

forall λ

An Introduction to Formal Logic

P.D. Magnus

University at Albany, State University of New York

fecundity.com/logic, version 1.11 [060124]

This book is offered under the Creative Commons license.
(Attribution-NonCommercial-ShareAlike 2.0)

I am thankful for considerable editorial assistance from Cristyn Magnus. Thanks also to Bin Kang and the students of Introduction to Logic, who detected various errors in previous versions of the book.

© 2005–2006 by P.D. Magnus. Some rights reserved.

You are free to copy this book, to distribute it, to display it, and to make derivative works, under the following conditions: (1) Attribution. You must give the original author credit. (2) Noncommercial. You may not use this work for commercial purposes. (3) Share Alike. If you alter, transform, or build upon this work, you may distribute the resulting work only under a license identical to this one. — For any reuse or distribution, you must make clear to others the license terms of this work. Any of these conditions can be waived if you get permission from the copyright holder. Your fair use and other rights are in no way affected by the above. — This is a human-readable summary of the full license, which is available on-line at <http://creativecommons.org/licenses/by-nc-sa/2.0/legalcode>

Typesetting was carried out entirely in L^AT_EX2 ϵ . The style for typesetting proofs is based on `fitch.sty` (v0.4) by Peter Selinger, University of Ottawa.

This copy of `forall χ` is current as of January 24, 2006. The most recent version is available on-line at <http://www.fecundity.com/logic>

Contents

1	What is logic?	5
1.1	Arguments	6
1.2	Sentences	6
1.3	Two ways that arguments can go wrong	7
1.4	Deductive validity	8
1.5	Other logical notions	10
1.6	Formal languages	12
	Practice Exercises	15
2	Sentential logic	17
2.1	Sentence letters	17
2.2	Connectives	19
2.3	Other symbolization	28
2.4	Sentences of SL	29
	Practice Exercises	33
3	Truth tables	37
3.1	Truth-functional connectives	37
3.2	Complete truth tables	38
3.3	Using truth tables	41
3.4	Partial truth tables	42
	Practice Exercises	44
4	Quantified logic	48
4.1	From sentences to predicates	48
4.2	Building blocks of QL	50
4.3	Quantifiers	54
4.4	Translating to QL	57
4.5	Sentences of QL	68
4.6	Identity	71
	Practice Exercises	76
5	Formal semantics	82
5.1	Semantics for SL	83

5.2	Interpretations and models in QL	87
5.3	Working with models	91
5.4	Semantics for identity	97
5.5	Truth in QL	98
	Practice Exercises	100
6	Proofs	104
6.1	Basic rules for SL	105
6.2	Derived rules	114
6.3	Rules of replacement	116
6.4	Rules for quantifiers	118
6.5	Rules for identity	123
6.6	Proof strategy	124
6.7	Proof-theoretic concepts	125
6.8	Proofs and models	127
6.9	Soundness and completeness	128
	Practice Exercises	130
A	Other symbolic notation	135
B	Solutions to selected exercises	138
C	Quick reference	149

Chapter 1

What is logic?

Logic is the business of evaluating arguments, sorting good ones from bad ones. In everyday language, we sometimes use the word ‘argument’ to refer to belligerent shouting matches. If you and a friend have an argument in this sense, things are not going well between the two of you.

In logic, we are not interested in the teeth-gnashing, hair-pulling kind of argument. A logical argument is structured to give someone a reason to believe some conclusion. Here is one such argument:

- (1) It is raining heavily.
- (2) If you do not take an umbrella, you will get soaked.
- ∴ You should take an umbrella.

The three dots on the third line of the argument mean ‘Therefore’ and they indicate that the final sentence is the *conclusion* of the argument. The other sentences are *premises* of the argument. If you believe the premises, then the argument provides you with a reason to believe the conclusion.

This chapter discusses some basic logical notions that apply to arguments in a natural language like English. It is important to begin with a clear understanding of what arguments are and of what it means for an argument to be valid. Later we will translate arguments from English into a formal language. We want formal validity, as defined in the formal language, to have at least some of the important features of natural-language validity.

1.1 Arguments

When people mean to give arguments, they typically often use words like ‘therefore’ and ‘because.’ When analyzing an argument, the first thing to do is to separate the premises from the conclusion. Words like these are a clue to what the argument is supposed to be, especially if— in the argument as given— the conclusion comes at the beginning or in the middle of the argument.

premise indicators: since, because, given that

conclusion indicators: therefore, hence, thus, then, so

To be perfectly general, we can define an ARGUMENT as a series of sentences. The sentences at the beginning of the series are premises. The final sentence in the series is the conclusion. If the premises are true and the argument is a good one, then you have a reason to accept the conclusion.

Notice that this definition is quite general. Consider this example:

There is coffee in the coffee pot.
There is a dragon playing bassoon on the armoire.
∴ Salvador Dali was a poker player.

It may seem odd to call this an argument, but that is because it would be a terrible argument. The two premises have nothing at all to do with the conclusion. Nevertheless, given our definition, it still counts as an argument— albeit a bad one.

1.2 Sentences

In logic, we are only interested in sentences that can figure as a premise or conclusion of an argument. So we will say that a SENTENCE is something that can be true or false.

You should not confuse the idea of a sentence that can be true or false with the difference between fact and opinion. Often, sentences in logic will express things that would count as facts— such as ‘Kierkegaard was a hunchback’ or ‘Kierkegaard liked almonds.’ They can also express things that you might think of as matters of opinion— such as, ‘Almonds are yummy.’

Also, there are things that would count as ‘sentences’ in a linguistics or grammar course that we will not count as sentences in logic.

Questions In a grammar class, ‘Are you sleepy yet?’ would count as an interrogative sentence. Although you might be sleepy or you might be alert, the question itself is neither true nor false. For this reason, questions will not count as sentences in logic. Suppose you answer the question: ‘I am not sleepy.’ This is either true or false, and so it is a sentence in the logical sense. Generally, *questions* will not count as sentences, but *answers* will.

‘What is this course about?’ is not a sentence. ‘Noone knows what this course is about’ is a sentence.

Imperatives Commands are often phrased as imperatives like ‘Wake up!’, ‘Sit up straight’, and so on. In a grammar class, these would count as imperative sentences. Although it might be good for you to sit up straight or it might not, the command is neither true nor false. Note, however, that commands are not always phrased as imperatives. ‘If you sit up straight, then you will get a cookie’ is either true or false, and so it counts as a sentence in the logical sense.

Exclamations ‘Ouch!’ is sometimes called an exclamatory sentence, but it is neither true nor false. We will treat ‘Ouch, I hurt my toe!’ as meaning the same thing as ‘I hurt my toe.’ The ‘ouch’ does not add anything that could be true or false.

1.3 Two ways that arguments can go wrong

Consider the argument that you should take an umbrella (on p. 5, above). If premise (1) is false— if it is sunny outside— then the argument gives you no reason to carry an umbrella. Even if it is raining outside, you might not need an umbrella. You might wear a rain pancho or keep to covered walkways. In these cases, premise (2) would be false, since you could go out without an umbrella and still avoid getting soaked.

Suppose for a moment that both the premises are true. You do not own a rain pancho. You need to go places where there are no covered walkways. Now does the argument show you that you should take an umbrella? Not necessarily. Perhaps you enjoy walking in the rain, and you would like to get soaked. In that case, even though the premises were true, the conclusion would be false.

For any argument, there are two ways that it could be weak. First, one or more of the premises might be false. An argument gives you a reason to believe its conclusion only if you believe its premises. Second, the premises might fail to entail the conclusion. Even if the premises were true, the form of the argument might be weak. The example we just considered is weak in both ways.

When an argument is weak in the second way, there is something wrong with the *logical form* of the argument: Premises of the kind given do not necessarily lead to a conclusion of the kind given. We will be interested primarily in the logical form of arguments.

Consider another example:

You are reading this book.
This is a logic book.
 \therefore You are a logic student.

This is not a terrible argument. Most people who read this book are logic students. Yet, it is possible for someone besides a logic student to read this book. If your roommate picked up the book and thumbed through it, they would not immediately become a logic student. So the premises of this argument, even though they are true, do not guarantee the truth of the conclusion. Its logical form is less than perfect.

An argument that had no weakness of the second kind would have perfect logical form. If its premises were true, then its conclusion would *necessarily* be true. We call such an argument ‘deductively valid’ or just ‘valid.’

Even though we might count the argument above as a good argument in some sense, it is not valid; that is, it is ‘invalid.’ One important task of logic is to sort valid arguments from invalid arguments.

1.4 Deductive validity

An argument is deductively VALID if and only if it is impossible for the premises to be true and the conclusion false.

The crucial thing about a valid argument is that it is impossible for the premises to be true *at the same time* that the conclusion is false. Consider this example:

Oranges are either fruits or musical instruments.
Oranges are not fruits.
 \therefore Oranges are musical instruments.

The conclusion of this argument is ridiculous. Nevertheless, it follows validly from the premises. This is a valid argument. *If* both premises were true, *then* the conclusion would necessarily be true.

This shows that a deductively valid argument does not need to have true premises or a true conclusion. Conversely, having true premises and a true conclusion is not enough to make an argument valid. Consider this example:

London is in England.
Beijing is in China.
∴ Paris is in France.

The premises and conclusion of this argument are, as a matter of fact, all true. This is a terrible argument, however, because the premises have nothing to do with the conclusion. Imagine what would happen if Paris declared independence from the rest of France. Then the conclusion would be false, even though the premises would both still be true. Thus, it is *logically possible* for the premises of this argument to be true and the conclusion false. The argument is invalid.

The important thing to remember is that validity is not about the actual truth or falsity of the sentences in the argument. Instead, it is about the form of the argument: The truth of the premises is incompatible with the falsity of the conclusion.

Inductive arguments

There can be good arguments which nevertheless fail to be deductively valid. Consider this one:

In January 1997, it rained in San Diego.
In January 1998, it rained in San Diego.
In January 1999, it rained in San Diego.
∴ It rains every January in San Diego.

This is an **INDUCTIVE** argument, because it generalizes from many cases to a conclusion about all cases.

Certainly, the argument could be made stronger by adding additional premises: In January 2000, it rained in San Diego. In January 2001... and so on. Regardless of how many premises we add, however, the argument will still not be deductively valid. It is possible, although unlikely, that it will fail to rain next January in San Diego. Moreover, we know that the weather can be fickle. No amount of evidence should convince us that it rains there *every* January. Who is to say that some year will not be a freakish year in which there is no rain in January in San Diego; even a single counter-example is enough to make the conclusion of the argument false.

Inductive arguments, even good inductive arguments, are not deductively valid. We will not be interested in inductive arguments in this book.

1.5 Other logical notions

In addition to deductive validity, we will be interested in some other logical concepts.

Truth-values

True or false is said to be the TRUTH-VALUE of sentence. We defined sentences as things that could be true or false; we could have said instead that sentences are things that can have truth-values.

Logical truth

In considering arguments formally, we care about what would be true *if* the premises were true. Generally, we are not concerned with the actual truth value of any particular sentences— whether they are *actually* true or false. Yet there are some sentences that must be true, just as a matter of logic.

Consider these sentences:

1. It is raining.
2. Either it is raining, or it is not.
3. It is both raining and not raining.

In order to know if sentence 1 is true, you would need to look outside or check the weather channel. Logically speaking, it might be either true or false. Sentences like this are called *contingent* sentences.

Sentence 2 is different. You do not need to look outside to know that it is true. Regardless of what the weather is like, it is either raining or not. This sentence is *logically true*; it is true merely as a matter of logic, regardless of what the world is actually like. A logically true sentence is called a TAUTOLOGY.

You do not need to check the weather to know about sentence 3, either. It must be false, simply as a matter of logic. It might be raining here and not raining across town, it might be raining now but stop raining even as you read this, but it is impossible for it to be both raining and not raining here at this moment.

The third sentence is *logically false*; it is false regardless of what the world is like. A logically false sentence is called a CONTRADICTION.

To be precise, we can define a CONTINGENT SENTENCE as a sentence that is neither a tautology nor a contradiction.

A sentence might *always* be true and still be contingent. For instance, if there never were a time when the universe contained fewer than seven things, then the sentence ‘At least seven things exist’ would always be true. Yet the sentence is contingent; its truth is not a matter of logic. There is no contradiction in considering a possible world in which there are fewer than seven things. The important question is whether the sentence *must* be true, just on account of logic.

Logical equivalence

We can also ask about the logical relations *between* two sentences. For example:

John went to the store after he washed the dishes.
John washed the dishes before he went to the store.

These two sentences are both contingent, since John might not have gone to the store or washed dishes at all. Yet they must have the same truth-value. If either of the sentences is true, then they both are; if either of the sentences is false, then they both are. When two sentences necessarily have the same truth value, we say that they are LOGICALLY EQUIVALENT.

Consistency

Consider these two sentences:

B1 My only brother is taller than I am.
B2 My only brother is shorter than I am.

Logic alone cannot tell us which, if either, of these sentences is true. Yet we can say that *if* the first sentence (B1) is true, *then* the second sentence (B2) must be false. And if B2 is true, then B1 must be false. It cannot be the case that both of these sentences are true.

If a set of sentences could not all be true at the same time, like B1–B2, they are said to be INCONSISTENT. Otherwise, they are CONSISTENT.

We can ask about the consistency of any number of sentences. For example, consider the following list of sentences:

- G1** There are at least four giraffes at the wild animal park.
- G2** There are exactly seven gorillas at the wild animal park.
- G3** There are not more than two martians at the wild animal park.
- G4** Every giraffe at the wild animal park is a martian.

G1 and G4 together imply that there are at least four martian giraffes at the park. This conflicts with G3, which implies that there are no more than two martian giraffes there. So the set of sentences G1–G4 is inconsistent. Notice that the inconsistency has nothing at all to do with G2. G2 just happens to be part of an inconsistent set.

Sometimes, people will say that an inconsistent set of sentences ‘contains a contradiction.’ By this, they mean that it would be logically impossible for all of the sentences to be true at once. A set can be inconsistent even when all of the sentences in it are either contingent or tautologous. When a single sentence is a contradiction, then that sentence alone cannot be true.

1.6 Formal languages

Here is a famous valid argument:

- Socrates is a man.
- All men are mortal.
- \therefore Therefore, Socrates is mortal.

This is an iron-clad argument. The only way you could challenge the conclusion is by denying one of the premises— the logical form is impeccable. What about this next argument?

- Socrates is a man.
- All men are carrots.
- \therefore Therefore, Socrates is a carrot.

This argument might be less interesting than the first, because the second premise is obviously false. There is no clear sense in which all men are carrots. Yet the argument is valid. To see this, notice that both arguments have this form:

S is M .
 All M s are C s.
 $\therefore S$ is C .

In both arguments S stands for Socrates and M stands for man. In the first argument, C stands for mortal; in the second, C stands for carrot. Both arguments have this form, and every argument of this form is valid. So both arguments are valid.

What we did here was replace words like ‘man’ or ‘carrot’ with symbols like ‘ M ’ or ‘ C ’ so as to make the logical form explicit. This is the central idea behind formal logic. We want to remove irrelevant or distracting features of the argument to make the logical form more perspicuous.

Starting with an argument in a *natural language* like English, we translate the argument into a *formal language*. Parts of the English sentences are replaced with letters and symbols. The goal is to reveal the formal structure of the argument, as we did with these two.

There are formal languages that work like the symbolization we gave for these two arguments. A logic like this was developed by Aristotle, a philosopher who lived in Greece during the 4th century BC. Aristotle was a student of Plato and the tutor of Alexander the Great. Aristotle’s logic, with some revisions, was the dominant logic in the western world for more than two millennia.

In Aristotelean logic, categories are replaced with capital letters. Every sentence of an argument is then represented as having one of four forms, which medieval logicians labeled in this way: (A) All A s are B s. (E) No A s are B s. (I) Some A is B . (O) Some A is not B .

It is then possible to describe valid *sylogisms*, two-line arguments like the two we considered above. Medieval logicians gave mnemonic names to all of the valid argument forms. The form of our two arguments, for instance, was called *Barbara*. The vowels in the name, all A s, represent the fact that the two premises and the conclusion are all (A) form sentences.

There are many limitations to Aristotelean logic. One is that it makes no distinction between kinds and individuals. So the first premise might just as well be written ‘All S s are M s’: All Socrateses are men. Despite its historical importance, Aristotelean logic has been superseded. The remainder of this book will develop two formal languages.

The first is SL, which stands for *sentential logic*. In SL, the smallest units are sentences themselves. Simple sentences are represented as letters and connected with logical connectives like ‘and’ and ‘not’ to make more complex sentences.

The second is QL, which stands for *quantified logic*. In QL, the basic units are objects, properties of objects, and relations between objects.

When we translate an argument into a formal language, we hope to make its logical structure clearer. We want to include enough of the structure of the English language argument so that we can judge whether the argument is valid or invalid. If we included every feature of the English language, all of the subtlety and nuance, then there would be no advantage in translating to a formal language. We might as well think about the argument in English.

At the same time, we would like a formal language that allows us to represent many kinds of English language arguments. This is one reason to prefer QL to Aristotelean logic; QL can represent every valid argument of Aristotelean logic and more.

So when deciding on a formal language, there is inevitably a tension between wanting to capture as much structure as possible and wanting a simple formal language— simpler formal languages leave out more. This means that there is no perfect formal language. Some will do a better job than others in translating particular English-language arguments.

In this book, we make the assumption that *true* and *false* are the only possible truth-values. Logical languages that make this assumption are called *bivalent*, which means *two-valued*. Aristotelean logic, SL, and QL are all bivalent, but there are limits to the power of bivalent logic. For instance, some philosophers have claimed that the future is not yet determined. If they are right, then sentences about *what will be the case* are not yet true or false. Some formal languages accommodate this by allowing for sentences that are neither true nor false, but something in between. Other formal languages, so-called paraconsistent logics, allow for sentences that are both true *and* false.

The languages presented in this book are not the only possible formal languages. However, most nonstandard logics extend on the basic formal structure of the bivalent logics discussed in this book. So this is a good place to start.

Summary of logical notions

- ▷ An argument is (deductively) VALID if it is impossible for the premises to be true and the conclusion false; it is INVALID otherwise.
- ▷ A TAUTOLOGY is a sentence that must be true, as a matter of logic.
- ▷ A CONTRADICTION is a sentence that must be false, as a matter of logic.
- ▷ A CONTINGENT SENTENCE is neither a tautology nor a contradiction.

- ▷ Two sentences are LOGICALLY EQUIVALENT if they necessarily have the same truth value.
- ▷ A set of sentences is CONSISTENT if it is logically possible for all the members of the set to be true at the same time; it is INCONSISTENT otherwise.

Practice Exercises

At the end of each chapter, you will find a series of practice problems that review and explore the material covered in the chapter. There is no substitute for actually working through some problems, because logic is more about a way of thinking than it is about memorizing facts. The answers to some of the problems are provided at the end of the book in appendix B; the problems that are solved in the appendix are marked with a \star .

Part A Which of the following are ‘sentences’ in the logical sense?

1. England is smaller than China.
2. Greenland is south of Jerusalem.
3. Is New Jersey east of Wisconsin?
4. The atomic number of helium is 2.
5. The atomic number of helium is π .
6. I hate overcooked noodles.
7. Blech! Overcooked noodles!
8. Overcooked noodles are disgusting.
9. Take your time.
10. This is the last question.

Part B For each of the following: Is it a tautology, a contradiction, or a contingent sentence?

1. Caesar crossed the Rubicon.
2. Someone once crossed the Rubicon.
3. No one has ever crossed the Rubicon.
4. If Caesar crossed the Rubicon, then someone has.
5. Even though Caesar crossed the Rubicon, no one has ever crossed the Rubicon.
6. If anyone has ever crossed the Rubicon, it was Caesar.

\star **Part C** Look back at the sentences G1–G4 on p. 12, and consider each of the following sets of sentences. Which are consistent? Which are inconsistent?

1. G2, G3, and G4
2. G1, G3, and G4
3. G1, G2, and G4
4. G1, G2, and G3

★ **Part D** Which of the following is possible? If it is possible, give an example. If it is not possible, explain why.

1. A valid argument that has one false premise and one true premise
2. A valid argument that has a false conclusion
3. A valid argument, the conclusion of which is a contradiction
4. An invalid argument, the conclusion of which is a tautology
5. A tautology that is contingent
6. Two logically equivalent sentences, both of which are tautologies
7. Two logically equivalent sentences, one of which is a tautology and one of which is contingent
8. Two logically equivalent sentences that together are an inconsistent set
9. A consistent set of sentences that contains a contradiction
10. An inconsistent set of sentences that contains a tautology

Chapter 2

Sentential logic

This chapter introduces a logical language called SL. It is a version of *sentential logic*, because the basic units of the language will represent entire sentences.

2.1 Sentence letters

In SL, capital letters are used to represent basic sentences. Considered only as a symbol of SL, the letter *A* could mean any sentence. So when translating from English into SL, it is important to provide a *symbolization key*. The key provides an English language sentence for each sentence letter used in the symbolization.

For example, consider this argument:

There is an apple on the desk.
If there is an apple on the desk, then Jenny made it to class.
 \therefore Jenny made it to class.

This is obviously a valid argument in English. In symbolizing it, we want to preserve the structure of the argument that makes it valid. What happens if we replace each sentence with a letter? Our symbolization key would look like this:

A: There is an apple on the desk.
B: If there is an apple on the desk, then Jenny made it to class.
C: Jenny made it to class.

We would then symbolize the argument in this way:

$$\begin{array}{l} A \\ B \\ \therefore C \end{array}$$

There is no necessary connection between some sentence A , which could be any sentence, and some other sentences B and C , which could be any sentences. The structure of the argument has been completely lost in this translation.

The important thing about the argument is that the second premise is not merely *any* sentence, logically divorced from the other sentences in the argument. The second premise contains the first premise and the conclusion *as parts*. Our symbolization key for the argument only needs to include meanings for A and C , and we can build the second premise from those pieces. So we symbolize the argument this way:

$$\begin{array}{l} A \\ \text{If } A, \text{ then } C. \\ \therefore C \end{array}$$

This preserves the structure of the argument that makes it valid, but it still makes use of the English expression ‘If... then...’ Although we ultimately want to replace all of the English expressions with logical notation, this is a good start.

The sentences that can be symbolized with sentence letters are called *atomic sentences*, because they are the basic building blocks out of which more complex sentences can be built. Whatever logical structure a sentence might have is lost when it is translated as an atomic sentence. From the point of view of SL, the sentence is just a letter. It can be used to build more complex sentences, but it cannot be taken apart.

There are only twenty-six letters of the alphabet, but there is no logical limit to the number of atomic sentences. We can use the same letter to symbolize different atomic sentences by adding a subscript, a small number written after the letter. We could have a symbolization key that looks like this:

A_1 : The apple is under the armoire.
 A_2 : Arguments in SL always contain atomic sentences.
 A_3 : Adam Ant is taking an airplane from Anchorage to Albany.
 \vdots
 A_{294} : Alliteration angers otherwise affable astronauts.

Keep in mind that each of these is a different sentence letter. When there are subscripts in the symbolization key, it is important to keep track of them.

2.2 Connectives

Logical connectives are used to build complex sentences from atomic components. There are five logical connectives in SL. This table summarizes them, and they are explained below.

symbol	what it is called	what it means
\neg	negation	'It is not the case that...'
$\&$	conjunction	'Both... and ...'
\vee	disjunction	'Either... or ...'
\rightarrow	conditional	'If ... then ...'
\leftrightarrow	biconditional	'... if and only if ...'

Negation

Consider how we might symbolize these sentences:

1. Mary is in Barcelona.
2. Mary is not in Barcelona.
3. Mary is somewhere besides Barcelona.

In order to symbolize sentence 1, we will need one sentence letter. We can provide a symbolization key:

B: Mary is in Barcelona.

