y M(x, y)$

1.7 Models and the semantics of quantifier logic

With propositional logic, truth tables gave us a complete specification of how a logical sentence might behave in any possible situation. We could exhaustively list out all possible combinations of inputs and compute the behavior for each possibility. In this way we could completely define the semantics—meaning—of a sentence of propositional logic.

Predicate logic gives us extra expressive power, but that comes at the cost of a more complicated semantics. In this section we will take a brief look at this topic. The goal is not to attain a complete mastery. That would take a dedicated course. Instead, the goal is get a sketch of the big ideas and thereby an appreciation of the difficulties and the work necessary to overcome these difficulties.

Here is the fundamental challenge. In propositional logic, a possible state of affairs can be summarized as a listing of truth values of the various variables. If you have n variables, then there are 2^n many possible states of affair. That can be a big number, it's a finite list you can exhaustively work through. (More realistically, if it's a big number you can program a computer to work through it for you.)

In predicate logic, an interpretation of the possible meaning of a sentence is more complicated. You need to specify the domain and you need to specify the meaning of the predicate symbols in the sentence. There's a lot of possibilities, infinitely many. It's hopeless to want to put all these possibilities into a single table.

Definition 1.7.1 Interpretations of predicate logic. An interpretation of a sentence (or sentences) of predicate logic is a possible meaning of the symbols in the sentence. More precisely, an interpretation consists of the following data.

- The **domain** or **universe**, a set U of individuals.
- For each name for an individual, a selection of which individual it picks out.
- If there are any unbound variables in the sentence, the interpretation also needs to specify which individuals they refer to.
- For each unary predicate $P(x)$, the set of which individuals in the domain U the predicate is true of. Similarly, for each binary predicate $Q(x, y)$ the set of which pairs of individuals from the domain the predicate is true of. And similarly for ternary, and so on predicates.



Example 1.7.2 Consider the sentence $M(\text{Socrates})$, whose intended interpretation is to express “Socrates is mortal”. An interpretation must specify the domain, which individual Socrates refers to, and which individuals have property M .

- Here's one interpretation, the intended one. The domain is the set of all humans, Socrates refers to the ancient Greek philosopher who was accused of corrupting the youth, sentenced to drink hemlock, etc., and every human has property M . (After all, all humans are mortal.) Under this interpretation, $M(\text{Socrates})$ is true.
- Here's another interpretation. The domain is the set of positive whole numbers, Socrates refers to the number seven, and $M(x)$ when x is as

even number. Under this interpretation, $M(\text{Socrates})$ is false.

As these examples show, the interpretation gives a lot of information. ▲

Given an interpretation, you can then work out whether a statement is true. The rules for connectives are the same as we saw with their truth tables, but we have the extra work of determining what's going on with quantifiers.

Definition 1.7.3 Truth and interpretations. Given an interpretation of a sentence (or sentences) of predicate logic, we can determine its truth according to the following rules.

- An atomic sentence $P(x)$ is true if the interpretation says that individual x has property P .
- If we already know what's going on with sentences A and B , then $A \wedge B$ is true when both A and B are true.
- If we already know what's going on with sentences A and B , then $A \vee B$ is true when at least one of A and B is true.
- If we already know what's going on with sentence A , then $\neg A$ is true when A is false.
- If we already know what's going on with sentences A and B , then $A \Rightarrow B$ is true when either A is false or B is true. That is, $A \Rightarrow B$ is true if the corresponding row in the truth table is true.
- If we already know what's going on with sentences A and B , then $A \iff B$ is true when A and B have the same truth value.
- If we already know what's going on with sentences $A(x)$, then $\forall x A(x)$ is true if for any choice c of individual we have that $A(x)$ is true if x is interpreted to refer to c .
- If we already know what's going on with sentences $A(x)$, then $\exists x A(x)$ is true if there is at least one choice c of individual so that $A(x)$ is true when x is interpreted to refer to c .