Note that here we are giving B a different interpretation than we did in the previous section. The symbolization key only specifies what B means *in a specific context*. It is vital that we continue to use this meaning of B so long as we are talking about Mary and Barcelona. Later, when we are symbolizing different sentences, we can write a new symbolization key and use B to mean something else.

Now, sentence 1 is simply B .

Since sentence 2 is obviously related to the sentence 1, we do not want to introduce a different sentence letter. To put it partly in English, the sentence means 'Not B .' In order to symbolize this, we need a symbol for logical negation. We will use ' \neg .' Now we can translate 'Not B ' to $\neg B$.

Sentence 3 is about whether or not Mary is in Barcelona, but it does not contain the word 'not.' Nevertheless, it is obviously logically equivalent to sentence 2.

They both mean: It is not the case that Mary is in Barcelona. As such, we can translate both sentence 2 and sentence 3 as $\neg B$.

A sentence can be symbolized as $\neg\mathcal{A}$ if it can be paraphrased in English as ‘It is not the case that \mathcal{A} .’

Consider these further examples:

4. The widget can be replaced if it breaks.
5. The widget is irreplaceable.
6. The widget is not irreplaceable.

If we let R mean ‘The widget is replaceable’, then sentence 4 can be translated as R .

What about sentence 5? Saying the widget is irreplaceable means that it is not the case that the widget is replaceable. So even though sentence 5 is not negative in English, we symbolize it using negation as $\neg R$.

Sentence 6 can be paraphrased as ‘It is not the case that the widget is irreplaceable.’ Using negation twice, we translate this as $\neg\neg R$. The two negations in a row each work as negations, so the sentence means ‘It is not the case that... it is not the case that... R .’ If you think about the sentence in English, it is logically equivalent to sentence 4. So when we define logical equivalence is SL, we will make sure that R and $\neg\neg R$ are logically equivalent.

More examples:

7. Elliott is happy.
8. Elliott is unhappy.

If we let H mean ‘Elliott is happy’, then we can symbolize sentence 7 as H .

However, it would be a mistake to symbolize sentence 8 as $\neg H$. If Elliott is unhappy, then he is not happy—but sentence 8 does not mean the same thing as ‘It is not the case that Elliott is happy.’ It could be that he is not happy but that he is not unhappy either. Perhaps he is somewhere between the two. In order to symbolize sentence 8, we would need a new sentence letter.

For any sentence \mathcal{A} : If \mathcal{A} is true, then $\neg\mathcal{A}$ is false. If $\neg\mathcal{A}$ is true, then \mathcal{A} is false. Using ‘T’ for true and ‘F’ for false, we can summarize this in a *characteristic truth table* for negation:

\mathcal{A}	$\neg\mathcal{A}$
T	F
F	T

We will discuss truth tables at greater length in the next chapter.

Conjunction

Consider these sentences:

9. Adam is athletic.
10. Barbara is athletic.
11. Adam is athletic, and Barbara is also athletic.

We will need separate sentence letters for 9 and 10, so we define this symbolization key:

- A:** Adam is athletic.
B: Barbara is athletic.

Sentence 9 can be symbolized as A .

Sentence 10 can be symbolized as B .

Sentence 11 can be paraphrased as ‘ A and B .’ In order to fully symbolize this sentence, we need another symbol. We will use ‘&.’ We translate ‘ A and B ’ as $A \& B$. The logical connective ‘&’ is called CONJUNCTION, and A and B are each called CONJUNCTS.

Notice that we make no attempt to symbolize ‘also’ in sentence 11. Words like ‘both’ and ‘also’ function to draw our attention to the fact that two things are being conjoined. They are not doing any further logical work, so we do not need to represent them in SL.

Some more examples:

12. Barbara is athletic and energetic.
13. Barbara and Adam are both athletic.
14. Although Barbara is energetic, she is not athletic.
15. Barbara is athletic, but Adam is more athletic than she is.

Sentence 12 is obviously a conjunction. The sentence says two things about Barbara, so in English it is permissible to refer to Barbara only once. It might

be tempting to try this when translating the argument: Since B means ‘Barbara is athletic’, one might paraphrase the sentences as ‘ B and energetic.’ This would be a mistake. Once we translate part of a sentence as B , any further structure is lost. B is an atomic sentence; it is nothing more than true or false. Conversely, ‘energetic’ is not a sentence; on its own it is neither true nor false. We should instead paraphrase the sentence as ‘ B and Barbara is energetic.’ Now we need to add a sentence letter to the symbolization key. Let E mean ‘Barbara is energetic.’ Now the sentence can be translated as $B \& E$.

A sentence can be symbolized as $\mathcal{A} \& \mathcal{B}$ if it can be paraphrased in English as ‘Both \mathcal{A} , and \mathcal{B} .’ Each of the conjuncts must be a sentence.

Sentence 13 has says one thing about two different subjects. It says of both Barbara and Adam that they are athletic, and in English we use the word ‘athletic’ only once. In translating to SL, it is important to realize that the sentence can be paraphrased as, ‘Barbara is athletic, and Adam is athletic.’ This translates as $B \& A$.

Sentence 14 is a bit more complicated. The word ‘although’ sets up a contrast between the first part of the sentence and the second part. Nevertheless, the sentence says both that Barbara is energetic and that she is not athletic. In order to make each of the conjuncts an atomic sentence, we need to replace ‘she’ with ‘Barbara.’

So we can paraphrase sentence 14 as, ‘Both Barbara is energetic, and Barbara is not athletic.’ The second conjunct contains a negation, so we paraphrase further: ‘Both Barbara is energetic and it is not the case that Barbara is athletic.’ This translates as $E \& \neg B$.

Sentence 15 contains a similar contrastive structure. It is irrelevant for the purpose of translating to SL, so we can paraphrase the sentence as ‘Both Barbara is athletic, and Adam is more athletic than Barbara.’ (Notice that we once again replace the pronoun ‘she’ with her name.) How should we translate the second conjunct? We already have the sentence letter A which is about Adam’s being athletic and B which is about Barbara’s being athletic, but neither is about one of them being more athletic than the other. We need a new sentence letter. Let R mean ‘Adam is more athletic than Barbara.’ Now the sentence translates as $B \& R$.

Sentences that can be paraphrased ‘ \mathcal{A} , but \mathcal{B} ’ or ‘Although \mathcal{A} , \mathcal{B} ’ are best symbolized using conjunction: $\mathcal{A} \& \mathcal{B}$

It is important to keep in mind that the sentence letters A , B , and R are atomic sentences. Considered as symbols of SL, they have no meaning beyond being

true or false. We have used them to symbolize different English language sentences that are all about people being athletic, but this similarity is completely lost when we translate to SL. No formal language can capture all the structure of the English language, but as long as this structure is not important to the argument there is nothing lost by leaving it out.

For any sentences \mathcal{A} and \mathcal{B} , $\mathcal{A} \& \mathcal{B}$ is true if and only if both \mathcal{A} and \mathcal{B} are true. We can summarize this in the characteristic truth table for conjunction:

\mathcal{A}	\mathcal{B}	$\mathcal{A} \& \mathcal{B}$
T	T	T
T	F	F
F	T	F
F	F	F

Conjunction is *symmetrical* because we can swap the conjuncts without changing the truth-value of the sentence. Regardless of what \mathcal{A} and \mathcal{B} are, $\mathcal{A} \& \mathcal{B}$ is logically equivalent to $\mathcal{B} \& \mathcal{A}$.

Disjunction

Consider these sentences:

16. Either Denison will play golf with me, or he will watch movies.
17. Either Denison or Ellery will play golf with me.

For these sentences we can use this symbolization key:

- D:** Denison will play golf with me.
- E:** Ellery will play golf with me.
- M:** Denison will watch movies.

Sentence 16 is ‘Either D or M .’ To fully symbolize this, we introduce a new symbol. The sentence becomes $D \vee M$. The ‘ \vee ’ connective is called DISJUNCTION, and D and M are called DISJUNCTS.

Sentence 17 is only slightly more complicated. There are two subjects, but the English sentence only gives the verb once. In translating, we can paraphrase it as. ‘Either Denison will play golf with me, or Ellery will play golf with me.’ Now it obviously translates as $D \vee E$.

A sentence can be symbolized as $\mathcal{A} \vee \mathcal{B}$ if it can be paraphrased in English as ‘Either \mathcal{A} , or \mathcal{B} .’ Each of the disjuncts must be a sentence.

Sometimes in English, the word ‘or’ excludes the possibility that both disjuncts are true. This is called an EXCLUSIVE OR. An *exclusive or* is clearly intended when it says, on a restaurant menu, ‘Entrees come with either soup or salad.’ You may have soup; you may have salad; but, if you want *both* soup *and* salad, then you have to pay extra.

At other times, the word ‘or’ allows for the possibility that both disjuncts might be true. This is probably the case with sentence 17, above. I might play with Denison, with Ellery, or with both Denison and Ellery. Sentence 17 merely says that I will play with *at least* one of them. This is called an INCLUSIVE OR.

The symbol ‘ \vee ’ represents an *inclusive or*. So $D \vee E$ is true if D is true, if E is true, or if both D and E are true. It is false only if both D and E are false. We can summarize this with the characteristic truth table for disjunction:

\mathcal{A}	\mathcal{B}	$\mathcal{A} \vee \mathcal{B}$
T	T	T
T	F	T
F	T	T
F	F	F

Like conjunction, disjunction is symmetrical. $\mathcal{A} \vee \mathcal{B}$ is logically equivalent to $\mathcal{B} \vee \mathcal{A}$.

These sentences are somewhat more complicated:

- 18. Either you will not have soup, or you will not have salad.
- 19. You will have neither soup nor salad.
- 20. You get either soup or salad, but not both.

We let S_1 mean that you get soup and S_2 mean that you get salad.

Sentence 18 can be paraphrased in this way: ‘Either *it is not the case that* you get soup, or *it is not the case that* you get salad.’ Translating this requires both disjunction and negation. It becomes $\neg S_1 \vee \neg S_2$.

Sentence 19 also requires negation. It can be paraphrased as, ‘*It is not the case that* either that you get soup or that you get salad.’ We need some way of indicating that the negation does not just negate the right or left disjunct, but rather negates the entire disjunction. In order to do this, we put parentheses

around the disjunction: ‘It is not the case that $(S_1 \vee S_2)$.’ This becomes simply $\neg(S_1 \vee S_2)$.

Notice that the parentheses are doing important work here. The sentence $\neg S_1 \vee S_2$ would mean ‘Either you will not have soup, or you will have salad.’

Sentence 20 is an *exclusive or*. We can break the sentence into two parts. The first part says that you get one or the other. We translate this as $(S_1 \vee S_2)$. The second part says that you do not get both. We can paraphrase this as, ‘It is not the case both that you get soup and that you get salad.’ Using both negation and conjunction, we translate this as $\neg(S_1 \& S_2)$. Now we just need to put the two parts together. As we saw above, ‘but’ can usually be translated as a conjunction. Sentence 20 can thus be translated as $(S_1 \vee S_2) \& \neg(S_1 \& S_2)$.

Although ‘ \vee ’ is an *inclusive or*, we can symbolize an *exclusive or* in SL. We just need more than one connective to do it.

Conditional

For the following sentences, let R mean ‘You will cut the red wire’ and B mean ‘The bomb will explode.’

21. If the you cut the red wire, then the bomb will explode.
22. The bomb will explode only if you cut the red wire.

Sentence 21 can be translated partially as ‘If R , then B .’ We will use the symbol ‘ \rightarrow ’ to represent logical entailment. The sentence becomes $R \rightarrow B$. The connective is called a **CONDITIONAL**. The sentence on the left-hand side of the conditional (R in this example) is called the **ANTECEDENT**. The sentence on the right-hand side (B) is called the **CONSEQUENT**.

Sentence 22 is also a conditional. Since the word ‘if’ appears in the second half of the sentence, it might be tempting to symbolize this in the same way as sentence 21. That would be a mistake.

The conditional $R \rightarrow B$ says that *if* R were true, *then* B would also be true. It does not say that your cutting the red wire is the *only* way that the bomb could explode. Someone else might cut the wire, or the bomb might be on a timer. The sentence $R \rightarrow B$ does not say anything about what to expect if R is false. Sentence 22 is different. It says that the only conditions under which the bomb will explode involve your having cut the red wire; i.e., if the bomb explodes, then you must have cut the wire. As such, sentence 22 should be symbolized as $B \rightarrow R$.

It is important to remember that the connective ‘ \rightarrow ’ says only that, if the antecedent is true, then the consequent is true. It says nothing about the *causal* connection between the two events. Translating sentence 22 as $B \rightarrow R$ does not mean that the bomb exploding would somehow have caused your cutting the wire. Both sentence 21 and 22 suggest that, if you cut the red wire, your cutting the red wire would be the cause of the bomb exploding. They differ on the *logical* connection. If sentence 22 were true, then an explosion would tell us—those of us safely away from the bomb—that you had cut the red wire. Without an explosion, sentence 22 tells us nothing.

The paraphrased sentence ‘ \mathcal{A} only if \mathcal{B} ’ is logically equivalent to ‘If \mathcal{A} , then \mathcal{B} .’

‘If \mathcal{A} then \mathcal{B} ’ means that if \mathcal{A} is true then so is \mathcal{B} . So we know that if the antecedent \mathcal{A} is true but the consequent \mathcal{B} is false, then the conditional ‘If \mathcal{A} then \mathcal{B} ’ is false. What is the truth value of ‘If \mathcal{A} then \mathcal{B} ’ under other circumstances? Suppose, for instance, that the antecedent \mathcal{A} happened to be false. ‘If \mathcal{A} then \mathcal{B} ’ would then not tell us anything about the actual truth value of the consequent \mathcal{B} , and it is unclear what the truth value of ‘If \mathcal{A} then \mathcal{B} ’ would be.

In English, the truth of conditionals often depends on what *would* be the case if the antecedent *were true*—even if, as a matter of fact, the antecedent is false. This poses a problem for translating conditionals into SL. Considered as sentences of SL, R and B in the above examples have nothing intrinsic to do with each other. In order to consider what the world would be like if R were true, we would need to analyze what R says about the world. Since R is an atomic symbol of SL, however, there is no further structure to be analyzed. When we replace a sentence with a sentence letter, we consider it merely as some atomic sentence that might be true or false.

In order to translate conditionals into SL, we will not try to capture all the subtleties of the English language ‘If... then...’ Instead, the symbol ‘ \rightarrow ’ will be a *material conditional*. This means that when \mathcal{A} is false, the conditional $\mathcal{A} \rightarrow \mathcal{B}$ is automatically true, regardless of the truth value of \mathcal{B} . If both \mathcal{A} and \mathcal{B} are true, then the conditional $\mathcal{A} \rightarrow \mathcal{B}$ is true.

In short, $\mathcal{A} \rightarrow \mathcal{B}$ is false if and only if \mathcal{A} is true and \mathcal{B} is false. We can summarize this with a characteristic truth table for the conditional.

\mathcal{A}	\mathcal{B}	$\mathcal{A} \rightarrow \mathcal{B}$
T	T	T
T	F	F
F	T	T
F	F	T

The conditional is *asymmetrical*. You cannot swap the antecedent and consequent without changing the meaning of the sentence, because $\mathcal{A} \rightarrow \mathcal{B}$ and $\mathcal{B} \rightarrow \mathcal{A}$ are not logically equivalent.

Not all sentences of the form ‘If... then...’ are conditionals. Consider this sentence:

23. If anyone wants to see me, then I will be on the porch.

If I say this, it means that I will be on the porch, regardless of whether anyone wants to see me or not— but if someone did want to see me, then they should look for me there. If we let P mean ‘I will be on the porch,’ then sentence 23 can be translated simply as P .

Biconditional

Consider these sentences:

24. The figure on the board is a triangle only if it has exactly three sides.
 25. The figure on the board is a triangle if it has exactly three sides.
 26. The figure on the board is a triangle if and only if it has exactly three sides.

Let T mean ‘The figure is a triangle’ and S mean ‘The figure has three sides.’

Sentence 24, for reasons discussed above, can be translated as $T \rightarrow S$.

Sentence 25 is importantly different. It can be paraphrased as, ‘If the figure has three sides, then it is a triangle.’ So it can be translated as $S \rightarrow T$.

Sentence 26 says that T is true *if and only if* S is true; we can infer S from T , and we can infer T from S . This is called a BICONDITIONAL, because it entails the two conditionals $S \rightarrow T$ and $T \rightarrow S$. We will use ‘ \leftrightarrow ’ to represent the biconditional; sentence 26 can be translated as $S \leftrightarrow T$.

We could abide without a new symbol for the biconditional. Since sentence 26 means ‘ $T \rightarrow S$ and $S \rightarrow T$,’ we could translate it as $(T \rightarrow S) \& (S \rightarrow T)$. We would need parentheses to indicate that $(T \rightarrow S)$ and $(S \rightarrow T)$ are separate conjuncts; the expression $T \rightarrow S \& S \rightarrow T$ would be ambiguous.

Because we could always write $(\mathcal{A} \rightarrow \mathcal{B}) \& (\mathcal{B} \rightarrow \mathcal{A})$ instead of $\mathcal{A} \leftrightarrow \mathcal{B}$, we do not strictly speaking *need* to introduce a new symbol for the biconditional. Nevertheless, logical languages usually have such a symbol. SL will have one, which makes it easier to translate phrases like ‘if and only if.’

$\mathcal{A} \leftrightarrow \mathcal{B}$ is true if and only if \mathcal{A} and \mathcal{B} have the same truth value. This is the characteristic truth table for the biconditional:

\mathcal{A}	\mathcal{B}	$\mathcal{A} \leftrightarrow \mathcal{B}$
T	T	T
T	F	F
F	T	F
F	F	T

2.3 Other symbolization

We have now introduced all of the connectives of SL. We can use them together to translate many kinds of sentences. Consider these examples of sentences that use the English-language connective ‘unless’:

27. Unless you wear a jacket, you will catch cold.
 28. You will catch cold unless you wear a jacket.

Let J mean ‘You will wear a jacket’ and let D mean ‘You will catch a cold.’

We can paraphrase sentence 27 as ‘Unless J , D .’ This means that if you do not wear a jacket, then you will catch cold; with this in mind, we might translate it as $\neg J \rightarrow D$. It also means that if you do not catch a cold, then you must have worn a jacket; with this in mind, we might translate it as $\neg D \rightarrow J$.

Which of these is the correct translation of sentence 27? Both translations are correct, because the two translations are logically equivalent in SL.

Sentence 28, in English, is logically equivalent to sentence 27. It can be translated as either $\neg J \rightarrow D$ or $\neg D \rightarrow J$.

When symbolizing sentences like sentence 27 and sentence 28, it is easy to get turned around. Since the conditional is not symmetric, it would be wrong to translate either sentence as $J \rightarrow \neg D$. Fortunately, there are other logically equivalent expressions. Both sentences mean that you will wear a jacket or—if you do not wear a jacket—then you will catch a cold. So we can translate them as $J \vee D$. (You might worry that the ‘or’ here should be an exclusive or. However, the sentences do not exclude the possibility that you might *both* wear a jacket *and* catch a cold; jackets do not protect you from all the possible ways that you might catch a cold.)

If a sentence can be paraphrased as ‘Unless \mathcal{A} , \mathcal{B} ,’ then it can be symbolized as $\mathcal{A} \vee \mathcal{B}$.

Symbolization of standard sentence types is summarized on p. 149.

2.4 Sentences of SL

The sentence ‘Apples are red, or berries are blue’ is a sentence of English, and the sentence ‘ $(A \vee B)$ ’ is a sentence of SL. Although we can identify sentences of English when we encounter them, we do not have a formal definition of ‘sentence of English’. In SL, it is possible to formally define what counts as a sentence. This is one respect in which a formal language like SL is more precise than a natural language like English.

It is important to distinguish between the logical language SL, which we are developing, and the language that we use to talk about SL. When we talk about a language, the language that we are talking about is called the OBJECT LANGUAGE. The language that we use to talk about the object language is called the METALANGUAGE.

The object language in this chapter is SL. The metalanguage is English— not conversational English, but English supplemented with some logical and mathematical vocabulary. The sentence ‘ $(A \vee B)$ ’ is a sentence in the object language, because it uses only symbols of SL. The word ‘sentence’ is not itself part of SL, however, so the sentence ‘This expression is a sentence of SL’ is not a sentence of SL. It is a sentence in the metalanguage, a sentence that we use to talk *about* SL.

In this section, we will give a formal definition for ‘sentence of SL.’ The definition itself will be given in mathematical English, the metalanguage.

Expressions

There are three kinds of symbols in SL:

sentence letters with subscripts, as needed	A, B, C, \dots, Z $A_1, B_1, Z_1, A_2, A_{25}, J_{375}, \dots$
connectives	$\neg, \&, \vee, \rightarrow, \leftrightarrow$
parentheses	$(,)$

We define an EXPRESSION OF SL as any string of symbols of SL. Take any of the symbols of SL and write them down, in any order, and you have an expression.

Well-formed formulae

Since any sequence of symbols is an expression, many expressions of SL will be gobbleddegook. A meaningful expression is called a *well-formed formula*. It is common to use the acronym *wff*; the plural is wffs.

Obviously, individual sentence letters like A and G_{13} will be wffs. We can form further wffs out of these by using the various connectives. Using negation, we can get $\neg A$ and $\neg G_{13}$. Using conjunction, we can get $A \& G_{13}$, $G_{13} \& A$, $A \& A$, and $G_{13} \& G_{13}$. We could also apply negation repeatedly to get wffs like $\neg\neg A$ or apply negation along with conjunction to get wffs like $\neg(A \& G_{13})$ and $\neg(G_{13} \& \neg G_{13})$. The possible combinations are endless, even starting with just these two sentence letters, and there are infinitely many sentence letters. So there is no point in trying to list all the wffs.

Instead, we will describe the process by which wffs can be constructed. Consider negation: Given any wff \mathcal{A} of SL, $\neg\mathcal{A}$ is a wff of SL. It is important here that \mathcal{A} is not the sentence letter A . Rather, it is a variable that stands in for any wff at all. Notice that this variable \mathcal{A} is not a symbol of SL, so $\neg\mathcal{A}$ is not an expression of SL. Instead, it is an expression of the metalanguage that allows us to talk about infinitely many expressions of SL: all of the expressions that start with the negation symbol. Because \mathcal{A} is part of the metalanguage, it is called a *metavariable*.

We can say similar things for each of the other connectives. For instance, if \mathcal{A} and \mathcal{B} are wffs of SL, then $(\mathcal{A} \& \mathcal{B})$ is a wff of SL. Providing clauses like this for all of the connectives, we arrive at the following formal definition for a well-formed formula of SL:

1. Every atomic sentence is a wff.
2. If \mathcal{A} is a wff, then $\neg\mathcal{A}$ is a wff of SL.
3. If \mathcal{A} and \mathcal{B} are wffs, then $(\mathcal{A} \& \mathcal{B})$, is a wff.
4. If \mathcal{A} and \mathcal{B} are wffs, then $(\mathcal{A} \vee \mathcal{B})$ is a wff.
5. If \mathcal{A} and \mathcal{B} are wffs, then $(\mathcal{A} \rightarrow \mathcal{B})$ is a wff.
6. If \mathcal{A} and \mathcal{B} are wffs, then $(\mathcal{A} \leftrightarrow \mathcal{B})$ is a wff.
7. All and only wffs of SL can be generated by applications of these rules.

Notice that we cannot immediately apply this definition to see whether an arbitrary expression is a wff. Suppose we want to know whether or not $\neg\neg\neg D$ is a wff of SL. Looking at the second clause of the definition, we know that $\neg\neg\neg D$ is a wff *if* $\neg\neg D$ is a wff. So now we need to ask whether or not $\neg\neg D$

is a wff. Again looking at the second clause of the definition, $\neg\neg D$ is a wff if $\neg D$ is. Again, $\neg D$ is a wff if D is a wff. Now D is a sentence letter, an atomic sentence of SL, so we know that D is a wff by the first clause of the definition. So for a compound formula like $\neg\neg\neg D$, we must apply the definition repeatedly. Eventually we arrive at the atomic sentences from which the wff is built up.

Definitions like this are called *recursive*. Recursive definitions begin with some specifiable base elements and define ways to indefinitely compound the base elements. To determine whether or not something meets the definition, you may have to refer back to the definition many times.

The connective that you look to first in decomposing a sentence is called the MAIN LOGICAL OPERATOR of that sentence. For example. The main logical operator of $\neg(E \vee (F \rightarrow G))$ is negation, \neg . The main logical operator of $(\neg E \vee (F \rightarrow G))$ is disjunction, \vee .

Sentences

Recall that a sentence is a meaningful expression that can be true or false. Since the meaningful expressions of SL are the wffs and since every wff of SL is either true or false, the definition for a sentence of SL is the same as the definition for a wff. Not every formal language will have this nice feature. In the language QL, which is developed later in the book, there are wffs which are not sentences.

The recursive structure of sentences in SL will be important when we consider the circumstances under which a particular sentence would be true or false. The sentence $\neg\neg\neg D$ is true if and only if the sentence $\neg\neg D$ is false, and so on through the structure of the sentence until we arrive at the atomic components: $\neg\neg\neg D$ is true if and only if the atomic sentence D is false. We will return to this point in the next chapter.

Notational conventions

A wff like $(Q \& R)$ must be surrounded by parentheses, because we might apply the definition again to use this as part of a more complicated sentence. If we negate $(Q \& R)$, we get $\neg(Q \& R)$. If we just had $Q \& R$ without the parentheses and put a negation in front of it, we would have $\neg Q \& R$. It is most natural to read this as meaning the same thing as $(\neg Q \& R)$, something very different than $\neg(Q \& R)$. The sentence $\neg(Q \& R)$ means that it is not the case that both Q and R are true; Q might be false or R might be false, but the sentence does not tell us which. The sentence $(\neg Q \& R)$ means specifically that Q is false and that R is true. As such, parentheses are crucial to the meaning of the sentence.

So, strictly speaking, $Q \& R$ without parentheses is *not* a sentence of SL. When using SL, however, we will often be able to relax the precise definition so as to make things easier for ourselves. We will do this in several ways.

First, we understand that $Q \& R$ means the same thing as $(Q \& R)$. As a matter of convention, we can leave off parentheses that occur *around the entire sentence*.

Second, it can sometimes be confusing to look at long sentences with many, nested pairs of parentheses. We adopt the convention of using square brackets '[' and ']' in place of parenthesis. There is no logical difference between $(P \vee Q)$ and $[P \vee Q]$, for example. The unwieldy sentence

$$(((H \rightarrow I) \vee (I \rightarrow H)) \& (J \vee K))$$

could be written in this way:

$$[(H \rightarrow I) \vee (I \rightarrow H)] \& (J \vee K)$$

Third, we will sometimes want to translate the conjunction of three or more sentences. For the sentence 'Alice, Bob, and Candice all went to the party', suppose we let A mean 'Alice went', B mean 'Bob went', and C mean 'Candice went.' The definition only allows us to form a conjunction out of two sentences, so we can translate it as $(A \& B) \& C$ or as $A \& (B \& C)$. There is no reason to distinguish between these, since the two translations are logically equivalent. There is no logical difference between the first, in which $(A \& B)$ is conjoined with C , and the second, in which A is conjoined with $(B \& C)$. So we might as well just write $A \& B \& C$. As a matter of convention, we can leave out parentheses when we conjoin three or more sentences.

Fourth, a similar situation arises with multiple disjunctions. 'Either Alice, Bob, or Candice went to the party' can be translated as $(A \vee B) \vee C$ or as $A \vee (B \vee C)$. Since these two translations are logically equivalent, we may write $A \vee B \vee C$.

These latter two conventions only apply to multiple conjunctions or multiple disjunctions. If a series of connectives includes both disjunctions and conjunctions, then the parentheses are essential; as with $(A \& B) \vee C$ and $A \& (B \vee C)$. The parentheses are also required if there is a series of conditionals or biconditionals; as with $(A \rightarrow B) \rightarrow C$ and $A \leftrightarrow (B \leftrightarrow C)$.

We have adopted these four rules as *notational conventions*, not as changes to the definition of a sentence. Strictly speaking, $A \vee B \vee C$ is still not a sentence. Instead, it is a kind of shorthand. We write it for the sake of convenience, but we really mean the sentence $(A \vee (B \vee C))$.

If we had given a different definition for a wff, then these could count as wffs. We might have written rule 3 in this way: "If \mathcal{A} , \mathcal{B} , ... \mathcal{Z} are wffs, then $(\mathcal{A} \& \mathcal{B} \& \dots \& \mathcal{Z})$, is a wff." This would make it easier to translate some English sentences, but would have the cost of making our formal language more

complicated. We would have to keep the complex definition in mind when we develop truth tables and a proof system. We want a logical language that is *expressively simple* and allows us to translate easily from English, but we also want a *formally simple* language. Adopting notational conventions is a compromise between these two desires.

Practice Exercises

★ **Part A** Using the symbolization key given, translate each English-language sentence into SL.

- M:** Those creatures are men in suits.
- C:** Those creatures are chimpanzees.
- G:** Those creatures are gorillas.

1. Those creatures are not men in suits.
2. Those creatures are men in suits, or they are not.
3. Those creatures are either gorillas or chimpanzees.
4. Those creatures are neither gorillas nor chimpanzees.
5. If those creatures are chimpanzees, then they are neither gorillas nor men in suits.
6. Unless those creatures are men in suits, they are either chimpanzees or they are gorillas.

Part B Using the symbolization key given, translate each English-language sentence into SL.

- A:** Mister Ace was murdered.
- B:** The butler did it.
- C:** The cook did it.
- D:** The Duchess is lying.
- E:** Mister Edge was murdered.
- F:** The murder weapon was a frying pan.

1. Either Mister Ace or Mister Edge was murdered.
2. If Mister Ace was murdered, then the cook did it.
3. If Mister Edge was murdered, then the cook did not do it.
4. Either the butler did it, or the Duchess is lying.
5. The cook did it only if the Duchess is lying.
6. If the murder weapon was a frying pan, then the culprit must have been the cook.

7. If the murder weapon was not a frying pan, then the culprit was either the cook or the butler.
8. Mister Ace was murdered if and only if Mister Edge was not murdered.
9. The Duchess is lying, unless it was Mister Edge who was murdered.
10. If Mister Ace was murdered, he was done in with a frying pan.
11. The cook did it, so the butler did not.
12. Of course the Duchess is lying!

★ **Part C** Using the symbolization key given, translate each english-language sentence into SL.

- E**₁: Ava is an electrician.
E₂: Harrison is an electrician.
F₁: Ava is a firefighter.
F₂: Harrison is a firefighter.
S₁: Ava is satisfied with her career.
S₂: Harrison is satisfied with his career.

1. Ava and Harrison are both electricians.
2. If Ava is a firefighter, then she is satisfied with her career.
3. Ava is a firefighter, unless she is an electrician.
4. Harrison is an unsatisfied electrician.
5. Neither Ava nor Harrison is an electrician.
6. Both Ava and Harrison are electricians, but neither of them find it satisfying.
7. Harrison is satisfied only if he is a firefighter.
8. If Ava is not an electrician, then neither is Harrison, but if she is, then he is too.
9. Ava is satisfied with her career if and only if Harrison is not satisfied with his.
10. If Harrison is both an electrician and a firefighter, then he must be satisfied with his work.
11. It cannot be that Harrison is both an electrician and a firefighter.
12. Harrison and Ava are both firefighters if and only if neither of them is an electrician.

★ **Part D** Give a symbolization key and symbolize the following sentences in SL.

1. Alice and Bob are both spies.
2. If either Alice or Bob is a spy, then the code has been broken.
3. If neither Alice nor Bob is a spy, then the code remains unbroken.
4. The German embassy will be in an uproar, unless someone has broken the code.

5. Either the code has been broken or it has not, but the German embassy will be in an uproar regardless.
6. Either Alice or Bob is a spy, but not both.

Part E Give a symbolization key and symbolize the following sentences in SL.

1. If Gregor plays first base, then the team will lose.
2. The team will lose unless there is a miracle.
3. The team will either lose or it won't, but Gregor will play first base regardless.
4. Gregor's mom will bake cookies if and only if Gregor plays first base.
5. If there is a miracle, then Gregor's mom will not bake cookies.

Part F For each argument, write a symbolization key and translate the argument as well as possible into SL.

1. If Dorothy plays the piano in the morning, then Roger wakes up cranky. Dorothy plays piano in the morning unless she is distracted. So if Roger does not wake up cranky, then Dorothy must be distracted.
2. It will either rain or snow on Tuesday. If it rains, Neville will be sad. If it snows, Neville will be cold. Therefore, Neville will either be sad or cold on Tuesday.
3. If Zoog remembered to do his chores, then things are clean but not neat. If he forgot, then things are neat but not clean. Therefore, things are either neat or clean— but not both.

★ **Part G** For each of the following: (a) Is it a wff of SL? (b) Is it a sentence of SL, allowing for notational conventions?

1. (A)
2. $J_{374} \vee \neg J_{374}$
3. $\neg\neg\neg\neg F$
4. $\neg \& S$
5. $(G \& \neg G)$
6. $\mathcal{A} \rightarrow \mathcal{A}$
7. $(A \rightarrow (A \& \neg F)) \vee (D \leftrightarrow E)$
8. $[(Z \leftrightarrow S) \rightarrow W] \& [J \vee X]$
9. $(F \leftrightarrow \neg D \rightarrow J) \vee (C \& D)$

Part H

1. Are there any wffs of SL that contain no sentence letters? Why or why not?

2. In the chapter, we symbolized an *exclusive or* using \vee , $\&$, and \neg . How could you translate an *exclusive or* using only two connectives? Is there any way to translate an *exclusive or* using only one connective?

Chapter 3

Truth tables

3.1 Truth-functional connectives

Any non-atomic sentence of SL is composed of atomic sentences with sentential connectives. The truth-value of the compound sentence depends only on the truth-value of the atomic sentences that comprise it. In order to know the truth-value of $(D \leftrightarrow E)$, for instance, you only need to know the truth-value of D and the truth-value of E . Connectives that work in this way are called TRUTH-FUNCTIONAL.

In this chapter, we will make use of the fact that all of the logical operators in SL are truth-functional— it makes it possible to construct truth tables to determine the logical features of sentences. You should realize, however, that this is not possible for all languages. In English, it is possible to form a new sentence from any simpler sentence \mathcal{X} by saying ‘It is possible that \mathcal{X} .’ The truth-value of this new sentence does not depend directly on the truth-value of \mathcal{X} . Even if \mathcal{X} is false, perhaps in some sense \mathcal{X} *could* have been true— then the new sentence would be true. Some formal languages, called *modal logics*, have an operator for possibility. In a modal logic, we could translate ‘It is possible that \mathcal{X} ’ as $\diamond\mathcal{X}$. However, the ability to translate sentences like these come at a cost: The \diamond operator is not truth-functional, and so modal logics are not amenable to truth tables.

\mathcal{A}	$\neg\mathcal{A}$	\mathcal{A}	\mathcal{B}	$\mathcal{A} \& \mathcal{B}$	$\mathcal{A} \vee \mathcal{B}$	$\mathcal{A} \rightarrow \mathcal{B}$	$\mathcal{A} \leftrightarrow \mathcal{B}$
T	F	T	T	T	T	T	T
T	F	T	F	F	T	F	F
F	T	F	T	F	T	T	F
F	T	F	F	F	F	T	T

Table 3.1: The characteristic truth tables for the connectives of SL.

3.2 Complete truth tables

The truth-value of sentences that contain only one connective is given by the characteristic truth table for that connective. To put them all in one place, the truth tables for the connectives of SL are repeated in table 3.1.

The characteristic truth table for conjunction, for example, gives the truth conditions for any sentence of the form $(\mathcal{A} \& \mathcal{B})$. Even if the conjuncts \mathcal{A} and \mathcal{B} are long, complicated sentences, the conjunction is true if and only if both \mathcal{A} and \mathcal{B} are true. Consider the sentence $(H \& I) \rightarrow H$. We consider all the possible combinations of true and false for H and I , which gives us four rows. We then copy the truth-values for the sentence letters and write them underneath the letters in the sentence.

H	I	$(H \& I) \rightarrow H$		
T	T	T	T	T
T	F	T	F	T
F	T	F	T	F
F	F	F	F	F

Now consider the subsentence $H \& I$. This is a conjunction $\mathcal{A} \& \mathcal{B}$ with H as \mathcal{A} and with I as \mathcal{B} . H and I are both true on the first row. Since a conjunction is true when both conjuncts are true, we write a T underneath the conjunction symbol. We continue for the other three rows and get this:

H	I	$(H \& I) \rightarrow H$		
		\mathcal{A}	$\&$	\mathcal{B}
T	T	T	T	T
T	F	T	F	F
F	T	F	F	T
F	F	F	F	F

The entire sentence is a conditional $\mathcal{A} \rightarrow \mathcal{B}$ with $(H \& I)$ as \mathcal{A} and with H as \mathcal{B} . On the second row, for example, $(H \& I)$ is false and H is true. Since a conditional is true when the antecedent is false, we write a T in the second row

underneath the conditional symbol. We continue for the other three rows and get this:

H	I	$(H \ \& \ I) \rightarrow H$	
		\mathcal{A}	$\rightarrow \mathcal{B}$
T	T	T	T T
T	F	F	T T
F	T	F	T F
F	F	F	T F

The column of Ts underneath the conditional tells us that the sentence $(H \ \& \ I) \rightarrow I$ is true regardless of the truth-values of H and I . They can be true or false in any combination, and the compound sentence still comes out true. It is crucial that we have considered all of the possible combinations. If we only had a two-line truth table, we could not be sure that the sentence was not false for some other combination of truth-values.

In this example, we have not repeated all of the entries in every successive table. When actually writing truth tables on paper, however, it is impractical to erase whole columns or rewrite the whole table for every step. Although it is more crowded, the truth table can be written in this way:

H	I	$(H \ \& \ I) \rightarrow H$			
T	T	T	T	T	T
T	F	T	F	F	T
F	T	F	F	T	T
F	F	F	F	F	T

Most of the columns underneath the sentence are only there for bookkeeping purposes. When you become more adept with truth tables, you will probably no longer need to copy over the columns for each of the sentence letters. In any case, the truth-value of the sentence on each row is just the column underneath the main logical operator of the sentence; in this case, the column underneath the conditional.

A COMPLETE TRUTH TABLE has a row for all the possible combinations of T and F for all of the sentence letters. The size of the complete truth table depends on the number of different sentence letters in the table. A sentence that contains only one sentence letter requires only two rows, as in the characteristic truth table for negation. This is true even if the same letter is repeated many times, as in the sentence $[(C \leftrightarrow C) \rightarrow C] \ \& \ \neg(C \rightarrow C)$. The complete truth table requires only two lines because there are only two possibilities: C can be true or it can be false. A single sentence letter can never be marked both T and F on the same row. The truth table for this sentence looks like this:

C	$[(C \leftrightarrow C) \rightarrow C] \& \neg(C \rightarrow C)$
T	T T T T T F F T T T
F	F T F F F F F F T F

Looking at the column underneath the main connective, we see that the sentence is false on both rows of the table; i.e., it is false regardless of whether C is true and false.

A sentence that contains two sentence letters requires four lines for a complete truth table, as in the characteristic truth tables and the table for $(H \& I) \rightarrow I$.

A sentence that contains three sentence letters requires eight lines. For example:

M	N	P	$M \& (N \vee P)$
T	T	T	T T T T T
T	T	F	T T T T F
T	F	T	T T F T T
T	F	F	T F F F F
F	T	T	F F T T T
F	T	F	F F T T F
F	F	T	F F F T T
F	F	F	F F F F F

From this table, we know that the sentence $M \& (N \vee P)$ might be true or false, depending on the truth-values of M , N , and P .

A complete truth table for a sentence that contains four different sentence letters requires 16 lines. Five letters, 32 lines. Six letters, 64 lines. And so on. To be perfectly general: If a complete truth table has n different sentence letters, then it must have 2^n rows.

In order to fill in the columns of a complete truth table, begin with the right-most sentence letter and alternate Ts and Fs. In the next column to the left, write two Ts, write two Fs, and repeat. For the third sentence letter, write four Ts followed by four Fs. This yields an eight line truth table like the one above. For a 16 line truth table, the next column of sentence letters should have eight Ts followed by eight Fs. For a 32 line table, the next column would have 16 Ts followed by 16 Fs. And so on.

3.3 Using truth tables

Tautologies, contradictions, and contingent sentences

Recall that an English sentence is a tautology if it must be true as a matter of logic. With a complete truth table, we consider all of the ways that the world might be. If the sentence is true on every line of a complete truth table, then it is true as a matter of logic, regardless of what the world is like.

So a sentence is a **TAUTOLOGY IN SL** if the column under its main connective is T on every row of a complete truth table.

Conversely, a sentence is a **CONTRADICTION IN SL** if the column under its main connective is F on every row of a complete truth table.

A sentence is **CONTINGENT IN SL** if it is neither a tautology nor a contradiction; i.e. if it is T on at least one row and F on at least one row.

From the truth tables in the previous section, we know that $(H \& I) \rightarrow H$ is a tautology, that $[(C \leftrightarrow C) \rightarrow C] \& \neg(C \rightarrow C)$ is a contradiction, and that $M \& (N \vee P)$ is contingent.

Logical equivalence

Two sentences are logically equivalent in English if they have the same truth value as a matter of logic. Once again, truth tables allow us to define an analogous concept for SL: Two sentences are **LOGICALLY EQUIVALENT IN SL** if they have the same truth-value on every row of a complete truth table.

Consider the sentences $\neg(A \vee B)$ and $\neg A \& \neg B$. Are they logically equivalent? To find out, we construct a truth table.

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

Look at the columns for the main connectives; negation for the first sentence, conjunction for the second. On the first three rows, both are F. On the final row, both are T. Since they match on every row, the two sentences are logically equivalent.

Consistency

A set of sentences in English is consistent if it is logically possible for them all to be true at once. A set of sentences is LOGICALLY CONSISTENT IN SL if there is at least one line of a complete truth table on which all of the sentences are true. It is INCONSISTENT otherwise.

Validity

An argument in English is valid if it is logically impossible for the premises to be true and for the conclusion to be false at the same time. An argument is VALID IN SL if there is no row of a complete truth table on which the premises are all T and the conclusion is F; an argument is INVALID IN SL if there is such a row.

Consider this argument:

$$\begin{array}{l} \neg L \rightarrow (J \vee L) \\ \neg L \\ \therefore J \end{array}$$

Is it valid? To find out, we construct a truth table.

J	L	$\neg L \rightarrow (J \vee L)$	$\neg L$	J
T	T	F	T	T
T	F	T	T	T
F	T	F	F	F
F	F	T	F	F

Yes, the argument is valid. The only row on which both the premises are T is the second row, and on that row the conclusion is also T.

3.4 Partial truth tables

In order to show that a sentence is a tautology, we need to show that it is T on every row. So we need a complete truth table. To show that a sentence is *not* a tautology, however, we only need one line: a line on which the sentence is F. Therefore, in order to show that something is not a tautology, it is enough to provide a one-line *partial truth table*— regardless of how many sentence letters the sentence might have in it.

Consider, for example, the sentence $(U \& T) \rightarrow (S \& W)$. We want to show that it is *not* a tautology by providing a partial truth table. We fill in F for the entire sentence. The main connective of the sentence is a conditional. In order for the conditional to be false, the antecedent must be true (T) and the consequent must be false (F). So we fill these in on the table:

S	T	U	W	$(U \& T) \rightarrow (S \& W)$
				T F F

In order for the $(U \& T)$ to be true, both U and T must be true.

S	T	U	W	$(U \& T) \rightarrow (S \& W)$
	T	T		T T T F F

Now we just need to make $(S \& W)$ false. To do this, we need to make at least one of S and W false. We can make both S and W false if we want. All that matters is that the whole sentence turns out false on this line. Making an arbitrary decision, we finish the table in this way:

S	T	U	W	$(U \& T) \rightarrow (S \& W)$
F	T	T	F	T T T F F F F

Showing that something is a contradiction requires a complete truth table. Showing that something is *not* a contradiction requires only a one-line partial truth table, where the sentence is true on that one line.

A sentence is contingent if it is neither a tautology nor a contradiction. So showing that a sentence is contingent requires a *two-line* partial truth table: The sentence must be true on one line and false on the other. For example, we can show that the sentence above is contingent with this truth table:

S	T	U	W	$(U \& T) \rightarrow (S \& W)$
F	T	T	F	T T T F F F F
F	T	F	F	F F T T F F F

Note that there are many combinations of truth values that would have made the sentence true, so there are many ways we could have written the second line.

Showing that a sentence is *not* contingent requires providing a complete truth table, because it requires showing that the sentence is a tautology or that it is a contradiction. If you do not know whether a particular sentence is contingent, then you do not know whether you will need a complete or partial truth table.

	YES	NO
tautology?	complete truth table	one-line partial truth table
contradiction?	complete truth table	one-line partial truth table
contingent?	two-line partial truth table	complete truth table
equivalent?	complete truth table	one-line partial truth table
consistent?	one-line partial truth table	complete truth table
valid?	complete truth table	one-line partial truth table

Table 3.2: Do you need a complete truth table or a partial truth table? It depends on what you are trying to show.

You can always start working on a complete truth table. If you complete rows that show the sentence is contingent, then you can stop. If not, then complete the truth table. Even though two carefully selected rows will show that a contingent sentence is contingent, there is nothing wrong with filling in more rows.

Showing that two sentences are logically equivalent requires providing a complete truth table. Showing that two sentences are *not* logically equivalent requires only a one-line partial truth table: Make the table so that one sentence true and the other false.

Showing that a set of sentences is consistent requires providing one row of a truth table on which all of the sentences are true. The rest of the table is irrelevant, so a one-line partial truth table will do. Showing that a set of sentences is inconsistent, on the other hand, requires a complete truth table: You must show that on every row of the table at least one of the sentences is false.

Showing that an argument is valid requires a complete truth table. Showing that an argument is *invalid* only requires providing a one-line truth table: If you can produce a line on which the premises are all true and the conclusion is false, then the argument is invalid.

Table 3.2 summarizes when a complete truth table is required and when a partial truth table will do.

Practice Exercises

If you want additional practice, you can construct truth tables for any of the sentences and arguments in the exercises for the previous chapter.

★ **Part A** Determine whether each sentence is a tautology, a contradiction, or a contingent sentence. Justify your answer with a complete or partial truth table

where appropriate.

1. $A \rightarrow A$
2. $\neg B \& B$
3. $C \rightarrow \neg C$
4. $\neg D \vee D$
5. $(A \leftrightarrow B) \leftrightarrow \neg(A \leftrightarrow \neg B)$
6. $(A \& B) \vee (B \& A)$
7. $(A \rightarrow B) \vee (B \rightarrow A)$
8. $\neg[A \rightarrow (B \rightarrow A)]$
9. $(A \& B) \rightarrow (B \vee A)$
10. $A \leftrightarrow [A \rightarrow (B \& \neg B)]$
11. $\neg(A \vee B) \leftrightarrow (\neg A \& \neg B)$
12. $\neg(A \& B) \leftrightarrow A$
13. $[(A \& B) \& \neg(A \& B)] \& C$
14. $A \rightarrow (B \vee C)$
15. $[(A \& B) \& C] \rightarrow B$
16. $(A \& \neg A) \rightarrow (B \vee C)$
17. $\neg[(C \vee A) \vee B]$
18. $(\neg B \& D) \leftrightarrow [A \leftrightarrow (A \vee C)]$

★ **Part B** Determine whether each pair of sentences is logically equivalent. Justify your answer with a complete or partial truth table where appropriate.