This looks complicated, but here's the summary. The truth of atomic sentences is given by the interpretation. Compare to how in predicate logic the truth of a variable was given by the row of a truth table. For the connectives, we use the same rules that define their truth tables. For the two quantifiers, we again use their definitions. Namely, $\forall x P(x)$ means every individual x has property P and $\exists x P(x)$ means at least one individual has property P . ◆

With propositional logic, you could always calculate the truth value of a sentence in a row of the truth table. If the sentence was long this might take some work, but it's a matter of applying enough time to repeatedly apply the rules for each connective. Here, it's not so simple. If the domain is infinite, then determining the rules for quantifiers requires surveying an infinite domain. That could take a very long time.

In general, there is no single uniform procedure you can use to determine whether a sentence is true in an interpretation. In fact, all of mathematics can be cast in predicate logic. That includes difficult, unsolved problems. If you had a method that always worked to determine truth (given an interpretation) in predicate logic you could solve every open problem in mathematics.

Once you have a notion of interpretation and truth in an interpretation, you can carry over all the semantic notions we had for propositional logic. There, interpretations were rows of a truth table. So we make the necessary changes

to swap in the new definition of an interpretation.

Here's a few examples.

Definition 1.7.4 Entailment. A conclusion C is **entailed** by premises $P_1, P_2, \dots$ if any interpretation which makes all premises true also makes C true. ♦

Definition 1.7.5 Consistency. A collection of sentences of predicate logic, possibly just one sentence, is **consistent** if there is an interpretation which makes all the sentences true. ♦

Definition 1.7.6 Equivalence. Two sentences of predicate logic are **logically equivalent** if under any interpretation they have the same truth values. ♦

There is a trade-off. Adding in predicates and quantifiers gave us more expressive power. We can express new arguments, like the one about Socrates being mortal. Indeed, this gives us enough power to do all of mathematics. But this extra power means that the system is harder to analyze. One place you see this is with entailment. With propositional logic, we could use truth tables to get a process for checking validity of arguments. With predicate logic, we don't have an algorithm because we'd have to look at infinitely many interpretations. How then can we check validity of arguments?

One thing you can do is isolate some basic one step deductions that are valid, so that any argument can be cast as a long sequence of these deductions. This leads to the big topic in logic we haven't talked about, namely **proofs**. A proof is a step-by-step deduction of a conclusion from premises, which makes clear that the conclusion is entailed by the premises. A proof is a **syntactic** object, something you can write down. Unlike the semantic notions of predicate logic which require surveying an infinite domain to check, you can sit down and check that each step in a proof is valid.

If you are interested in learning more about logic, this would be the first major thing to study.

Exercises

1. Consider the following statements of predicate logic. For each of them, come up with an interpretation in which the statement is true. Give an intuitive explanation for why your interpretation makes it true.

(a) $\forall x P(x)$

(b) $\exists x Q(x) \wedge \neg R(x)$

(c) $\forall x S(x, x)$

(d) $\forall x \exists y T(x, y)$

2. If you interpret $E(x, y)$ to mean " x and y are the same individual" (we call it E for "equality") then no matter your domain the following two sentences will be true.

(a) $\forall x E(x, x)$

(b) $\forall x \forall y E(x, y) \Rightarrow E(y, x)$

Explain why those are both true for the equality interpretation. Come up with two new interpretations for $E(x, y)$, one where both of these are false and one where exactly one is true. Give an intuitive explanation for why those interpretations do what you want.

Hint. Try to first translate the two sentences into ordinary language to understand what they are saying.

Chapter 2

Statistics

Check back later!

2.1 Visualizing data

Check back later!

Exercises

1.

2.2 Averages

Check back later!

Exercises

1.

2.3 Range and quantiles

Check back later!

Exercises

1.

2.4 Box and whisker plots

Check back later!

Exercises

1.

2.5 The standard deviation

Check back later!

Exercises

1.

2.6 Normal distributions

Check back later!

Exercises

1.

2.7 Hypothesis testing

Check back later!

Exercises

1.

Chapter 3

Modeling

Check back later!

3.1 Functions

Check back later!

Exercises

1.

3.2 Linear models

Check back later!

Exercises

1.

3.3 Exponential functions

Check back later!

Exercises

1.

3.4 Exponential models

Check back later!

Exercises

1.

3.5 Logistic models

Check back later!

Exercises

- 1.

Colophon

This book was authored in PreTeXt.