1. $A, \neg A$
2. $A, (A \vee A)$
3. $A \rightarrow A, A \leftrightarrow A$
4. $(A \vee \neg B), (A \rightarrow B)$
5. $A \& \neg A, \neg B \leftrightarrow B$
6. $\neg(A \& B), (\neg A \vee \neg B)$
7. $\neg(A \rightarrow B), (\neg A \rightarrow \neg B)$
8. $(A \rightarrow B), (\neg B \rightarrow \neg A)$
9. $[(A \vee B) \vee C], [A \vee (B \vee C)]$
10. $[(A \vee B) \& C], [A \vee (B \& C)]$

★ **Part C** Determine whether each set of sentences is consistent or inconsistent. Justify your answer with a complete or partial truth table where appropriate.

1. $A \rightarrow A, \neg A \rightarrow \neg A, A \& A, A \vee A$
2. $A \& B, C \rightarrow \neg B, C$
3. $A \vee B, A \rightarrow C, B \rightarrow C$
4. $A \rightarrow B, B \rightarrow C, A, \neg C$
5. $B \& (C \vee A), A \rightarrow B, \neg(B \vee C)$
6. $A \vee B, B \vee C, C \rightarrow \neg A$

7. $A \leftrightarrow (B \vee C), C \rightarrow \neg A, A \rightarrow \neg B$
8. $A, B, C, \neg D, \neg E, F$

★ **Part D** Determine whether each argument is valid or invalid. Justify your answer with a complete or partial truth table where appropriate.

1. $A \rightarrow A, \therefore A$
2. $A \vee [A \rightarrow (A \leftrightarrow A)], \therefore A$
3. $A \rightarrow (A \& \neg A), \therefore \neg A$
4. $A \leftrightarrow \neg(B \leftrightarrow A), \therefore A$
5. $A \vee (B \rightarrow A), \therefore \neg A \rightarrow \neg B$
6. $A \rightarrow B, B, \therefore A$
7. $A \vee B, B \vee C, \neg A, \therefore B \& C$
8. $A \vee B, B \vee C, \neg B, \therefore A \& C$
9. $(B \& A) \rightarrow C, (C \& A) \rightarrow B, \therefore (C \& B) \rightarrow A$
10. $A \leftrightarrow B, B \leftrightarrow C, \therefore A \leftrightarrow C$

★ **Part E** Answer each of the questions below and justify your answer.

1. Suppose that \mathcal{A} and \mathcal{B} are logically equivalent. What can you say about $\mathcal{A} \leftrightarrow \mathcal{B}$?
2. Suppose that $(\mathcal{A} \& \mathcal{B}) \rightarrow \mathcal{C}$ is contingent. What can you say about the argument “ $\mathcal{A}, \mathcal{B}, \therefore \mathcal{C}$ ”?
3. Suppose that $\{\mathcal{A}, \mathcal{B}, \mathcal{C}\}$ is inconsistent. What can you say about $(\mathcal{A} \& \mathcal{B} \& \mathcal{C})$?
4. Suppose that \mathcal{A} is a contradiction. What can you say about the argument “ $\mathcal{A}, \mathcal{B}, \therefore \mathcal{C}$ ”?
5. Suppose that \mathcal{C} is a tautology. What can you say about the argument “ $\mathcal{A}, \mathcal{B}, \therefore \mathcal{C}$ ”?
6. Suppose that \mathcal{A} and \mathcal{B} are logically equivalent. What can you say about $(\mathcal{A} \vee \mathcal{B})$?
7. Suppose that \mathcal{A} and \mathcal{B} are *not* logically equivalent. What can you say about $(\mathcal{A} \vee \mathcal{B})$?

Part F We could leave out the biconditional (\leftrightarrow) out of the language. If we did that, we could still write ‘ $A \leftrightarrow B$ ’ so as to make sentences easier to read, but that would be shorthand for $(A \rightarrow B) \& (B \rightarrow A)$. The resulting language would be formally equivalent to SL, since $A \leftrightarrow B$ and $(A \rightarrow B) \& (B \rightarrow A)$ are logically equivalent in SL. If we valued formal simplicity over expressive richness, we could replace more of the connectives with notational conventions and still have a language equivalent to SL.

There are a number of equivalent languages with only two connectives. It would be enough to have only negation and the material conditional. Show this by writing sentences that are logically equivalent to each of the following using only parentheses, sentence letters, negation (\neg), and the material conditional (\rightarrow).

- * 1. $A \vee B$
- * 2. $A \& B$
- * 3. $A \leftrightarrow B$

We could have a language that is equivalent to SL with only negation and disjunction as connectives. Show this: Using only parentheses, sentence letters, negation (\neg), and disjunction (\vee), write sentences that are logically equivalent to each of the following.

- 4. $A \& B$
- 5. $A \rightarrow B$
- 6. $A \leftrightarrow B$

The *Sheffer stroke* is a logical connective with the following characteristic truth table:

\mathcal{A}	\mathcal{B}	$\mathcal{A} \mathcal{B}$
T	T	F
T	F	T
F	T	T
F	F	T

- 7. Write a sentence using the connectives of SL that is logically equivalent to $(A|B)$.

Every sentence written using a connective of SL can be rewritten as a logically equivalent sentence using one or more Sheffer strokes. Using only the Sheffer stroke, write sentences that are equivalent to each of the following.

- 8. $\neg A$
- 9. $(A \& B)$
- 10. $(A \vee B)$
- 11. $(A \rightarrow B)$
- 12. $(A \leftrightarrow B)$

Chapter 4

Quantified logic

4.1 From sentences to predicates

Consider the following argument, which is obviously valid in English:

If everyone knows logic, then either noone will be confused or everyone will. Everyone will be confused only if we try to believe a contradiction. This is a logic class, so everyone knows logic.
 \therefore If we don't try to believe a contradiction, then noone will be confused.

In order to symbolize this in SL, we will need a symbolization key.

L: Everyone knows logic.
N: Noone will be confused.
E: Everyone will be confused.
B: We try to believe a contradiction.

Notice that N and E are both about people being confused, but they are two separate sentence letters. We could not replace E with $\neg N$. Why not? $\neg N$ means 'It is not the case that noone will be confused.' This would be the case if even one person were confused, so it is a long way from saying that *everyone* will be confused.

Once we have separate sentence letters for N and E , however, we erase any connection between the two. They are just two atomic sentences which might be true or false independently. In English, it could never be the case that

both noone and everyone was confused. As sentences of SL, however, there is a truth-value assignment for which N and E are both true.

Expressions like ‘noone’, ‘everyone’, and ‘anyone’ are called *quantifiers*. By translating N and E as separate atomic sentences, we leave out the *quantifier structure* of the sentences. Fortunately, the quantifier structure is not what makes this argument valid. As such, we can safely ignore it. To see this, we translate the argument to SL:

$$\begin{array}{l} L \rightarrow (N \vee E) \\ E \rightarrow B \\ L \\ \therefore \neg B \rightarrow N \end{array}$$

This is a valid argument in SL. (You can do a truth table to check this.)

Now consider another argument. This one is also valid in English.

Willard is a logician. All logicians wear funny hats.
 \therefore Willard wears a funny hat.

To symbolize it in SL, we define a symbolization key:

L: Willard is a logician.
A: All logicians wear funny hats.
F: Willard wears a funny hat.

Now we symbolize the argument:

$$\begin{array}{l} L \\ A \\ \therefore F \end{array}$$

This is *invalid* in SL. (Again, you can confirm this with a truth table.) There is something very wrong here, because this is clearly a valid argument in English. The symbolization in SL leaves out all the important structure. Once again, the translation to SL overlooks quantifier structure: The sentence ‘All logicians wear funny hats’ is about both logicians and hat-wearing. By not translating this structure, we lose the connection between Willard’s being a logician and Willard’s wearing a hat.

Some arguments with quantifier structure can be captured in SL, like the first example, even though SL ignores the quantifier structure. Other arguments are

completely botched in SL, like the second example. Notice that the problem is not that we have made a mistake while symbolizing the second argument. These are the best symbolizations we can give for these arguments *in SL*.

Generally, if an argument containing quantifiers comes out *valid in SL*, then the English language argument is valid. If it comes out *invalid in SL*, then we cannot say the English language argument is invalid. The argument might be valid because of quantifier structure which the natural language argument has and which the argument in SL lacks.

Similarly, if a sentence with quantifiers comes out as a *tautology in SL*, then the English sentence is logically true. If it comes out as *contingent in SL*, then this might be because of the structure of the quantifiers that gets removed when we translate into the formal language.

In order to symbolize arguments that rely on quantifier structure, we need to develop a different logical language. We will call this language quantified logic, QL.

4.2 Building blocks of QL

Just as sentences were the basic unit of sentential logic, predicates will be the basic unit of quantified logic. A predicate is an expression like ‘is a dog.’ This is not a sentence on its own. It is neither true nor false. In order to be true or false, we need to specify something: Who or what is it that is a dog?

The details of this will be explained in the rest of the chapter, but here is the basic idea: In QL, we will represent predicates with capital letters. For instance, we might let D stand for ‘_____ is a dog.’ We will use lower-case letters as the names of specific things. For instance, we might let b stand for Bertie. The expression Db will be a sentence in QL. It is a translation of the sentence ‘Bertie is a dog.’

In order to represent quantifier structure, we will also have symbols that represent quantifiers. For instance, ‘ \exists ’ will mean ‘There is some_____.’ So to say that there is a dog, we can write $\exists xDx$; that is: There is some x such that x is a dog.

That will come later. We start by defining singular terms and predicates.

Singular Terms

In English, a SINGULAR TERM is a word or phrase that refers to a *specific* person, place, or thing. The word ‘dog’ is not a singular term, because there are a great many dogs. The phrase ‘Philip’s dog Bertie’ is a singular term, because it refers to a specific little terrier.

A PROPER NAME is a singular term that picks out an individual without describing it. The name ‘Emerson’ is a proper name, and the name alone does not tell you anything about Emerson. Of course, some names are traditionally given to boys and other are traditionally given to girls. If ‘Jack Hathaway’ is used as a singular term, you might guess that it refers to a man. However, the name does not necessarily mean that the person referred to is a man— or even that the creature referred to is a person. Jack might be a giraffe for all you could tell just from the name. There is a great deal of philosophical action surrounding this issue, but the important point here is that a name is a singular term because it picks out a single, specific individual.

Other singular terms more obviously convey information about the thing to which they refer. For instance, you can tell without being told anything further that ‘Philip’s dog Bertie’ is a singular term that refers to a dog. A DEFINITE DESCRIPTION picks out an individual by means of a unique description. In English, definite descriptions are often phrases of the form ‘the such-and-so.’ They refer to *the* specific thing that matches the given description. For example, ‘the tallest member of Monty Python’ and ‘the first emperor of China’ are definite descriptions. A description that does not pick out a specific individual is not a definite description. ‘A member of Monty Python’ and ‘an emperor of China’ are not definite descriptions.

We can use proper names and definite descriptions to pick out the same thing. The proper name ‘Mount Rainier’ names the location picked out by the definite description ‘the highest peak in Washington state.’ The expressions refer to the same place in different ways. You learn nothing from my saying that I am going to Mount Rainier, unless you already know some geography. You could guess that it is a mountain, perhaps, but even this is not a sure thing; for all you know it might be a college, like Mount Holyoke. Yet if I were to say that I was going to the highest peak in Washington state, you would know immediately that I was going to a mountain in Washington state.

In English, the specification of a singular term may depend on context; ‘Willard’ means a specific person and not just someone named Willard; ‘P.D. Magnus’ as a logical singular term means *me* and not the other P.D. Magnus. We live with this kind of ambiguity in English, but it is important to keep in mind that singular terms in QL must refer to just one specific thing.

In QL, we will symbolize singular terms with lower-case letters *a* through *w*.

We can add subscripts if we want to use some letter more than once. So $a, b, c, \dots, w, a_1, f_{32}, j_{390}$, and m_{12} are all terms in QL.

Note that x, y , and z are not terms in QL. These will be *variables*, letters which do not stand for any specific thing. We will need them when we introduce quantifiers.

Predicates

The simplest predicates are properties of individuals. They are things you can say about an object. ‘_____ is a dog’ and ‘_____ is a member of Monty Python’ are both predicates. In translating English sentences, the term will not always come at the beginning of the sentence: ‘A piano fell on _____’ is also a predicate. Predicates like these are called ONE-PLACE or MONADIC, because there is only one blank to fill in. A one-place predicate and a singular term combine to make a sentence.

Other predicates are about the *relation* between two things. For instance, ‘_____ is bigger than _____’, ‘_____ is to the left of _____’, and ‘_____ owes money to _____.’ These are TWO-PLACE or DYADIC predicates, because they need to be filled in with two terms in order to make a sentence.

In general, you can think about predicates as schematic sentences that need to be filled out with some number of terms. Conversely, you can start with sentences and make predicates out of them by removing terms. Consider the sentence, ‘Vinnie borrowed the family car from Nunzio.’ By removing a singular term, we can recognize this sentence as using any of three different monadic predicates:

_____ borrowed the family car from Nunzio.
Vinnie borrowed _____ from Nunzio.
Vinnie borrowed the family car from _____.

By removing two singular terms, we can recognize three different dyadic predicates:

Vinnie borrowed _____ from _____.
_____ borrowed the family car from _____.
_____ borrowed _____ from Nunzio.

By removing all three singular terms, we can recognize one THREE-PLACE or TRIADIC predicate:

_____ borrowed _____ from _____.

If we are translating this sentence into QL, should we translate it with a one-, two-, or three-place predicate? It depends on what we want to be able to say. If the only thing that we will discuss being borrowed is the family car, then the generality of the three-place predicate is unnecessary. If the only borrowing we need to symbolize is different people borrowing the family car from Nunzio, then a one-place predicate will be enough.

In general, we can have predicates with as many places as we need. Predicates with more than one place are called POLYADIC. Predicates with n places, for some number n , are called N-PLACE or N-ADIC.

In QL, we symbolize predicates with capital letters A through Z , with or without subscripts. When we give a symbolization key for predicates, we will not use blanks; instead, we will use variables. By convention, constants are listed at the end of the key. So we might write a key that looks like this:

Ax: x is angry.
Hx: x is happy.
T₁xy: x is as tall or taller than y .
T₂xy: x is as tough or tougher than y .
Bxyz: y is between x and z .
d: Donald
g: Gregor
m: Marybeth

We can symbolize sentences that use any combination of these predicates and terms. For example:

1. Donald is angry.
2. If Donald is angry, then so are Gregor and Marybeth.
3. Marybeth is at least as tall and as tough as Gregor.
4. Donald is shorter than Gregor.
5. Gregor is between Donald and Marybeth.

Sentence 1 is straightforward: Ad . The ' x ' in the key entry ' Ax ' is just a placeholder; we can replace it with other terms when translating.

Sentence 2 can be paraphrased as, 'If Ad , then Ag and Am .' QL has all the truth-functional connectives of SL, so we translate this as $Ad \rightarrow (Ag \& Am)$.

Sentence 3 can be translated as $T_1mg \& T_2mg$.

Sentence 4 might seem as if it requires a new predicate. If we only needed to symbolize this sentence, we could define a predicate like Sxy to mean ' x is shorter than y .' However, this would ignore the logical connection between

‘shorter’ and ‘taller.’ Considered only as symbols of QL, there is no connection between S and T_1 . They might mean anything at all. Instead of introducing a new predicate, we paraphrase sentence 4 using predicates already in our key: ‘It is not the case that Donald is as taller or taller than Gregor.’ We can translate it as $\neg T_1 dg$.

Sentence 5 requires that we pay careful attention to the order of terms in the key. It becomes $Bdgm$.

4.3 Quantifiers

We are now ready to introduce quantifiers. Consider these sentences:

6. Everyone is happy.
7. Everyone is at least as tough as Donald.
8. Someone is angry.

It might be tempting to translate sentence 6 as $Hd \& Hg \& Hm$. Yet this would only say that Donald, Gregor, and Marybeth are happy. We want to say that *everyone* is happy, even if we have not defined a constant to name them. In order to do this, we introduce the ‘ \forall ’ symbol. This is called the UNIVERSAL QUANTIFIER.

A quantifier must always be followed by a variable and a formula that includes that variable. We can translate sentence 6 as $\forall x Hx$. Paraphrased in English, this means ‘For all x , x is happy.’ We call $\forall x$ an *x-quantifier*. The formula that follows the quantifier is called the *scope* of the quantifier. We will give a formal definition of scope later, but intuitively it is the part of the sentence that the quantifier quantifies over. In $\forall x Hx$, the scope of the universal quantifier is Hx .

Sentence 7 can be paraphrased as, ‘For all x , x is at least as tough as Donald.’ This translates as $\forall x T_2 xd$.

In these quantified sentences, the variable x is serving as a kind of placeholder. The expression $\forall x$ means that you can pick anyone and put them in as x . There is no special reason to use x rather than some other variable. The sentence $\forall x Hx$ means exactly the same thing as $\forall y Hy$, $\forall z Hz$, and $\forall x_5 Hx_5$.

To translate sentence 8, we introduce another new symbol: the EXISTENTIAL QUANTIFIER, \exists . Like the universal quantifier, the existential quantifier requires a variable. Sentence 8 can be translated as $\exists x Ax$. This means that there is some x which is angry. More precisely, it means that there is *at least one* angry person. Once again, the variable is a kind of placeholder; we could just as easily have translated sentence 8 as $\exists z Az$.

Consider these further sentences:

9. Noone is angry.
10. There is someone who is not happy.
11. Not everyone is happy.

Sentence 9 can be paraphrased as, ‘It is not the case that someone is angry.’ This can be translated using negation and an existential quantifier: $\neg\exists xAx$. Yet sentence 9 could also be paraphrased as, ‘Everyone is not angry.’ With this in mind, it can be translated using negation and a universal quantifier: $\forall x\neg Ax$. Both of these are acceptable translations, because they are logically equivalent. The critical thing is whether the negation comes before or after the quantifier.

In general, $\forall x\mathcal{A}$ is logically equivalent to $\neg\exists x\neg\mathcal{A}$. This means that any sentence which can be symbolized with a universal quantifier can be symbolized with an existential quantifier, and vice versa. One translation might seem more natural than the other, but there is no logical different in translating with one quantifier rather than the other. For some sentences, it will simply be a matter of taste.

Sentence 10 is most naturally paraphrased as, ‘There is some x such that x is not happy.’ This becomes $\exists x\neg Hx$. Equivalently, we could write $\neg\forall xHx$.

Sentence 11 is most naturally translated as $\neg\forall xHx$. This is logically equivalent to sentence 10 and so could also be translated as $\exists x\neg Hx$.

Although we have two quantifiers in QL, we could have an equivalent formal language with only one quantifier. We could proceed with only the universal quantifier, for instance, and treat the existential quantifier as a notational convention. We use square brackets $[\]$ to make some sentences more readable, but we know that these are really just parentheses $()$. In the same way, we could write ‘ $\exists x$ ’ knowing that this is just shorthand for ‘ $\neg\forall x\neg$.’ There is a choice between making logic formally simple and making it expressively simple. With QL, we opt for expressive simplicity. Both \forall and \exists will be symbols of QL.

Universe of Discourse

Given the symbolization key we have been using, $\forall xHx$ means ‘Everyone is happy.’ Who is included in this *everyone*? When we use sentences like this in English, we usually do not mean everyone now alive on the Earth. We certainly do not mean everyone who was ever alive or who will ever live. We mean something more modest: everyone in the building, everyone in the class, or everyone in the room.

In order to eliminate this ambiguity, we will need to specify a UNIVERSE OF DISCOURSE— abbreviated UD. The UD is the set of things that we are talking

about. So if we want to talk about people in Chicago, we define the UD to be people in the Chicago. We write this at the beginning of the symbolization key, like this:

UD: people in Chicago

The quantifiers *range over* the universe of discourse. Given this UD, $\forall x$ means ‘Everyone in Chicago’ and $\exists x$ means ‘Someone in Chicago.’ Each constant names some member of the UD, so we can only use this UD with the symbolization key above if Donald, Gregor, and Marybeth are all in Chicago. If we want to talk about people in places besides Chicago, then we need to include those people in the UD.

In QL, the UD must be *non-empty*; that is, it must contain at least one member. It is possible to construct formal languages that allow for empty UD’s, but this introduces complications.

Even allowing for a UD with just one member can produce some strange results. Suppose we have this as a symbolization key:

UD: the Eiffel Tower
Px: x is in Paris.

The sentence $\forall xPx$ might be paraphrased in English as ‘Everything is in Paris.’ Yet that would be misleading. It means that everything *in the UD* is in Paris. This UD contains only the Eiffel Tower, so with this symbolization key $\forall xPx$ just means that the Eiffel Tower is in Paris.

Non-referring terms

In QL, each constant must pick out exactly one member of the UD. A constant cannot refer to more than one thing— it is a *singular* term. Each constant must still pick out *something*. This is connected to a classic philosophical problem: the so-called problem of non-referring terms.

Medieval philosophers typically used sentences about the *chimera* to exemplify this problem. Chimera is a mythological creature; it does not really exist. Consider these two sentences:

12. Chimera is angry.
13. Chimera is not angry.

It is tempting just to define a constant to mean ‘chimera.’ The symbolization key would look like this:

UD: creatures on Earth
Ax: x is angry.
c: chimera

We could then translate sentence 12 as Ac and sentence 13 as $\neg Ac$.

Problems will arise when we ask whether these sentences are true or false.

One option is to say that sentence 12 is not true, because there is no chimera. If sentence 12 is false because it talks about a non-existent thing, then sentence 13 is false for the same reason. Yet this would mean that Ac and $\neg Ac$ would both be false. Given the truth conditions for negation, this cannot be the case.

Since we cannot say that they are both false, what should we do? Another option is to say that sentence 12 is *meaningless* because it talks about a non-existent thing. So Ac would be a meaningful expression in QL for some interpretations but not for others. Yet this would make our formal language hostage to particular interpretations. Since we are interested in logical form, we want to consider the logical force of a sentence like Ac apart from any particular interpretation. If Ac were sometimes meaningful and sometimes meaningless, we could not do that.

This is the *problem of non-referring terms*, and we will return to it later (see p. 74.) The important point for now is that each constant of QL *must* refer to something in the UD, although the UD can be any set of things that we like. If we want to symbolize arguments about mythological creatures, then we must define a UD that includes them. This option is important if we want to consider the logic of stories. We can translate a sentence like ‘Sherlock Holmes lived at 221B Baker Street’ by including fictional characters like Sherlock Holmes in our UD.

4.4 Translating to QL

We now have all of the pieces of QL. Translating more complicated sentences will only be a matter of knowing the right way to combine predicates, constants, quantifiers, and connectives. Consider these sentences:

14. Every coin in my pocket is a quarter.
15. Some coin on the table is a dime.
16. Not all the coins on the table are dimes.
17. None of the coins in my pocket are dimes.

In providing a symbolization key, we need to specify a UD. Since we are talking about coins in my pocket and on the table, the UD must at least contain all of

those coins. Since we are not talking about anything besides coins, we let the UD be all coins. Since we are not talking about any specific coins, we do not need to define any constants. So we define this key:

- UD:** all coins
- Px:** x is in my pocket.
- Tx:** x is on the table.
- Qx:** x is a quarter.
- Dx:** x is a dime.

Sentence 14 is most naturally translated with a universal quantifier. The universal quantifier says something about everything in the UD, not just about the coins in my pocket. Sentence 14 means that, for any coin, *if* that coin is in my pocket *then* it is a quarter. So we can translate it as $\forall x(Px \rightarrow Qx)$.

Since sentence 14 is about coins that are both in my pocket *and* that are quarters, it might be tempting to translate it using a conjunction. However, the sentence $\forall x(Px \& Qx)$ would mean that everything in the UD is both in my pocket and a quarter: All the coins that exist are quarters in my pocket. This is would be a crazy thing to say, and it means something very different than sentence 14.

Sentence 15 is most naturally translated with an existential quantifier. It says that there is some coin which is both on the table and which is a dime. So we can translate it as $\exists x(Tx \& Dx)$.

Notice that we needed to use a conditional with the universal quantifier, but we used a conjunction with the existential quantifier. What would it mean to write $\exists x(Tx \rightarrow Dx)$? Probably not what you think. It means that there is some member of the UD which would satisfy the subformula; roughly speaking, there is some a such that $(Ta \rightarrow Da)$ is true. In SL, $\mathcal{A} \rightarrow \mathcal{B}$ is logically equivalent to $\neg\mathcal{A} \vee \mathcal{B}$, and this will also hold in QL. So $\exists x(Tx \rightarrow Dx)$ is true if there is some a such that $(\neg Ta \vee Da)$; i.e., it is true if some coin is *either* not on the table *or* is a dime. Of course there is a coin that is not the table— there are coins lots of other places. So $\exists x(Tx \rightarrow Dx)$ is trivially true. A conditional will usually be the natural connective to use with a universal quantifier, but a conditional within the scope of an existential quantifier can do very strange things. As a general rule, do not put conditionals in the scope of existential quantifiers unless you are sure that you need one.

Sentence 16 can be paraphrased as, ‘It is not the case that every coin on the table is a dime.’ So we can translate it as $\neg\forall x(Tx \rightarrow Dx)$. You might look at sentence 16 and paraphrase it instead as, ‘Some coin on the table is not a dime.’ You would then translate it as $\exists x(Tx \& \neg Dx)$. Although it is probably not obvious, these two translations are logically equivalent. (This is due to

the logical equivalence between $\neg\forall x\mathcal{A}$ and $\exists x\neg\mathcal{A}$, along with the equivalence between $\neg(\mathcal{A} \rightarrow \mathcal{B})$ and $\mathcal{A} \& \neg\mathcal{B}$.)

Sentence 17 can be paraphrased as, ‘It is not the case that there is some dime in my pocket.’ This can be translated as $\neg\exists x(Px \& Dx)$. It might also be paraphrased as, ‘Everything in my pocket is a non-dime,’ and then could be translated as $\forall x(Px \rightarrow \neg Dx)$. Again the two translations are logically equivalent. Both are correct translations of sentence 17.

We can now translate the argument from p. 49, the one that motivated the need for quantifiers:

Willard is a logician. All logicians wear funny hats.
 \therefore Willard wears a funny hat.

UD: people
Lx: x is a logician.
Fx: x wears a funny hat.
w: Willard

Translating, we get:

Lw
 $\forall x(Lx \rightarrow Fx)$
 $\therefore Fw$

This captures the structure that was left out of the SL translation of this argument, and this is a valid argument in QL.

Empty predicates

A predicate need not apply to anything in the UD. A predicate that applies to nothing in the UD is called an **EMPTY** predicate.

Suppose we want to symbolize these two sentences:

18. Every monkey knows sign language.
19. Some monkey knows sign language.

It is possible to write the symbolization key for these sentences in this way:

UD: animals

Mx: x is a monkey.

Sx: x knows sign language.

Sentence 18 can now be translated as $\forall x(Mx \rightarrow Sx)$.

Sentence 19 becomes $\exists x(Mx \& Sx)$.

It is tempting to say that sentence 18 entails sentence 19; that is: if every monkey knows sign language, then it must be that some monkey knows sign language. This is a valid inference in Aristotelean logic: All M s are S , \therefore some M is S . However, the entailment does not hold in QL. It is possible for the sentence $\forall x(Mx \rightarrow Sx)$ to be true even though the sentence $\exists x(Mx \& Sx)$ is false.

How can this be? The answer comes from considering whether these sentences would be true or false *if there were no monkeys*.

We have defined \forall and \exists in such a way that $\forall \mathcal{A}$ is equivalent to $\neg \exists \neg \mathcal{A}$. As such, the universal quantifier doesn't involve the existence of anything— only non-existence. If sentence 18 is true, then there are *no* monkeys who don't know sign language. If there were no monkeys, then $\forall x(Mx \rightarrow Sx)$ would be true and $\exists x(Mx \& Sx)$ would be false.

We allow empty predicates because we want to be able to say things like, 'I do not know if there are any monkeys, but any monkeys that there are know sign language.' That is, we want to be able to have predicates that do not (or might not) refer to anything.

What happens if we add an empty predicate R to the interpretation above? For example, we might define Rx to mean ' x is a refrigerator.' Now the sentence $\forall x(Rx \rightarrow Mx)$ will be true. This is counterintuitive, since we do not want to say that there are a whole bunch of refrigerator monkeys. It is important to remember, though, that $\forall x(Rx \rightarrow Mx)$ means that any member of the UD that is a refrigerator is a monkey. Since the UD is animals, there are no refrigerators in the UD and so the sentence is trivially true.

If you were actually translating the sentence 'All refrigerators are monkeys', then you would want to include appliances in the UD. Then the predicate R would not be empty and the sentence $\forall x(Rx \rightarrow Mx)$ would be false.

A UD must have *at least* one member.
 A predicate may apply to some, all, or no members of the UD.
 A constant must pick out *exactly* one member of the UD

Picking a Universe of Discourse

The appropriate symbolization of an English language sentence in QL will depend on the symbolization key. In some ways, this is obvious: It matters whether Dx means ‘ x is dainty’ or ‘ x is dangerous.’ The meaning of sentences in QL also depends on the UD.

Let Rx mean ‘ x is a rose,’ let Tx mean ‘ x has a thorn,’ and consider this sentence:

20. Every rose has a thorn.

It is tempting to say that sentence 20 should be translated as $\forall x(Rx \rightarrow Tx)$. If the UD contains all roses, that would be correct. Yet if the UD is merely *things on my kitchen table*, then $\forall x(Rx \rightarrow Tx)$ would only mean that every rose on my kitchen table has a thorn. If there are no roses on my kitchen table, the sentence would be trivially true.

The universal quantifier only ranges over members of the UD, so we need to include all roses in the UD in order to translate sentence 20. We have two options. First, we can restrict the UD to include all roses but *only* roses. Then sentence 20 becomes $\forall xTx$. This means that everything in the UD has a thorn; since the UD just is the set of roses, this means that every rose has a thorn. This option can save us trouble if every sentence that we want to translate using the symbolization key is about roses.

Second, we can let the UD contain things besides roses: rhododendrons, rats, rifles, and whatall else. Then sentence 20 must be $\forall x(Rx \rightarrow Tx)$.

If we want the universal quantifier to mean *everything*, without restriction, then we need a UD that contains everything. This can be somewhat ambiguous, because it is unclear whether *everything* includes things that do not actually exist—like fictional characters. Also, it is a matter of philosophical debate whether or not it makes sense to talk about *everything* in an unrestricted way. Generally, if there is any ambiguity, it helps to specify the UD more carefully. For example, if you mean to talk about plants, people, and cities, then the UD might be ‘living things and places.’

Suppose that we want to translate sentence 20 and, with the same symbolization key, translate these sentences:

- 21. Esmerelda has a rose in her hair.
- 22. Everyone is cross with Esmerelda.

We need a UD that includes roses (so that we can symbolize sentence 20) and a

UD that includes people (so we can translate sentence 21–22.) Here is a suitable key:

UD: people and plants
Px: x is a person.
Rx: x is a rose.
Tx: x has a thorn.
Cxy: x is cross with y .
Hxy: x has y in their hair.
e: Esmerelda

Since we do not have a predicate that means ‘... has a rose in her hair’, translating sentence 21 will require paraphrasing. The sentence says that there is a rose in Esmerelda’s hair; that is, there is something which is both a rose and is in Esmerelda’s hair. So we get: $\exists x(Rx \& Hxe)$.

It is tempting to translate sentence 22 as $\forall xCxe$. Unfortunately, this would mean that every member of the UD is cross with Esmerelda—both people and plants. It would mean, for instance, that the rose in Esmerelda’s hair is cross with her. Of course, sentence 22 does not mean that.

‘Everyone’ means every person, not every member of the UD. So we can paraphrase sentence 22 as, ‘Every person is cross with Esmerelda.’ We know how to translate sentences like this: $\forall x(Px \rightarrow Cxe)$

In general, the universal quantifier can be used to mean ‘everyone’ if the UD contains only people. If there are people and other things in the UD, then ‘everyone’ must be treated as ‘every person.’

Translating pronouns

When translating to QL, it is important to understand the structure of the sentences you want to translate. What matters is the final translation in QL, and sometimes you will be able to move from an English language sentence directly to a sentence of QL. Other times, it helps to paraphrase the sentence one or more times. Each successive paraphrase should move from the original sentence closer to something that you can translate directly into QL.

For the next several examples, we will use this symbolization key:

UD: people
Gx: x can play guitar.
Rx: x is a rock star.
l: Lemmy

Now consider these sentences:

- 23. If Lemmy can play guitar, then he is a rock star.
- 24. If a person can play guitar, then he is a rock star.

Sentence 23 and sentence 24 have the same consequent ('... he is a rock star'), but they cannot be translated in the same way. It helps to paraphrase the original sentences, replacing pronouns with explicit references.

Sentence 23 can be paraphrased as, 'If Lemmy can play guitar, then *Lemmy* is a rockstar.' This can obviously be translated as $Gl \rightarrow Rl$.

Sentence 24 must be paraphrased differently: 'If a person can play guitar, then *that person* is a rock star.' This sentence is not about any particular person, so we need a variable. Translating halfway, we can paraphrase the sentence as, 'For any person x , if x can play guitar, then x is a rockstar.' Now this can be translated as $\forall x(Gx \rightarrow Rx)$. This is the same as, 'Everyone who can play guitar is a rock star.'

Consider these further sentences:

- 25. If anyone can play guitar, then Lemmy can.
- 26. If anyone can play guitar, then he or she is a rock star.

These two sentences have the same antecedent ('If anyone can play guitar...'), but they have different logical structures.

Sentence 25 can be paraphrased, 'If someone can play guitar, then Lemmy can play guitar.' The antecedent and consequent are separate sentences, so it can be symbolized with a conditional as the main logical operator: $\exists xGx \rightarrow Gl$.

Sentence 26 can be paraphrased, 'For anyone, if that one can play guitar, then that one is a rock star.' It would be a mistake to symbolize this with an existential quantifier, because it is talking about everybody. The sentence is equivalent to 'All guitar players are rock stars.' It is best translated as $\forall x(Gx \rightarrow Rx)$.

The English words 'any' and 'anyone' should typically be translated using quantifiers. As these two examples show, they sometimes call for an existential quantifier (as in sentence 25) and sometimes for a universal quantifier (as in sentence 26). If you have a hard time determining which is required, paraphrase the sentence with an English language sentence that uses words besides 'any' or 'anyone.'

Quantifiers and scope

In the sentence $\exists xGx \rightarrow Gl$, the scope of the existential quantifier is the expression Gx . Would it matter if the scope of the quantifier were the whole sentence? That is, does the sentence $\exists x(Gx \rightarrow Gl)$ mean something different?

With the key given above, $\exists xGx \rightarrow Gl$ means that if there is some guitarist, then Lemmy is a guitarist. $\exists x(Gx \rightarrow Gl)$ would mean that there is some person such that if that person were a guitarist, then Lemmy would be a guitarist. Recall that the conditional here is a material conditional; the conditional is true if the antecedent is false. Let the constant p denote the author of this book, someone who is certainly not a guitarist. The sentence $Gp \rightarrow Gl$ is true because Gp is false. Since someone (namely p) satisfies the sentence, then $\exists x(Gx \rightarrow Gl)$ is true. The sentence is true because there is a non-guitarist, regardless of Lemmy's skill with the guitar.

Something strange happened when we changed the scope of the quantifier, because the conditional in QL is a material conditional. In order to keep the meaning the same, we would have to change the quantifier: $\exists xGx \rightarrow Gl$ means the same thing as $\forall x(Gx \rightarrow Gl)$, and $\exists x(Gx \rightarrow Gl)$ means the same thing as $\forall xGx \rightarrow Gl$.

This oddity does not arise with other connectives or if the variable is in the consequent of the conditional. For example, $\exists xGx \& Gl$ means the same thing as $\exists x(Gx \& Gl)$, and $Gl \rightarrow \exists xGx$ means the same things as $\exists x(Gl \rightarrow Gx)$.

Ambiguous predicates

Suppose we just want to translate this sentence:

27. Adina is a skilled surgeon.

Let the UD be people, let Kx mean ' x is a skilled surgeon', and let a mean Adina. Sentence 27 is simply Ka .

Suppose instead that we want to translate this argument:

The hospital will only hire a skilled surgeon. All surgeons are greedy.
Billy is a surgeon, but is not skilled. Therefore, Billy is greedy, but
the hospital will not hire him.

We need to distinguish being a *skilled surgeon* from merely being a *surgeon*. So we define this symbolization key:

UD: people
Gx: x is greedy.
Hx: The hospital will hire x .
Rx: x is a surgeon.
Kx: x is skilled.
b: Billy

Now the argument can be translated in this way:

$$\begin{array}{l}
 \forall x [\neg(Rx \ \& \ Kx) \rightarrow \neg Hx] \\
 \forall x (Rx \rightarrow Gx) \\
 Rb \ \& \ \neg Kb \\
 \therefore Gb \ \& \ \neg Hb
 \end{array}$$

Next suppose that we want to translate this argument:

Carol is a skilled surgeon and a tennis player. Therefore, Carol is a surgeon and a skilled tennis player.

If we start with the symbolization key we used for the previous argument, we could add a predicate (let Tx mean ' x is a tennis player') and a constant (let c mean Carol). Then the argument becomes:

$$\begin{array}{l}
 (Rc \ \& \ Kc) \ \& \ Tc \\
 \therefore Tc \ \& \ Kc
 \end{array}$$

This translation is a disaster! It takes what in English is a terrible argument and translates it as a valid argument in QL. The problem is that there is a difference between being *skilled as a surgeon* and *skilled as a tennis player*. Translating this argument correctly requires two separate predicates, one for each type of skill. If we let K_1x mean ' x is skilled as a surgeon' and K_2x mean ' x is skilled as a tennis player,' then we can symbolized the argument in this way:

$$\begin{array}{l}
 (Rc \ \& \ K_1c) \ \& \ Tc \\
 \therefore Tc \ \& \ K_2c
 \end{array}$$

Like the English language argument it translates, this is invalid.

The moral of these examples is that you need to be careful of symbolizing predicates in an ambiguous way. Similar problems can arise with predicates like *good*, *bad*, *big*, and *small*. Just as skilled surgeons and skilled tennis players have different skills, big dogs, big mice, and big problems are big in different ways.

Is it enough to have a predicate that means ‘ x is a skilled surgeon’, rather than two predicates ‘ x is skilled’ and ‘ x is a surgeon’? Sometimes. As sentence 27 shows, sometimes we do not need to distinguish between skilled surgeons and other surgeons.

Must we always distinguish between different ways of being skilled, good, bad, or big? No. As the argument about Billy shows, sometimes we only need to talk about one kind of skill. If you are translating an argument that is just about dogs, it is fine to define a predicate that means ‘ x is big.’ If the UD includes dogs and mice, however, it is probably best to make the predicate mean ‘ x is big for a dog.’

Multiple quantifiers

Consider this following symbolization key and the sentences that follow it:

UD: People and dogs
 Dx: x is a dog.
 Fxy: x is a friend of y .
 Oxy: x owns y .
 f: Fifi
 g: Gerald

28. Fifi is a dog.
29. Gerald is a dog owner.
30. Someone is a dog owner.
31. All of Gerald’s friends are dog owners.
32. Every dog owner is the friend of a dog owner.

Sentence 28 is easy: Df .

Sentence 29 can be paraphrased as, ‘There is a dog that Gerald owns.’ This can be translated as $\exists x(Dx \& Ogx)$.

Sentence 30 can be paraphrased as, ‘There is some y such that y is a dog owner.’ The subsentence ‘ y is a dog owner’ is just like sentence 29, except that it is about y rather than being about Gerald. So we can translate sentence 30 as $\exists y \exists x(Dx \& Oyx)$.

Sentence 31 can be paraphrased as, ‘Every friend of Gerald is a dog owner.’ Translating part of this sentence, we get $\forall x(Fxg \rightarrow ‘x$ is a dog owner’). Again, it is important to recognize that ‘ x is a dog owner’ is structurally just like sentence 29. Since we already have an x-quantifier, we will need a different

variable for the existential quantifier. Any other variable will do. Using z , sentence 31 can be translated as $\forall x[Fxg \rightarrow \exists z(Dz \& Oxz)]$.

Sentence 32 can be paraphrased as ‘For any x that is a dog owner, there is a dog owner who is x ’s friend.’ Partially translated, this becomes

$$\forall x[x \text{ is a dog owner} \rightarrow \exists y(y \text{ is a dog owner} \& Fxy)].$$

Completing the translation, sentence 32 becomes

$$\forall x[\exists z(Dz \& Oxz) \rightarrow \exists y(\exists z(Dz \& Oyz) \& Fxy)].$$

Consider this symbolization key and these sentences:

UD: people
Lxy: x likes y .
i: Imre.
k: Karl.

33. Imre likes everyone that Karl likes.

34. There is someone who likes everyone who likes everyone that he likes.

Sentence 33 can be partially translated as $\forall x(\text{Karl likes } x \rightarrow \text{Imre likes } x)$. This becomes $\forall x(Lkx \rightarrow Lix)$.

Sentence 34 is almost a tongue-twister. There is little hope of writing down the whole translation immediately, but we can proceed by small steps. An initial, partial translation might look like this:

$$\exists x \text{ everyone who likes everyone that } x \text{ likes is liked by } x$$

The part that remains in English is a universal sentence, so we translate further:

$$\exists x \forall y (y \text{ likes everyone that } x \text{ likes} \rightarrow x \text{ likes } y).$$

The antecedent of the conditional is structurally just like sentence 33, with y and x in place of Imre and Karl. So sentence 34 can be completely translated in this way

$$\exists x \forall y [\forall z (Lxz \rightarrow Lyz) \rightarrow Lxy]$$

When symbolizing sentences with multiple quantifiers, it is best to proceed by small steps. Paraphrase the English sentence so that the logical structure is readily symbolized in QL. Then translate piecemeal, replacing the daunting task of translating a long sentence with the simpler task of translating shorter formulae.

4.5 Sentences of QL

In this section, we provide a formal definition for a *well-formed formula* (wff) and *sentence* of QL.

Expressions

There are six kinds of symbols in QL:

predicates with subscripts, as needed	A, B, C, \dots, Z $A_1, B_1, Z_1, A_2, A_{25}, J_{375}, \dots$
constants with subscripts, as needed	a, b, c, \dots, w $a_1, w_4, h_7, m_{32}, \dots$
variables with subscripts, as needed	x, y, z $x_1, y_1, z_1, x_2, \dots$
connectives	$\neg, \&, \vee, \rightarrow, \leftrightarrow$
parentheses	$(,)$
quantifiers	\forall, \exists

We define an EXPRESSION OF QL as any string of symbols of SL. Take any of the symbols of QL and write them down, in any order, and you have an expression.

Well-formed formulae

By definition, a TERM OF QL is either a constant or a variable.

An ATOMIC FORMULA OF QL is an n -place predicate followed by n terms.

Just as we did for SL, we will give a *recursive* definition for a wff of QL. In fact, most of the definition will look like the definition of for a wff of SL: Every atomic formula is a wff, and you can build new wffs by applying the sentential connectives.

We could just add a rule for each of the quantifiers and be done with it. For instance: If \mathcal{A} is a wff, then $\forall x\mathcal{A}$ and $\exists x\mathcal{A}$ are wffs. However, this would allow for bizarre sentences like $\forall x\exists xDx$ and $\forall xDw$. What could these possibly mean? We could adopt some interpretation of such sentences, but instead we will write the definition of a wff so that such abominations do not even count as well-formed.

In order for $\forall x\mathcal{A}$ to be a wff, \mathcal{A} must contain the variable x and must not already contain an x -quantifier. $\forall xDw$ will not count as a wff because ‘ x ’ does

not occur in Dw , and $\forall x\exists xDx$ will not count as a wff because $\exists xDx$ contains an x-quantifier

1. Every atomic formula is a wff.
2. If \mathcal{A} is a wff, then $\neg\mathcal{A}$ is a wff of SL.
3. If \mathcal{A} and \mathcal{B} are wffs, then $(\mathcal{A} \& \mathcal{B})$, is a wff.
4. If \mathcal{A} and \mathcal{B} are wffs, $(\mathcal{A} \vee \mathcal{B})$ is a wff.
5. If \mathcal{A} and \mathcal{B} are wffs, then $(\mathcal{A} \rightarrow \mathcal{B})$ is a wff.
6. If \mathcal{A} and \mathcal{B} are wffs, then $(\mathcal{A} \leftrightarrow \mathcal{B})$ is a wff.
7. If \mathcal{A} is a wff, χ is a variable, \mathcal{A} contains at least one occurrence of χ , and \mathcal{A} contains no χ -quantifiers, then $\forall\chi\mathcal{A}$ is a wff.
8. If \mathcal{A} is a wff, χ is a variable, \mathcal{A} contains at least one occurrence of χ , and \mathcal{A} contains no χ -quantifiers, then $\exists\chi\mathcal{A}$ is a wff.
9. All and only wffs of QL can be generated by applications of these rules.

Notice that the ‘ χ ’ that appears in the definition above is not the variable x . It is a *meta-variable* that stands in for any variable of QL. So $\forall xAx$ is a wff, but so are $\forall yAy$, $\forall zAz$, $\forall x_4Ax_4$, and $\forall z_9Az_9$.

We can now give a formal definition for scope: The SCOPE of a quantifier is the subformula for which the quantifier is the main logical operator.

Sentences

A sentence is something that can be either true or false. In SL, every wff was a sentence. This will not be the case in QL. Consider the following symbolization key:

UD: people
Lxy: x loves y
b: Boris

Consider the expression Lzz . It is an atomic formula: a two-place predicate followed by two terms. All atomic formula are wffs, so Lzz is a wff. Does it mean anything? You might think that it means that z loves himself, in the same way that Lbb means that Boris loves himself. Yet z is a variable; it does not name some person the way a constant would. The wff Lzz does not tell

us how to interpret z . Does it mean everyone? anyone? someone? If we had a z -quantifier, it would tell us how to interpret z . For instance, $\exists zLzz$ would mean that someone loves themselves.

Some formal languages treat a wff like Lzz as implicitly having a universal quantifier in front. We will not do this for QL. If you mean to say that everyone loves themselves, then you need to write the quantifier: $\forall zLzz$

In order to make sense of a variable, we need a quantifier to tell us how to interpret that variable. The scope of an x -quantifier, for instance, is the part of the formula where the quantifier tells how to interpret x .

In order to be precise about this, we define a BOUND VARIABLE to be an occurrence of a variable χ that is within the scope of an χ -quantifier. A FREE VARIABLE is an occurrence of a variable that is not bound.

For example, consider the wff $\forall x(Ex \vee Dy) \rightarrow \exists z(Ex \rightarrow Lzx)$. The scope of the universal quantifier $\forall x$ is $(Ex \vee Dy)$, so the first x is bound by the universal quantifier but the second and third x s are free. There is no y -quantifier, so the y is free. The scope of the existential quantifier $\exists z$ is $(Ex \rightarrow Lzx)$, so both occurrences of z are bound by it.

We define a SENTENCE of QL as a wff of QL that contains no free variables.

Notational conventions

We will adopt the same notational conventions that we did for SL (p. 31.) First, we may leave off the outermost parentheses of a formula. Second, we will use square brackets '[' and ']' in place of parentheses to increase the readability of formulae. Third, we will leave out parentheses between each pair of conjuncts when writing long series of conjunctions. Fourth, we will leave out parentheses between each pair of disjuncts when writing long series of disjunctions.

Substitution instance

If \mathcal{A} is a wff, c a constant, and x a variable, then $\mathcal{A}[c/x]$ is the wff made by replacing each occurrence of x in \mathcal{A} with c . This is called a SUBSTITUTION INSTANCE of $\forall x\mathcal{A}$ and $\exists x\mathcal{A}$; c is called the INSTANTIATING CONSTANT.

For example: $Aa \rightarrow Ba$, $Af \rightarrow Bf$, and $Ak \rightarrow Bk$ are all substitution instances of $\forall x(Ax \rightarrow Bx)$; the instantiating constants are a , f , and k , respectively.

Raj , Rdj , and Rjj are substitution instances of $\exists zRzj$; the instantiating constants are a , d , and j , respectively.

4.6 Identity

Consider this sentence:

35. Pavel owes money to everyone else.

Let the UD be people; this will allow us to translate ‘everyone’ as a universal quantifier. Let Oxy mean ‘ x owes money to y ’, and let p mean Pavel. Now we can symbolize sentence 35 as $\forall xOpx$. Unfortunately, this translation has some odd consequences. It says that Pavel owes money to every member of the UD, including Pavel; it entails that Pavel owes money to himself. However, sentence 35 does not say that Pavel owes money to himself; he owes money to everyone *else*. This is a problem, because $\forall xOpx$ is the best translation we can give of this sentence into QL.

The solution is to add another symbol to QL. The symbol ‘=’ is a two-place predicate. Since it has a special logical meaning, we write it a bit differently: For two terms t_1 and t_2 , $t_1 = t_2$ is an atomic formula.

The formula $x = y$ means ‘ x is identical to y .’ This does not mean merely that x and y are indistinguishable or that all of the same predicates are true of them. Rather, it means that x and y refer to a single thing.

When we write $x \neq y$, we mean that x and y are not identical. There is no reason to introduce this as an additional predicate. Instead, $x \neq y$ is an abbreviation of $\neg(x = y)$.

Now suppose we want to symbolize this sentence:

36. Pavel is Mister Checkov.

Let the constant c mean Mister Checkov. Sentence 36 can be symbolized as $p = c$. This means that the constants p and c both refer to the same guy.

This is all well and good, but how does it help with sentence 35? That sentence can be paraphrased as, ‘Everyone who is not Pavel is owed money by Pavel.’ This is a sentence structure we already know how to symbolize: ‘For all x , if x is not Pavel, then x is owed money by Pavel.’ In QL with identity, this becomes $\forall x(x \neq p \rightarrow Opx)$.

In addition to sentences that use the word ‘else’, identity will be helpful when symbolizing some sentences that contain the words ‘besides’ and ‘only.’ Consider these examples:

37. No one besides Pavel owes money to Hikaru.

38. Only Pavel owes Hikaru money.

We add the constant h , which means Hikaru.

Sentence 37 can be paraphrased as, ‘No one who is not Pavel owes money to Hikaru.’ This can be translated as $\neg\exists x(x \neq p \ \& \ O x h)$.

Sentence 38 can be paraphrased as, ‘Pavel owes Hikaru *and* no one besides Pavel owes Hikaru money.’ We have already translated one of the conjuncts, and the other is straightforward. Sentence 38 becomes $O p h \ \& \ \neg\exists x(x \neq p \ \& \ O x h)$.

Expressions of quantity

We can also use identity to say how many things there are of a particular kind. For example, consider these sentences:

- 39. There is at least one apple on the table.
- 40. There are at least two apples on the table.
- 41. There are at least three apples on the table.

Let the UD be *things on the table*, and let Ax mean ‘ x is an apple.’

Sentence 39 does not require identity. It can be translated adequately as $\exists x Ax$: There is some apple on the table—perhaps many, but at least one.

It might be tempting to also translate sentence 40 without identity. Yet consider the sentence $\exists x \exists y (Ax \ \& \ Ay)$. It means that there is some apple x in the UD and some apple y in the UD. Since nothing precludes x and y from picking out the same member of the UD, this would be true even if there were only one apple. In order to make sure that there are two *different* apples, we need an identity predicate. Sentence 40 needs to say that the two apples that exist are not identical, so it can be translated as $\exists x \exists y (Ax \ \& \ Ay \ \& \ x \neq y)$.

Sentence 41 requires talking about three different apples. It can be translated as $\exists x \exists y \exists z (Ax \ \& \ Ay \ \& \ Az \ \& \ x \neq y \ \& \ y \neq z \ \& \ x \neq z)$.

Continuing in this way, we could translate ‘There are at least n apples on the table.’ There is a summary of how to symbolize sentences like these on p. 150.

Now consider these sentences:

- 42. There is at most one apple on the table.
- 43. There are at most two apples on the table.

Sentence 42 can be paraphrased as, ‘It is not the case that there are at least *two* apples on the table.’ This is just the negation of sentence 40:

$$\neg \exists x \exists y (Ax \& Ay \& x \neq y)$$

Sentence 42 can also be approached in another way. It means that any apples that there are on the table must be the selfsame apple, so it can be translated as $\forall x \forall y [(Ax \& Ay) \rightarrow x = y]$. The two translations are logically equivalent, so both are correct.

In a similar way, sentence 43 can be translated in two equivalent ways. It can be paraphrased as, ‘It is not the case that there are *three* or more distinct apples’, so it can be translated as the negation of sentence 41. Using universal quantifiers, it can also be translated as

$$\forall x \forall y \forall z [(Ax \& Ay \& Az) \rightarrow (x = y \vee x = z \vee y = z)].$$

See p. 150 for the general case.

The examples above are sentences about apples, but the logical structure of the sentences translates mathematical inequalities like $a \geq 3$, $a \leq 2$, and so on. We also want to be able to translate statements of equality which say exactly how many things there are. For example:

- 44. There is exactly one apple on the table.
- 45. There are exactly two apples on the table.

Sentence 44 can be paraphrased as, ‘There is *at least* one apple on the table, and there is *at most* one apple on the table.’ This is just the conjunction of sentence 39 and sentence 42: $\exists x Ax \& \forall x \forall y [(Ax \& Ay) \rightarrow x = y]$. This is a somewhat complicated way of going about it. It is perhaps more straightforward to paraphrase sentence 44 as, ‘There is a thing which is the only apple on the table.’ Thought of in this way, the sentence can be translated $\exists x [Ax \& \neg \exists y (Ay \& x \neq y)]$.

Similarly, sentence 45 may be paraphrased as, ‘There are two different apples on the table, and these are the only apples on the table.’ This can be translated as $\exists x \exists y [Ax \& Ay \& x \neq y \& \neg \exists z (Az \& x \neq z \& y \neq z)]$.

Finally, consider this sentence:

- 46. There are at most two things on the table.

It might be tempting to add a predicate so that Tx would mean ‘ x is a thing on the table.’ However, this is unnecessary. Since the UD is the set of things

on the table, all members of the UD are on the table. If we want to talk about a *thing on the table*, we need only use a quantifier. Sentence 46 can be symbolized like sentence 43 (which said that there were at most two apples), but leaving out the predicate entirely. That is, sentence 46 can be translated as $\forall x\forall y\forall z(x = y \vee x = z \vee y = z)$.

Techniques for symbolizing expressions of quantity ('at most', 'at least', and 'exactly') are summarized on p. 150.

Definite descriptions

Recall that a constant of QL must refer to some member of the UD. This constraint allows us to avoid the problem of non-referring terms. Given a UD that included only actually existing creatures but a constant c that meant 'chimera' (a mythical creature), sentences containing c would become impossible to evaluate.

The most widely influential solution to this problem was introduced by Bertrand Russell in 1905. Russell asked how we should understand this sentence:

47. The present king of France is bald.

The phrase 'the present king of France' is supposed to pick out an individual by means of a definite description. However, there was no king of France in 1905 and there is none now. Since the description is a non-referring term, we cannot just define a constant to mean 'the present king of France' and translate the sentence as Kf .

Russell's idea was that sentences that contain definite descriptions have a different logical structure than sentences that contain proper names, even though they share the same grammatical form. What do we mean when we use an unproblematic, referring description, like 'the highest peak in Washington state'? We mean that there is such a peak, because we could not talk about it otherwise. We also mean that it is the only such peak. If there was another peak in Washington state of exactly the same height as Mount Rainier, then Mount Rainier would not be *the* highest peak.

According to this analysis, sentence 47 is saying three things. First, it makes an *existence* claim: There is some present king of France. Second, it makes a *uniqueness* claim: This guy is the only present king of France. Third, it makes a claim of *predication*: This guy is bald.

In order to symbolize definite descriptions in this way, we need the identity predicate. Without it, we could not translate the uniqueness claim which (according to Russell) is implicit in the definite description.

Let the UD be *people actually living*, let Kx mean ' x is the present king of France', and let Bx mean ' x is bald.' Sentence 47 can then be translated as $\exists x[Fx \& \neg \exists y(Fy \& x \neq y) \& Bx]$. This says that there is some guy who is the present king of France, he is the only present king of France, and he is bald.

Understood in this way, sentence 47 is meaningful but false. It says that this guy exists, but he does not.

The problem of non-referring terms is most vexing when we try to translate negations. So consider this sentence:

48. The present king of France is not bald.

According to Russell, this sentence is ambiguous in English. It could mean either of two things:

48a. It is not the case that the present king of France is bald.

48b. The present king of France is non-bald.

Both possible meanings negate sentence 47, but they put the negation in different places.

Sentence 48a is called a WIDE-SCOPE NEGATION, because it negates the entire sentence. It can be translated as $\neg \exists x[Fx \& \neg \exists y(Fy \& x \neq y) \& Bx]$. This does not say anything about the present king of France, but rather says that some sentence about the present king of France is false. Since sentence 47 is false, sentence 48a is true.

Sentence 48b says something about the present king of France. It says that he lacks the property of baldness. Like sentence 47, it makes an existence claim and a uniqueness claim; it just denies the claim of predication. This is called NARROW-SCOPE NEGATION. It can be translated as $\exists x[Fx \& \neg \exists y(Fy \& x \neq y) \& \neg Bx]$. Since there is no present king of France, this sentence is false.

Russell's theory of definite descriptions resolves the problem of non-referring terms and also explains why it seemed so paradoxical. Before we distinguished between the wide-scope and narrow-scope negations, it seemed that sentences like 48 should be both true and false. By showing that such sentences are ambiguous, Russell showed that they are true understood one way but false understood another way.

For a more detailed discussion of Russell's theory of definite descriptions, including objections to it, see the Peter Ludlow's entry 'descriptions' in *The Stanford Encyclopedia of Philosophy*: Summer 2005 edition, edited by Edward N. Zalta, <http://plato.stanford.edu/archives/sum2005/entries/descriptions/>

Practice Exercises

Part A Identify which variables are bound and which are free.

1. $\exists xLxy \ \& \ \forall yLyx$
2. $\forall xAx \ \& \ Bx$
3. $\forall x(Ax \ \& \ Bx) \ \& \ \forall y(Cx \ \& \ Dy)$
4. $\forall x\exists y[Rxy \rightarrow (Jz \ \& \ Kx)] \vee Ryx$
5. $\forall x_1(Mx_2 \leftrightarrow Lx_2x_1) \ \& \ \exists x_2Lx_3x_2$

*** Part B** Using the symbolization key given, translate each English-language sentence into QL.

- UD:** all animals
Ax: x is an alligator.
Mx: x is a monkey.
Rx: x is a reptile.
Zx: x lives at the zoo.
Lxy: x loves y .
a: Amos
b: Bouncer
c: Cleo

1. Amos, Bouncer, and Cleo all live at the zoo.
2. Bouncer is a reptile, but not an alligator.
3. If Cleo loves Bouncer, then Bouncer is a monkey.
4. If both Bouncer and Cleo are alligators, then Amos loves them both.
5. Some reptile lives at the zoo.
6. Every alligator is a reptile.
7. Any animal that lives at the zoo is either a monkey or an alligator.
8. There are reptiles which are not alligators.
9. Cleo loves a reptile.
10. Bouncer loves all the monkeys that live at the zoo.
11. All the monkeys that Amos loves love him back.
12. If any animal is a reptile, then Amos is.
13. If any animal is an alligator, then it is a reptile.
14. Every monkey that Cleo loves is also loved by Amos.
15. There is a monkey that loves Bouncer, but sadly Bouncer does not reciprocate this love.

Part C These are syllogistic figures identified by Aristotle and his successors, along with their medieval names. Translate each argument into QL.

- Barbara** All B s are C s. All A s are B s. \therefore All A s are C s.
- Baroco** All C s are B s. Some A is not B . \therefore Some A is not C .
- Bocardo** Some B is not C . All A s are B s. \therefore Some A is not C .
- Celantes** No B s are C s. All A s are B s. \therefore No C s are A s.
- Celarent** No B s are C s. All A s are B s. \therefore No A s are C s.
- Cemestres** No C s are B s. No A s are B s. \therefore No A s are C s.
- Cesare** No C s are B s. All A s are B s. \therefore No A s are C s.
- Dabitis** All B s are C s. Some A is B . \therefore Some C is A .
- Darii** All B s are C s. Some A is B . \therefore Some A is C .
- Datisi** All B s are C s. Some A is B . \therefore Some A is C .
- Disamis** Some B is C . All A s are B s. \therefore Some A is C .
- Ferison** No B s are C s. Some A is B . \therefore Some A is not C .
- Ferio** No B s are C s. Some A is B . \therefore Some A is not C .
- Festino** No C s are B s. Some A is B . \therefore Some A is not C .
- Baralippton** All B s are C s. All A s are B s. \therefore Some C is A .
- Frisosomorum** Some B is C . No A s are B s. \therefore Some C is not A .

Part D Using the symbolization key given, translate each English-language sentence into QL.

- UD:** all animals
Dx: x is a dog.
Sx: x likes samurai movies.
Lxy: x is larger than y .
b: Bertie
e: Emerson
f: Fergis

1. Bertie is a dog who likes samurai movies.
2. Bertie, Emerson, and Fergis are all dogs.
3. Emerson is larger than Bertie, and Fergis is larger than Emerson.
4. All dogs like samurai movies.
5. Only dogs like samurai movies.
6. There is a dog that is larger than Emerson.
7. If there is a dog larger than Fergis, then there is a dog larger than Emerson.

8. No animal that likes samurai movies is larger than Emerson.
9. Fergis is the largest dog.
10. Any animal that dislikes samurai movies is larger than Bertie.
11. There is an animal that is between Bertie and Emerson in size.
12. There is no dog that is between Bertie and Emerson in size.
13. No dog is larger than itself.
14. Every dog is larger than some dog.
15. There is an animal that is smaller than every dog.
16. If there is an animal that is larger than any dog, then that animal does not like samurai movies.

Part E For each argument, write a symbolization key and translate the argument into QL.

1. Nothing on my desk escapes my attention. There is a computer on my desk. As such, there is a computer that does not escape my attention.
2. All my dreams are black and white. Old TV shows are in black and white. Therefore, some of my dreams are old TV shows.
3. Neither Holmes nor Watson has been to Australia. A person could see a kangaroo only if they had been to Australia or to a zoo. Although Watson has not seen a kangaroo, Holmes has. Therefore, Holmes has been to a zoo.
4. No one expects the Spanish Inquisition. No one knows the troubles I've seen. Therefore, anyone who expects the Spanish Inquisition knows the troubles I've seen.
5. An antelope is bigger than a bread box. I am thinking of something that is no bigger than a bread box, and it is either an antelope or a cantaloupe. As such, I am thinking of a cantaloupe.
6. All babies are illogical. Nobody who is illogical can manage a crocodile. Berthold is a baby. Therefore, Berthold is unable to manage a crocodile.

★ **Part F** Using the symbolization key given, translate each English-language sentence into QL.

UD: candies

Cx: x has chocolate in it.

Mx: x has marzipan in it.

Sx: x has sugar in it.

Tx: Boris has tried x .

Bxy: x is better than y .

1. Boris has never tried any candy.
2. Marzipan is always made with sugar.
3. Some candy is sugar-free.

4. The very best candy is chocolate.
5. No candy is better than itself.
6. Boris has never tried sugar-free chocolate.
7. Boris has tried marzipan and chocolate, but never together.
8. Any candy with chocolate is better than any candy without it.
9. Any candy with chocolate and marzipan is better than any candy that lacks both.

Part G Using the symbolization key given, translate each English-language sentence into QL.

UD: people and dishes at a potluck

Rx: x has run out.

Tx: x is on the table.

Fx: x is food.

Px: x is a person.

Lxy: x likes y .

e: Eli

f: Francesca

g: the guacamole

1. All the food is on the table.
2. If the guacamole has not run out, then it is on the table.
3. Everyone likes the guacamole.
4. If anyone likes the guacamole, then Eli does.
5. Francesca only likes the dishes that have run out.
6. Francesca likes no one, and no one likes Francesca.
7. Eli likes anyone who likes the guacamole.
8. Eli likes anyone who likes the people that he likes.
9. If there is a person on the table already, then all of the food must have run out.

* **Part H** Using the symbolization key given, translate each English-language sentence into QL.

UD: people

Dx: x dances ballet.

Fx: x is female.

Mx: x is male.

Cxy: x is a child of y .

Sxy: x is a sibling of y .

e: Elmer

j: Jane

p: Patrick

1. All of Patrick's children are ballet dancers.
2. Jane is Patrick's daughter.
3. Patrick has a daughter.
4. Jane is an only child.
5. All of Patrick's daughters dance ballet.
6. Patrick has no sons.
7. Jane is Elmer's niece.
8. Patrick is Elmer's brother.
9. Patrick's brothers have no children.
10. Jane is an aunt.
11. Everyone who dances ballet has a sister who also dances ballet.
12. Every man who dances ballet is the child of someone who dances ballet.

Part I Using the symbolization key given, translate each English-language sentence into QL with identity. The last sentence is ambiguous and can be translated two ways; you should provide both translations. (Hint: Identity is only required for the last four sentences.)

UD: people
Kx: x knows the combination to the safe.
Sx: x is a spy.
Vx: x is a vegetarian.
Txy: x trusts y .
h: Hofthor
i: Ingmar

1. Hofthor is a spy, but no vegetarian is a spy.
2. No one knows the combination to the safe unless Ingmar does.
3. No spy knows the combination to the safe.
4. Neither Hofthor nor Ingmar is a vegetarian.
5. Hofthor trusts a vegetarian.
6. Everyone who trusts Ingmar trusts a vegetarian.
7. Everyone who trusts Ingmar trusts someone who trusts a vegetarian.
8. Only Ingmar knows the combination to the safe.
9. Ingmar trusts Hofthor, but no one else.
10. The person who knows the combination to the safe is a vegetarian.
11. The person who knows the combination to the safe is not a spy.

★ **Part J** Using the symbolization key given, translate each English-language sentence into QL with identity. The last two sentences are ambiguous and can be translated two ways; you should provide both translations for each.

UD: cards in a standard deck
Bx: x is black.

Cx: x is a club.
Dx: x is a deuce.
Jx: x is a jack.
Mx: x is a man with an axe.
Ox: x is one-eyed.
Wx: x is wild.

1. All clubs are black cards.
2. There are no wild cards.
3. There are at least two clubs.
4. There is more than one one-eyed jack.
5. There are at most two one-eyed jacks.
6. There are two black jacks.
7. There are four deuces.
8. The deuce of clubs is a black card.
9. One-eyed jacks and the man with the axe are wild.
10. If the deuce of clubs is wild, then there is exactly one wild card.
11. The man with the axe is not a jack.
12. The deuce of clubs is not the man with the axe.

Part K Using the symbolization key given, translate each English-language sentence into QL with identity. The last two sentences are ambiguous and can be translated two ways; you should provide both translations for each.

UD: animals in the world
Bx: x is in Farmer Brown's field.
Hx: x is a horse.
Px: x is a Pegasus.
Wx: x has wings.

1. There are at least three horses in the world.
2. There are at least three animals in the world.
3. There is more than one horse in Farmer Brown's field.
4. There are three horses in Farmer Brown's field.
5. There is a single winged creature in Farmer Brown's field.
6. The Pegasus is a winged horse.
7. The animal in Farmer Brown's field is not a horse.
8. The horse in Farmer Brown's field does not have wings.

Chapter 5

Formal semantics

In this chapter, we describe a *formal semantics* for SL and for QL. The word ‘semantics’ comes from the greek word for ‘mark’ and means ‘related to meaning.’ So a formal semantics will be a mathematical account of meaning in the formal language.

A formal, logical language is built from two kinds of elements: logical symbols and non-logical symbols. Connectives (like ‘&’) and quantifiers (like ‘ \forall ’) are logical symbols, because their meaning is specified within the formal language. When writing a symbolization key, you are not allowed to change the meaning of the logical symbols. You cannot say, for instance, that the ‘ \neg ’ symbol will mean ‘not’ in one argument and ‘perhaps’ in another. The ‘ \neg ’ symbol always means logical negation. It is used to translate the English language word ‘not’, but it is a symbol of a formal language and is defined by its truth conditions.

The sentences letters in SL are non-logical symbols, because their meaning is not defined by the logical structure of SL. When we translate an argument from English to SL, for example, the sentence letter M does not have its meaning fixed in advance; instead, we provide a symbolization key that says how M should be interpreted in that argument. In QL, the predicates and constants are non-logical symbols.

In translating from English to a formal language, we provided symbolization keys which were interpretations of all the non-logical symbols we used in the translation. An INTERPRETATION gives a meaning to all the non-logical elements of the language.

It is possible to provide different interpretations that make no formal difference. In SL, for example, we might say that D means ‘Today is Tuesday’; we might say instead that D means ‘Today is the day after Monday.’ These are two different

interpretations, because they use different English sentences for the meaning of D . Yet, formally, there is no difference between them. All that matters once we have symbolized these sentences is whether they are true or false. In order to characterize what makes a difference in the formal language, we need to know what makes sentences true or false. For this, we need a formal characterization of *truth*.

When we gave definitions for a sentence of SL and for a sentence of QL, we distinguished between the OBJECT LANGUAGE and the METALANGUAGE. The object language is the language that we are *talking about*: either SL or QL. The metalanguage is the language that we use to talk about the object language: English, supplemented with some mathematical jargon. It will be important to keep this distinction in mind.

5.1 Semantics for SL

This section provides a rigorous, formal characterization of *truth in SL* which builds on what we already know from doing truth tables. We were able to use truth tables to reliably test whether a sentence was a tautology in SL, whether two sentences were equivalent, whether an argument was valid, and so on. For instance: \mathcal{A} is a tautology in SL if it is T on every line of complete truth table.

This worked because each line of a truth table corresponds to a way the world might be. We considered all the possible combinations of T and F for the sentence letters that made a difference to the sentences we cared about. The truth table allowed us to determine what would happen given these different combinations.

Once we construct a truth table, the symbols ‘T’ and ‘F’ are divorced from their metalinguistic meaning of ‘true’ and ‘false’. We *interpret* ‘T’ as meaning ‘true’, but the formal properties of T are defined by the characteristic truth tables for the various connectives. The tables would be the same if we had used the symbols ‘1’ and ‘0’, and computers can be programmed to fill out truth tables without having any sense that 1 means true and 0 means false.

Formally, what we want is a function that assigns a 1 or 0 to each of the sentences of SL. We can interpret this function as a definition of truth for SL if it assigns 1 to all of the true sentences of SL and 0 to all of the false sentences of SL. Call this function ‘ v ’ (for ‘valuation’). We want v to be a function such that for any sentence \mathcal{A} , $v(\mathcal{A}) = 1$ if \mathcal{A} is true and $v(\mathcal{A}) = 0$ if \mathcal{A} is false.

Recall that the recursive definition of a wff for SL had two stages: The first step said that atomic sentences (solitary sentence letters) are wffs. The second stage allowed for wffs to be constructed out of more basic wffs. There were clauses of

the definition for all of the sentential connectives. For example, if \mathcal{A} is a wff, then $\neg\mathcal{A}$ is a wff.

Our strategy for defining the truth function, v , will also be in two steps. The first step will handle truth for atomic sentences; the second step will handle truth for compound sentences.

Truth in SL

How can we define truth for an atomic sentence of SL? Consider, for example, the sentence M . Without an interpretation, we cannot say whether M is true or false. It might mean anything. If we use M to symbolize ‘The moon orbits the Earth’, then M is true. If use M to symbolize ‘The moon is a giant turnip’, then M false.

Moreover, the way you would discover whether or not M is true depends on what M means. If M means ‘It is monday,’ then you would need to check a calendar. If M means ‘Jupiter’s moon IO has significant volcanic activity,’ then you would need to check an astronomy text— and astronomers know because they sent satellites to observe IO.

When we give a symbolization key for SL, we provide an interpretation of the sentence letters that we use. The key gives an English language sentence for each sentence letter that we use. In this way, the interpretation specifies what each of the sentence letters *means*. However, this not enough to determine whether or not that sentence is true. The sentences about the moon, for instance, require that you know some rudimentary astronomy. Imagine a small child who became convinced that the moon is a giant turnip. She could understand what the sentence ‘The moon is a giant turnip’ means, but mistakenly think that it was true.

Consider another example: If M means ‘It is morning now’, then whether it is true or not depends on when you are reading this. I know what the sentence means, but— since I do not know when you will be reading this— I do not know whether it is true or false.

So an interpretation alone does not determine whether a sentence is true or false. Truth or falsity depends also on what the world is like. If M meant ‘The moon is a giant turnip’ and the real moon were a giant turnip, then M would be true. To put the point in a general way, truth or falsity is determined by an interpretation *plus* a way that the world is.

INTERPRETATION + STATE OF THE WORLD \implies TRUTH/FALSITY

In providing a logical definition of truth, we will not be able to give an account of how an atomic sentence is made true or false by the world. Instead, we will introduce a *truth value assignment*. Formally, this will be a function that tells us the truth value of all the atomic sentences. Call this function ‘ a ’ (for ‘assignment’). We define a for all sentence letters \mathcal{P} , such that

$$a(\mathcal{P}) = \begin{cases} 1 & \text{if } \mathcal{P} \text{ is true,} \\ 0 & \text{otherwise.} \end{cases}$$

The function a is determined by the meaning of the sentence letters together with the state of the world. If D means ‘It is dark outside’, then $a(D) = 1$ at night or during a heavy storm, while $a(D) = 0$ on a clear day.

You can think of a as being like a row of a truth table. Whereas a truth table row assigns a truth value to a few atomic sentences, the truth value assignment assigns a value to every atomic sentence of SL. There are infinitely many sentence letters, and the truth value assignment gives a value to each of them. When constructing a truth table, we only care about sentence letters that affect the truth value of sentences that interest us. As such, we ignore the rest. Strictly speaking, every row of a truth table gives a *partial* truth value assignment.

It is important to note that the truth value assignment, a , is not part of the language SL. Rather, it is part of the mathematical machinery that we are using to describe SL. It encodes which atomic sentences are true and which are false.

We now define the truth function, v , using the same recursive structure that we used to define a wff of SL.

1. If \mathcal{A} is a sentence letter, then $v(\mathcal{A}) = a(\mathcal{A})$.
2. If \mathcal{A} is $\neg\mathcal{B}$ for some sentence \mathcal{B} , then

$$v(\mathcal{A}) = \begin{cases} 1 & \text{if } v(\mathcal{B}) = 0, \\ 0 & \text{otherwise.} \end{cases}$$

3. If \mathcal{A} is $\mathcal{B} \& \mathcal{C}$ for some sentences \mathcal{B}, \mathcal{C} , then

$$v(\mathcal{A}) = \begin{cases} 1 & \text{if } v(\mathcal{B}) = 1 \text{ and } v(\mathcal{C}) = 1, \\ 0 & \text{otherwise.} \end{cases}$$

It might seem as if this definition is circular, because it uses the word ‘and’ in trying to define ‘and.’ Notice, however, that this is not a definition of the English word ‘and’; it is a definition of truth for sentences of SL containing the logical symbol ‘ $\&$.’ We define truth for such sentences using the metalanguage word ‘and.’ There is nothing circular about that.

4. If \mathcal{A} is $(\mathcal{B} \vee \mathcal{C})$ for some sentences \mathcal{B}, \mathcal{C} , then

$$v(\mathcal{A}) = \begin{cases} 0 & \text{if } v(\mathcal{B}) = 0 \text{ and } v(\mathcal{C}) = 0, \\ 1 & \text{otherwise.} \end{cases}$$

5. If \mathcal{A} is $\mathcal{B} \rightarrow \mathcal{C}$ for some sentences \mathcal{B}, \mathcal{C} , then

$$v(\mathcal{A}) = \begin{cases} 0 & \text{if } v(\mathcal{B}) = 1 \text{ and } v(\mathcal{C}) = 0, \\ 1 & \text{otherwise.} \end{cases}$$

6. If \mathcal{A} is $\mathcal{B} \leftrightarrow \mathcal{C}$ for some sentences \mathcal{B}, \mathcal{C} , then

$$v(\mathcal{A}) = \begin{cases} 1 & \text{if } v(\mathcal{B}) = v(\mathcal{C}), \\ 0 & \text{otherwise.} \end{cases}$$

Since the definition of v has the same structure as the definition of a wff, we know that v assigns a value to *every* wff of SL. Since the sentences of SL and the wffs of SL are the same, this means that v returns the truth value of every sentence of SL.

Truth in SL is always truth *relative to* some truth value assignment, because the definition of truth for SL does not say whether a given sentence is true or false. Rather, it says how the truth of that sentence relates to a truth value assignment.

Other concepts in SL

Working with SL so far, we have done without a precise definition of ‘tautology’, ‘contradiction’, and so on. Truth tables provided a way to *check if* a sentence was a tautology in SL, but they did not *define* what it means to be a tautology in SL. We will give definitions of these concepts for SL in terms of entailment.

The relation of semantic entailment, ‘ \mathcal{A} entails \mathcal{B} ’, means that there is no truth value assignment for which \mathcal{A} is true and \mathcal{B} is false. Put differently, it means that \mathcal{B} is true for any and all truth value assignments for which \mathcal{A} is true.

We abbreviate this with a symbol called the *double turnstile*: $\mathcal{A} \models \mathcal{B}$ means ‘ \mathcal{A} semantically entails \mathcal{B} .’

We can talk about entailment between more than two sentences:

$$\{\mathcal{A}_1, \mathcal{A}_2, \mathcal{A}_3, \dots\} \models \mathcal{B}$$

means that there is no truth value assignment for which all of the sentences in the set $\{\mathcal{A}_1, \mathcal{A}_2, \mathcal{A}_3, \dots\}$ are true and \mathcal{B} is false.

We can also use symbol with just one sentence: $\models C$ means that C is true for all truth value assignments. This is equivalent to saying that that the sentence is entailed by anything.

Now we can give concise definitions for various concepts of SL:

A TAUTOLOGY IN SL is a sentence \mathcal{A} such that $\models \mathcal{A}$.

A CONTRADICTION IN SL is a sentence \mathcal{A} such that $\models \neg\mathcal{A}$.

A sentence is CONTINGENT IN SL if and only if it is neither a tautology nor a contradiction.

An argument “ $\mathcal{P}_1, \mathcal{P}_2, \dots, \therefore C$ ” is VALID IN SL if and only if $\{\mathcal{P}_1, \mathcal{P}_2, \dots\} \models C$.

Two sentences \mathcal{A} and \mathcal{B} are LOGICALLY EQUIVALENT IN SL if and only if both $\mathcal{A} \models \mathcal{B}$ and $\mathcal{B} \models \mathcal{A}$.

Logical consistency is somewhat harder to define in terms of semantic entailment. Instead, we will define it in this way:

The set $\{\mathcal{A}_1, \mathcal{A}_2, \mathcal{A}_3, \dots\}$ is CONSISTENT IN SL if and only if there is a least one truth value assignment for which all of the sentences are true. The set is INCONSISTENT IN SL if and if only there is no such assignment.

5.2 Interpretations and models in QL

In SL, an interpretation or symbolization key specifies what each of the sentence letters means. The interpretation of a sentence letter along with the state of the world determines whether the sentence letter is true or false. Since the basic units are sentence letters, an interpretation only matters insofar as it makes sentence letters true or false. Formally, the semantics for SL is strictly in terms of truth value assignments. Two interpretations are the same, formally, if they make for the same truth value assignment.

What is an interpretation in QL? Like a symbolization key for QL, an interpretation requires a UD, a schematic meaning for each of the predicates, and an object that is picked out by each constant. For example:

UD: comic book characters

Fx: x fights crime.

b: the Batman

w: Bruce Wayne

Consider the sentence Fb . The sentence is true on this interpretation, but—just as in SL—the sentence is not true *just because* of the interpretation. Most people in our culture know that Batman fights crime, but this requires a modicum of knowledge about comic books. The sentence Fb is true because of the interpretation *plus* some facts about comic books. This is especially obvious when we consider Fw . Bruce Wayne is the secret identity of the Batman in the comic books—the identity claim $b = w$ is true—so Fw is true. Since it is a *secret* identity, however, other characters do not know that Fw is true even though they know that Fb is true.

We could try to characterize this as a truth value assignment, as we did for SL. The truth value assignment would assign 0 or 1 to each atomic wff: Fb , Fw , and so on. If we were to do that, however, we might just as well translate the sentences from QL to SL by replacing Fb and Fw with sentence letters. We could then rely on the definition of truth for SL, but at the cost of ignoring all the logical structure of predicates and terms. In writing a symbolization key for QL, we do not give separate definitions for Fb and Fw . Instead, we give a meaning to F , b , and w . This is essential because we want to be able to use quantifiers. There is no adequate way to translate $\forall xFx$ into SL.

So we want a formal counterpart to an interpretation for predicates and constants, not just for sentences. We cannot use a truth value assignment for this, because a predicate is neither true nor false. In the interpretation given above, F is true *of* the Batman (i.e., Fb is true), but it makes no sense at all to ask whether F on its own is true. It would be like asking whether the English language fragment ‘. . . fights crime’ is true.

What does an interpretation do for a predicate, if it does not make it true or false? An interpretation helps to pick out the objects to which the predicate applies. Interpreting Fx to mean ‘ x fights crime’ picks out Batman, Superman, Spiderman, and other heroes as the things that are F s. Formally, this is a set of members of the UD to which the predicate applies; this set is called the EXTENSION of the predicate.

Many predicates have indefinitely large extensions. It would be impractical to try and write down all of the comic book crime fighters individually, so instead we use an English language expression to interpret the predicate. This is somewhat imprecise, because the interpretation alone does not tell you which members of the UD are in the extension of the predicate. In order to figure out whether a particular member of the UD is in the extension of the predicate (to figure out whether Black Lightning fights crime, for instance), you need to know about comic books. In general, the extension of a predicate is the result of an interpretation *along with* some facts.

Sometimes it is possible to list all of the things that are in the extension of a predicate. Instead of writing a schematic English sentence, we can write down

the extension as a set of things. Suppose we wanted to add a one-place predicate M to the key above. We want Mx to mean ‘ x lives in Wayne Manor’, so we write the extension as a set of characters:

$$\text{extension}(M) = \{\text{Bruce Wayne, Alfred the butler, Dick Grayson}\}$$

You do not need to know anything about comic books to be able to determine that, on this interpretation, Mb is true: Bruce Wayne is just specified to be one of the things that is M . Similarly, $\exists xMx$ is obviously true on this interpretation: There is at least one member of the UD that is an M — in fact, there are three of them.

What about the sentence $\forall xMx$? The sentence is false, because it is not true that all members of the UD are M . It requires the barest minimum of knowledge about comic books to know that there are other characters besides just these three. Although we specified the extension of M in a formally precise way, we still specified the UD with an English language description. Formally speaking, a UD is just a set of members.

The formal significance of a predicate is determined by its extension, but what should we say about constants like b and w ? The meaning of a constant determines which member of the UD is picked out by the constant. The individual that the constant picks out is called the REFERENT of the constant. Both b and w have the same referent, since they both refer to the same comic book character. You can think of a constant letter as a name and the referent as the thing named. In English, we can use the different names ‘Batman’ and ‘Bruce Wayne’ to refer to the same comic book character. In this interpretation, we can use the different constants ‘ b ’ and ‘ w ’ to refer to the same member of the UD.

Sets

We use curly brackets ‘{’ and ‘}’ to denote sets. The members of the set can be listed in any order, separated by commas. The fact that sets can be in any order is important, because it means that {foo, bar} and {bar, foo} are the same set.

It is possible to have a set with no members in it. This is called the EMPTY SET. The empty set is sometimes written as {}, but usually it is written as the single symbol \emptyset .

Models

As we have seen, an interpretation in QL is only formally significant insofar as it determines a UD, an extension for each predicate, and a referent for each constant. We call this formal structure a MODEL for QL.

To see how this works, consider this symbolization key:

UD: People who played as part of the ‘Three Stooges’
Hx: x had head hair.
f: Mister Fine

If you do not know anything about this three stooges, you will not be able to say which sentences of QL are true on this interpretation. Perhaps you just remember Larry, Curly, and Moe. Is the sentence Hf true or false? It depends on which of the stooges is Mister Fine.

What is the model that corresponds to this interpretation? The UD will have six members: Larry Fine, Moe Howard, and Curly Howard, Shemp Howard, Joe Besser, and Curly Joe DeRita. Curly, Joe, and Curly Joe were the only completely bald stooges. The result is this model:

$$\begin{aligned} \text{UD} &= \{\text{Larry, Curly, Moe, Shemp, Joe, Curly Joe}\} \\ \text{extension}(H) &= \{\text{Larry, Moe, Shemp}\} \\ \text{referent}(f) &= \text{Larry} \end{aligned}$$

You do not need to know anything about the Three Stooges in order to evaluate whether sentences are true or false in this *model*. Hf is true, since the referent of f (Larry) is in the extension of H . Both $\exists xHx$ and $\exists x\neg Hx$ are true, since there is at least one member of the UD that is in the extension of H and at least one member that is not in the extension of H . In this way, the model captures all of the formal significance of the interpretation.

Now consider this interpretation:

UD: whole numbers less than 10
Ex: x is even.
Nx: x is negative.
Lxy: x is less than y .
Txyz: x times y equals z .

What is the model that goes with this interpretation? The UD is the set $\{1, 2, 3, 4, 5, 6, 7, 8, 9\}$.

The extension of a one-place predicate like E or N is just the subset of the UD of which the predicate is true. Roughly speaking, the extension of the predicate

E is the set of E s in the UD. The extension of E is the subset $\{2, 4, 6, 8\}$. There are many even numbers besides these four, but these are the only members of the UD that are even. There are no negative numbers in the UD, so N has an empty extension; i.e. $\text{extension}(N) = \emptyset$.

The extension of a two-place predicate like L is somewhat vexing. It seems as if the extension of L ought to contain 1, since 1 is less than all the other numbers; it ought to contain 2, since 2 is less than all of the other numbers besides 1; and so on. Every member of the UD besides 9 is less than some member of the UD. What would happen if we just wrote $\text{extension}(L) = \{1, 2, 3, 4, 5, 6, 7, 8\}$?

The problem is that sets can be written in any order, so this would be the same as writing $\text{extension}(L) = \{8, 7, 6, 5, 4, 3, 2, 1\}$. This does not tell us which of the members of the set are less than which other members.

We need some way of showing that 1 is less than 8 but that 8 is not less than 1. The solution is to have the extension of L consist of pairs of numbers. An ORDERED PAIR is like a set with two members, except that the order *does* matter. We write ordered pairs with angle brackets ' \langle ' and ' \rangle '. The ordered pair $\langle \text{foo}, \text{bar} \rangle$ is different than the ordered pair $\langle \text{bar}, \text{foo} \rangle$. The extension of L is a collection of ordered pairs, all of the pairs of numbers in the UD such that the first number is less than the second. Writing this out completely:

$$\begin{aligned} \text{extension}(L) = \{ & \langle 1,2 \rangle, \langle 1,3 \rangle, \langle 1,4 \rangle, \langle 1,5 \rangle, \langle 1,6 \rangle, \langle 1,7 \rangle, \\ & \langle 1,8 \rangle, \langle 1,9 \rangle, \langle 2,3 \rangle, \langle 2,4 \rangle, \langle 2,5 \rangle, \langle 2,6 \rangle, \langle 2,7 \rangle, \langle 2,8 \rangle, \langle 2,9 \rangle, \\ & \langle 3,4 \rangle, \langle 3,5 \rangle, \langle 3,6 \rangle, \langle 3,7 \rangle, \langle 3,8 \rangle, \langle 3,9 \rangle, \langle 4,5 \rangle, \langle 4,6 \rangle, \langle 4,7 \rangle, \\ & \langle 4,8 \rangle, \langle 4,9 \rangle, \langle 5,6 \rangle, \langle 5,7 \rangle, \langle 5,8 \rangle, \langle 5,9 \rangle, \langle 6,7 \rangle, \langle 6,8 \rangle, \langle 6,9 \rangle, \\ & \langle 7,8 \rangle, \langle 7,9 \rangle, \langle 8,9 \rangle \} \end{aligned}$$

Three-place predicates will work similarly; the extension of a three-place predicate is a set of ordered triples where the predicate is true of those three things *in that order*. So the extension of T in this model will contain ordered triples like $\langle 2,4,8 \rangle$, because $2 \times 4 = 8$.

Generally, the extension of an n -place predicate is a set of all ordered n -tuples $\langle a_1, a_2, \dots, a_n \rangle$ such that a_1 – a_n are members of the UD and the predicate is true of a_1 – a_n in that order.

5.3 Working with models

We will use the double turnstile symbol for QL much as we did for SL. When \mathcal{A} and \mathcal{B} are two sentences of QL, $\mathcal{A} \models \mathcal{B}$ means that there is no model in which \mathcal{A} is true and \mathcal{B} is false. $\models \mathcal{A}$ means that \mathcal{A} is true in every model.

We can now give definitions for various concepts in QL. Because we are using the same symbol, these definitions will look similar to the definitions in SL. Remember, however, that the definitions in QL are in terms of *models* rather than in terms of truth value assignments.

A TAUTOLOGY IN QL is a sentence \mathcal{A} that is true in every model; i.e., $\models \mathcal{A}$.

A CONTRADICTION IN QL is a sentence \mathcal{A} that is false in every model; i.e., $\models \neg \mathcal{A}$.

A sentence is CONTINGENT IN QL if and only if it is neither a tautology nor a contradiction.

An argument “ $\mathcal{P}_1, \mathcal{P}_2, \dots, \therefore \mathcal{C}$ ” is VALID IN QL if and only if there is no model in which all of the premises are true and the conclusion is false; i.e., $\{\mathcal{P}_1, \mathcal{P}_2, \dots\} \models \mathcal{C}$. It is INVALID IN QL otherwise.

Two sentences \mathcal{A} and \mathcal{B} are LOGICALLY EQUIVALENT IN SL if and only if both $\mathcal{A} \models \mathcal{B}$ and $\mathcal{B} \models \mathcal{A}$.

The set $\{\mathcal{A}_1, \mathcal{A}_2, \mathcal{A}_3, \dots\}$ is CONSISTENT IN QL if and only if there is a least model in which all of the sentences are true. The set is INCONSISTENT IN QL if and only if there is no such model.

Constructing models

Suppose we want to show that $\forall x Axx \rightarrow Bd$ is *not* a tautology. This requires showing that the sentence is not true in every model; i.e., that it is false in some model. If we can provide just one model in which the sentence is false, then we will have shown that the sentence is not a tautology.

What would such a model look like? In order for $\forall x Axx \rightarrow Bd$ to be false, the antecedent ($\forall x Axx$) must be true, and the consequent (Bd) must be false.

To construct such a model, we start with a UD. It will be easier to specify extensions for predicates if we have a small UD, so start with a UD that has just one member. Formally, this single member might be anything. Let's say it is the city of Paris.

We want $\forall x Axx$ to be true, so we want all members of the UD to be paired with themselves in the extension of A ; this means that the extension of A must be $\{\langle \text{Paris}, \text{Paris} \rangle\}$.

We want Bd to be false, so the referent of d must not be in the extension of B . We give B an empty extension.

Since Paris is the only member of the UD, it must be the referent of d . The model we have constructed looks like this:

$$\begin{aligned} \text{UD} &= \{\text{Paris}\} \\ \text{extension}(A) &= \{\langle \text{Paris}, \text{Paris} \rangle\} \\ \text{extension}(B) &= \emptyset \\ \text{referent}(d) &= \text{Paris} \end{aligned}$$

Strictly speaking, a model specifies an extension for *every* predicate of QL and a referent for *every* constant. As such, it is generally impossible to write down a complete model. That would require writing down infinitely many extensions and infinitely many referents. However, we do not need to consider every predicate in order to show that there are models in which $\forall xAxx \rightarrow Bd$ is false. Predicates like H and constants like f_{13} make no difference to the truth or falsity of this sentence. It is enough to specify extensions for A and B and a referent for d , as we have done. This provides a *partial model* in which the sentence is false.

Perhaps you are wondering: What does the predicate A mean in English? The partial model could correspond to an interpretation like this one:

UD: Paris
 Axy : x is in the same country as y .
 Bx : x was founded in the 20th century.
 d : the City of Lights

However, all that the partial model tells us is that A is a predicate which is true of Paris and Paris. There are indefinitely many predicates in English that have this extension. Axy might instead translate ' x is the same size as y ' or ' x and y are both cities.' Similarly, Bx is some predicate that does not apply to Paris; it might instead translate ' x is on an island' or ' x is a subcompact car.' When we specify the extensions of A and B , we do not specify what English predicates A and B should be used to translate. We are concerned with whether the $\forall xAxx \rightarrow Bd$ comes out true or false, and all that matters for truth and falsity in QL is the information in the model: the UD, the extensions of predicates, and the referents of constants.

We can just as easily show that $\forall xAxx \rightarrow Bd$ is not a contradiction. We need only specify a model in which $\forall xAxx \rightarrow Bd$ is true; i.e., a model in which either $\forall xAxx$ is false or Bd is true. Here is one such partial model:

$$\begin{aligned} \text{UD} &= \{\text{Paris}\} \\ \text{extension}(A) &= \{\langle \text{Paris}, \text{Paris} \rangle\} \\ \text{extension}(B) &= \{\text{Paris}\} \\ \text{referent}(d) &= \text{Paris} \end{aligned}$$

We have now shown that $\forall xAxx \rightarrow Bd$ is neither a tautology nor a contradic-

tion. By the definition of ‘contingent in QL,’ this means that $\forall xAxx \rightarrow Bd$ is contingent. In general, showing that a sentence is contingent will require two models: one in which the sentence is true and another in which the sentence is false.

Suppose we want to show that $\forall xSx$ and $\exists xSx$ are not logically equivalent. We need to construct a model in which the two sentences have different truth values; we want one of them to be true and the other to be false. We start by specifying a UD. Again, we make the UD small so that we can specify extensions easily. We will need at least two members. Let the UD be {Duke, Miles}. (If we chose a UD with only one member, the two sentences would end up with the same truth value. In order to see why, try constructing some partial models with one-member UD's.)

We can make $\exists xSx$ true by including something in the extension of S , and we can make $\forall xSx$ false by leaving something out of the extension of S . It does not matter which one we include and which one we leave out. Making Duke the only S , we get a partial model that looks like this:

$$\begin{aligned} \text{UD} &= \{\text{Duke, Miles}\} \\ \text{extension}(S) &= \{\text{Duke}\} \end{aligned}$$

This partial model shows that the two sentences are *not* logically equivalent.

Back on p. 65, we said that this argument would be invalid in QL:

$$(Rc \& K_1c) \& Tc$$

$$\therefore Tc \& K_2c$$

In order to show that it is invalid, we need to show that there is some model in which the premises are true and the conclusion is false. We can construct such a model deliberately. Here is one way to do it:

$$\begin{aligned} \text{UD} &= \{\text{Björk}\} \\ \text{extension}(T) &= \{\text{Björk}\} \\ \text{extension}(K_1) &= \{\text{Björk}\} \\ \text{extension}(K_2) &= \emptyset \\ \text{extension}(R) &= \{\text{Björk}\} \\ \text{referent}(c) &= \text{Björk} \end{aligned}$$

Similarly, we can show that a set of sentences is consistent by constructing a model in which all of the sentences are true.

	YES	NO
Is \mathcal{A} a tautology?	show that \mathcal{A} must be true in any model	<i>construct a model</i> in which \mathcal{A} is false
Is \mathcal{A} a contradiction?	show that \mathcal{A} must be false in any model	<i>construct a model</i> in which \mathcal{A} is true
Is \mathcal{A} contingent?	<i>construct two models:</i> one in which \mathcal{A} is true and another in which \mathcal{A} is false	either show that \mathcal{A} is a tautology or that it is a contradiction
Are \mathcal{A} and \mathcal{B} equivalent?	show that \mathcal{A} and \mathcal{B} must have the same truth value in any model	<i>construct a model</i> in which \mathcal{A} and \mathcal{B} have different truth values
Is the set \mathbb{A} consistent?	<i>construct a model</i> in which all the sentences in \mathbb{A} are true	show that the sentences could not all be true in any model
Is the argument ' $\mathcal{P}, \therefore \mathcal{C}$ ' valid?	show that any model in which \mathcal{P} is true must be a model in which \mathcal{C} is true	<i>construct a model</i> in which \mathcal{P} is true and \mathcal{C} is false

Table 5.1: It is relatively easy to answer a question if you can do it by constructing a model or two. It is much harder if you need to reason about all possible models. This table shows when constructing models is enough.

Reasoning about all models

As we have seen, to show that a sentence is *not* a tautology requires just one carefully specified model: a model in which the sentence is false. To show that something is a tautology, on the other hand, it would not be enough to construct ten, one hundred, or even a thousand models in which the sentence is true. It is only a tautology if it is true in *every* model, and there are infinitely many models. This cannot be avoided by constructing partial models, because there are infinitely many partial models. Even if the sentence contains no constants and only a single, one-place predicate, there are two different partial models that have a 1-member UD, four different partial models that have a 2-member UD, eight different partial models that have a 3-member UD, and so on to infinity. In order to show that something is a tautology, we need to show something about all of these models.

Suppose that we want to show that $(\forall x)(Px \rightarrow Px)$ is a tautology. We can reason in this way:

Consider any member of the UD; call it α . It must be that α is either in the extension of P or it is not. If it is in the extension of P , then P is true of it, the consequent of $P\alpha \rightarrow P\alpha$ is true, and so $P\alpha \rightarrow P\alpha$ is true. If it is not in the extension of P , then P is false of it, the antecedent of $P\alpha \rightarrow P\alpha$ is false, and so $P\alpha \rightarrow P\alpha$ is true. This is true for any member of the UD, regardless of exact UD and regardless of the exact extension of P . Therefore $(\forall x)(Px \rightarrow Px)$ is true in any model; it is a tautology.

Here we are forced to reason in English about possible models. This is never an easy thing to do. ($P\alpha \rightarrow P\alpha$ is not even a sentence of QL, because α is not a symbol of QL; α is just a metalinguistic symbol used in this argument.)

Similar problems arise if we want to show that a sentence is a contradiction, that two sentences are equivalent, or that set of sentences is inconsistent, or that an argument is valid. For example, consider this argument: $\forall x(Hx \& Jx) \therefore \forall xHx$. We can show that this argument is valid by reasoning in this way:

Consider a model in which the premise $\forall x(Hx \& Jx)$ is true. Since the conjunction holds for every object in the UD, every object in the UD must be in the extension of both H and J . Since every object is in the extension of H , $(\forall x)Hx$ is true in this model. Since we did not assume anything about the model besides $\forall x(Hx \& Jx)$ being true, this must be the case for any model in which $\forall x(Hx \& Jx)$ is true. So $\forall x(Hx \& Jx) \models \forall xHx$.

Here we can only show that the argument is valid by considering what must be

true in every model in which the premise is true. The reasoning is somewhat complicated, even for a simple argument like this one.

In conclusion: There are some things that can only be shown by reasoning about all possible models. There are other things we can show by carefully constructing a model or two. Table 5.1 summarizes which things are which.

5.4 Semantics for identity

Identity is a special predicate of QL. We write it a bit differently than other two-place predicates: $x = y$ instead of Ixy . We also do not need to include it in a symbolization key. The sentence $x = y$ always means ‘ x is identical to y ,’ and it cannot be interpreted to mean anything else. In the same way, when you construct a model, you do not get to pick and choose which ordered pairs go into the extension of the identity predicate. It always contains just the ordered pair of each object in the UD with itself.

The sentence $\forall x Ixx$, which contains an ordinary two-place predicate, is contingent. Whether it is true for an interpretation depends on how you interpret I , and whether it is true in a model depends on the extension of I .

The sentence $\forall x x = x$ is a tautology. The extension of identity will always make it true.

Notice that although identity always has the same interpretation, it does not always have the same extension. The extension of identity depends on the UD. If the UD in a model is the set $\{\text{Doug}\}$, then $\text{extension}(=)$ in that model is $\{\langle \text{Doug}, \text{Doug} \rangle\}$. If the UD is the set $\{\text{Doug}, \text{Omar}\}$, then $\text{extension}(=)$ in that model is $\{\langle \text{Doug}, \text{Doug} \rangle, \langle \text{Omar}, \text{Omar} \rangle\}$. And so on.

If the referent of two constants is the same, then anything which is true of one is true of the other. For example, if $\text{referent}(a) = \text{referent}(b)$, then $Aa \leftrightarrow Ab$, $Ba \leftrightarrow Bb$, $Ca \leftrightarrow Cb$, $Rca \leftrightarrow Rcb$, $\forall x Rxa \leftrightarrow \forall x Rxb$, and so on for any two sentences containing a and b . However, the reverse is not true.

It is possible that anything which is true of a is also true of b , yet for a and b still to have different referents. This may seem puzzling, but it is easy to construct a model that shows this. Consider this model:

$$\begin{aligned} \text{UD} &= \{\text{Rosencrantz}, \text{Guildenstern}\} \\ \text{referent}(a) &= \text{Rosencrantz} \\ \text{referent}(b) &= \text{Guildenstern} \\ \text{for all predicates } \mathcal{P}, \text{extension}(\mathcal{P}) &= \emptyset \\ \text{extension}(=) &= \{\langle \text{Rosencrantz}, \text{Rosencrantz} \rangle, \\ &\quad \langle \text{Guildenstern}, \text{Guildenstern} \rangle\} \end{aligned}$$

This specifies an extension for every predicate of QL: All the infinitely-many predicates are empty. This means that both Aa and Ab are false, and they are equivalent; both Ba and Bb are false; and so on for any two sentences that contain a and b . Yet a and b refer to different things. We have written out the extension of identity to make this clear: The ordered pair $\langle \text{referent}(a), \text{referent}(b) \rangle$ is not in it. In this model, $a = b$ is false and $a \neq b$ is true.

5.5 Truth in QL

We now want to define truth in QL. For SL, we split the task into two parts: a truth value assignment for sentence letters (a) and a truth function for all sentences (v) that covered the way that complex sentences can be built out of sentence letters and connectives.

Truth in QL will not be this straightforward, because there are wffs of QL (well-formed formulae) that are not sentences. The wff $(\forall x)Px$ is a sentence, and so it is either true or false in any given model. The sub-formula Px is also a wff, but it is not a sentence. Since x is a variable and not a constant, Px is neither true nor false. Without a quantifier to specify how we should understand the variable, Px does not express a complete thought. This means that we cannot define the truth of $(\forall x)Px$ in terms of the truth of Px . We need another notion that does apply to Px . This notion is SATISFACTION.

We will define ‘satisfaction’ for all wffs of QL and then define ‘truth’ in terms of satisfaction. This will allow for a recursive definition, even though ‘truth’ is not defined for all wffs.

Satisfaction

We begin by defining a VARIABLE ASSIGNMENT. For each variable, the variable assignment designates an object from the UD that will stand in for that variable. We define the variable assignment as a function ‘ a ’ (for ‘assignment’) such that, for any variable χ , $a(\chi)$ is the member of the UD picked out to stand in for χ . There are infinitely many variables in QL, and a assigns members of the UD for all of them.

Let z be a variable and let α be some member of the UD. We now define $a[\alpha|z]$ as a function that obeys the the variable assignment a for every variable *except* z ; it assigns α to z , regardless of what a assigned to z . We use the greek letter alpha, α , to indicate that it is neither a symbol of QL nor a metavariable that ranges over symbols of QL: α is an actual member of the UD, not a constant of QL or a name of the member.

If we plug in x_7 for z , we get these values: $a[\alpha|x_7](y) = a(y)$, $a[\alpha|x_7](z) = a(z)$, $a[\alpha|x_7](x_2) = a(x_2)$, and so on, but $a[\alpha|x_7](x_7) = \alpha$. In general terms:

$$a[\alpha|z](\chi) = \begin{cases} \alpha & \text{if } \chi = z, \\ a(\chi) & \text{otherwise.} \end{cases}$$

We can now define satisfaction in terms of a variable assignment. A variable assignment a SATISFIES a wff \mathcal{A} in a model \mathbb{M} if and only if

1. If \mathcal{A} is an atomic wff of the form $\mathcal{P}t_1 \dots t_n$, then \mathcal{A} is satisfied by a in \mathbb{M} if and only if $\langle \alpha_1 \dots \alpha_n \rangle$ is in the extension of \mathcal{P} in \mathbb{M} ; where

$$\alpha_i = \begin{cases} \text{referent}(t_i) & \text{if } t_i \text{ is a constant,} \\ a(t_i) & \text{if } t_i \text{ is a variable.} \end{cases}$$

For each term t_i , α_i is the object picked out by that term. \mathcal{A} is satisfied if the n -tuple of objects picked out by $t_1 \dots t_n$ is in the extension of the predicate \mathcal{P} .

Next, we begin to define satisfaction for the sentential connectives.

2. If \mathcal{A} is of the form $\neg\mathcal{B}$ for some wff \mathcal{B} , then a satisfies \mathcal{A} in \mathbb{M} if and only if a does not satisfy \mathcal{B} in \mathbb{M} .

We define satisfaction for logical negation in part by using the English language word ‘not.’ This may seem circular, but it is not; remember that ‘not’ is a word of the metalanguage that we use to explicate the object language symbol ‘ \neg .’

Satisfaction for the other sentential connectives proceeds in a similar way:

3. If \mathcal{A} is of the form $(\mathcal{B} \& \mathcal{C})$ for some wffs \mathcal{B} and \mathcal{C} , then a satisfies \mathcal{A} in \mathbb{M} if and only if a satisfies both \mathcal{B} and \mathcal{C} in \mathbb{M} .
4. If \mathcal{A} is of the form $(\mathcal{B} \vee \mathcal{C})$ for some wffs \mathcal{B} and \mathcal{C} , then a satisfies \mathcal{A} in \mathbb{M} if and only if a satisfies \mathcal{B} , \mathcal{C} , or both \mathcal{B} and \mathcal{C} in \mathbb{M} .
5. If \mathcal{A} is of the form $(\mathcal{B} \rightarrow \mathcal{C})$ for some wffs \mathcal{B} and \mathcal{C} , then a satisfies \mathcal{A} in \mathbb{M} if and only if a satisfies \mathcal{C} or does not satisfy \mathcal{B} (or both) in \mathbb{M} .
6. If \mathcal{A} is of the form $(\mathcal{B} \leftrightarrow \mathcal{C})$ for some wffs \mathcal{B} and \mathcal{C} , then a satisfies \mathcal{A} in \mathbb{M} if and only if a satisfies both or neither \mathcal{B} and \mathcal{C} in \mathbb{M} .

Now we need to consider quantifiers. If Px would be satisfied in \mathbb{M} regardless of what is assigned to x , then $\forall xPx$ would be satisfied in \mathbb{M} . In considering $\forall xPx$, we effectively ignore what a assigns to x . Putting this in a general way, we say:

7. If \mathcal{A} is $\forall\chi\mathcal{B}$ for some wff \mathcal{B} and some variable χ , then a satisfies \mathcal{A} in \mathbb{M} if and only if $a[\alpha|\chi]$ satisfies \mathcal{B} in M for every member α of the UD.

You can think of this as saying that a satisfies \mathcal{A} if it does so regardless of what object you instantiate for x .

8. If \mathcal{A} is $\exists\chi\mathcal{B}$ for some wff \mathcal{B} and some variable χ , then a satisfies \mathcal{A} in \mathbb{M} if and only if $a[\alpha|\chi]$ satisfies \mathcal{B} in M for at least one member α of the UD.

This definition follows the same structure as the definition of a wff for QL, so we know that every wff of QL will be covered by this definition. For a model \mathbb{M} and a variable assignment a , any wff will either be satisfied or not not be satisfied. No wffs are left out or assigned conflicting values.

Truth

Recall that a sentence of QL is a wff with no free variables. Consider a simple sentence like $(\forall x)Px$. By part 7 in the definition of satisfaction, this sentence is satisfied if $a[\alpha|x]$ satisfies Px in \mathbb{M} for every α in the UD. By part 1 of the definition, this will be the case if every α is in the extension of P . Whether $(\forall x)Px$ is satisfied does not depend on the particular variable assignment a . If this sentence is satisfied, then it is true. This is a formalization of what we have said all along: $\forall xPx$ is true if everything in the UD is in the extension of P .

We can give this general definition: A sentence \mathcal{A} is TRUE IN \mathbb{M} if and only if some variable assignment satisfies \mathcal{A} in M ; \mathcal{A} is FALSE IN \mathbb{M} otherwise.

(It would be equivalent to say \mathcal{A} is TRUE IN \mathbb{M} if and only if *every* variable assignment satisfies \mathcal{A} in M . Can you explain why?)

Truth in QL is *truth in a model*. Sentences of QL are not flat-footedly true or false as mere symbols, but only relative to a model. A model provides the meaning of the symbols, insofar as it makes any difference to truth and falsity.

Practice Exercises

★ **Part A** Determine whether each sentence is true or false in the model given.

$UD = \{\text{Corwin, Benedict}\}$
 $\text{extension}(A) = \{\text{Corwin, Benedict}\}$
 $\text{extension}(B) = \{\text{Benedict}\}$
 $\text{extension}(N) = \emptyset$
 $\text{referent}(c) = \text{Corwin}$

1. Bc
2. $Ac \leftrightarrow \neg Nc$
3. $Nc \rightarrow (Ac \vee Bc)$
4. $\forall xAx$
5. $\forall x\neg Bx$
6. $\exists x(Ax \& Bx)$
7. $\exists x(Ax \rightarrow Nx)$
8. $\forall x(Nx \vee \neg Nx)$
9. $\exists xBx \rightarrow \forall xAx$

★ **Part B** Determine whether each sentence is true or false in the model given.

$UD = \{\text{Waylan, Willy, Johnny}\}$
 $\text{extension}(H) = \{\text{Waylan, Willy, Johnny}\}$
 $\text{extension}(W) = \{\text{Waylan, Willy}\}$
 $\text{extension}(R) = \{\langle \text{Waylan, Willy} \rangle, \langle \text{Willy, Johnny} \rangle, \langle \text{Johnny, Waylan} \rangle\}$
 $\text{referent}(m) = \text{Johnny}$

1. $\exists x(Rxm \& Rmx)$
2. $\forall x(Rxm \vee Rmx)$
3. $\forall x(Hx \leftrightarrow Wx)$
4. $\forall x(Rxm \rightarrow Wx)$
5. $\forall x[Wx \rightarrow (Hx \& Wx)]$
6. $\exists xRxx$
7. $\exists x\exists yRxy$
8. $\forall x\forall yRxy$
9. $\forall x\forall y(Rxy \vee Ryx)$
10. $\forall x\forall y\forall z[(Rxy \& Ryz) \rightarrow Rxz]$

Part C Show that each of the following is contingent.

- ★ 1. $Da \& Db$
- ★ 2. $\exists xTxh$
- ★ 3. $Pm \& \neg\forall xPx$
4. $\forall zJz \leftrightarrow \exists yJy$
5. $\forall x(Wxmn \vee \exists yLxy)$
6. $\exists x(Gx \rightarrow \forall yMy)$

★ **Part D** Show that the following pairs of sentences are not logically equivalent.

1. Ja, Ka
2. $\exists xJx, Jm$
3. $\forall xRxx, \exists xRxx$
4. $\exists xPx \rightarrow Qc, \exists x(Px \rightarrow Qc)$
5. $\forall x(Px \rightarrow \neg Qx), \exists x(Px \& \neg Qx)$
6. $\exists x(Px \& Qx), \exists x(Px \rightarrow Qx)$
7. $\forall x(Px \rightarrow Qx), \forall x(Px \& Qx)$
8. $\forall x\exists yRxy, \exists x\forall yRxy$
9. $\forall x\exists yRxy, \forall x\exists yRyx$

Part E Show that the following sets of sentences are consistent.

1. $\{Ma, \neg Na, Pa, \neg Qa\}$
2. $\{Lee, Lef, \neg Lfe, \neg Lff\}$
3. $\{\neg(Ma \& \exists xAx), Ma \vee Fa, \forall x(Fx \rightarrow Ax)\}$
4. $\{Ma \vee Mb, Ma \rightarrow \forall x\neg Mx\}$
5. $\{\forall yGy, \forall x(Gx \rightarrow Hx), \exists y\neg Iy\}$
6. $\{\exists x(Bx \vee Ax), \forall x\neg Cx, \forall x[(Ax \& Bx) \rightarrow Cx]\}$
7. $\{\exists xXx, \exists xYx, \forall x(Xx \leftrightarrow \neg Yx)\}$
8. $\{\forall x(Px \vee Qx), \exists x\neg(Qx \& Px)\}$
9. $\{\exists z(Nz \& Ozz), \forall x\forall y(Oxy \rightarrow Oyx)\}$
10. $\{\neg\exists x\forall yRxy, \forall x\exists yRxy\}$

Part F Construct models to show that the following arguments are invalid.

1. $\forall x(Ax \rightarrow Bx), \therefore \exists xBx$
2. $\forall x(Rx \rightarrow Dx), \forall x(Rx \rightarrow Fx), \therefore \exists x(Dx \& Fx)$
3. $\exists x(Px \rightarrow Qx), \therefore \exists xPx$
4. $Na \& Nb \& Nc, \therefore \forall xNx$
5. $Rde, \exists xRxd, \therefore Red$
6. $\exists x(Ex \& Fx), \exists xFx \rightarrow \exists xGx, \therefore \exists x(Ex \& Gx)$
7. $\forall xOxc, \forall xOcx, \therefore \forall xOxx$
8. $\exists x(Jx \& Kx), \exists x\neg Kx, \exists x\neg Jx, \therefore \exists x(\neg Jx \& \neg Kx)$
9. $Lab \rightarrow \forall xLxb, \exists xLxb, \therefore Lbb$

Part G

- ★ 1. Show that $\{\neg Raa, \forall x(x = a \vee Rxa)\}$ is consistent.
- ★ 2. Show that $\{\forall x\forall y\forall z(x = y \vee y = z \vee x = z), \exists x\exists y x \neq y\}$ is consistent.
- ★ 3. Show that $\{\forall x\forall y x = y, \exists x x \neq a\}$ is inconsistent.
4. Show that $\exists x(x = h \& x = i)$ is contingent.

5. Show that $\{\exists x\exists y(Zx \& Zy \& x = y), \neg Zd, d = s\}$ is consistent.
6. Show that $\forall x(Dx \rightarrow \exists yTyx) \therefore \exists y\exists z y \neq z$ is invalid.

Part H Many logic books define consistency and inconsistency in this way: “A set $\{\mathcal{A}_1, \mathcal{A}_2, \mathcal{A}_3, \dots\}$ is inconsistent if and only if $\{\mathcal{A}_1, \mathcal{A}_2, \mathcal{A}_3, \dots\} \models (\mathcal{B} \& \neg\mathcal{B})$ for some sentence \mathcal{B} . A set is consistent if it is not inconsistent.”

Does this definition lead to any different sets being consistent than the definition on p. 87? Explain your answer.

Chapter 6

Proofs

Consider two arguments in SL:

Argument A

$$\begin{array}{l} P \vee Q \\ \neg P \\ \therefore Q \end{array}$$

Argument B

$$\begin{array}{l} P \rightarrow Q \\ P \\ \therefore Q \end{array}$$

Clearly, these are valid arguments. You can confirm that they are valid by constructing four-line truth tables. Argument A makes use of an inference form that is always valid: Given a disjunction and the negation of one of the disjuncts, the other disjunct follows as a valid consequence. This rule is called *disjunctive syllogism*.

Argument B makes use of a different valid form: Given a conditional and its antecedent, the consequent follows as a valid consequence. This is called *modus ponens*.

When we construct truth tables, we do not need to give names to different inference forms. There is no reason to distinguish modus ponens from a disjunctive syllogism. For this same reason, however, the method of truth tables does not clearly show *why* an argument is valid. If you were to do a 1028-line truth table for an argument *why* that contains ten sentence letters, then you could check to see if there were any lines on which the premises were all true and the conclusion were false. If you did not see such a line and provided you made no mistakes in constructing the table, then you would know that the argument was valid. Yet you would not be able to say anything further about why this particular argument was a valid argument form.

The aim of a *proof system* is to show that particular arguments are valid in a way that allows us to understand the reasoning involved in the argument. We begin with basic argument forms, like disjunctive syllogism and modus ponens. These forms can then be combined to make more complicated arguments, like this one:

- $$\begin{array}{l} (1) \neg L \rightarrow (J \vee L) \\ (2) \neg L \\ \therefore J \end{array}$$

By modus ponens, (1) and (2) entail $J \vee L$. This is an *intermediate conclusion*. It follows logically from the premises, but it is not the conclusion we want. Now $J \vee L$ and (2) entail J , by disjunctive syllogism. We do not need a new rule for this argument. The proof of the argument shows that it is really just a combination of rules we have already introduced.

Formally, a PROOF is a sequence of sentences. The first sentences of the sequence are assumptions; these are the premises of the argument. Every sentence later in the sequence follows from earlier sentences by one of the rules of proof. The final sentence of the sequence is the conclusion of the argument.

This chapter begins with a proof system for SL, which is then extended to cover QL and QL plus identity.

6.1 Basic rules for SL

In designing a proof system, we could just start with disjunctive syllogism and modus ponens. Whenever we discovered a valid argument which could not be proven with rules we already had, we could introduce new rules. Proceeding in this way, we would have an unsystematic grab bag of rules. We might accidentally add some strange rules, and we would surely end up with more rules than we need.

Instead, we will develop what is called a NATURAL DEDUCTION system. In a natural deduction system, there will be two rules for each logical operator: an INTRODUCTION rule that allows us to prove a sentence that has it as the main logical operator and an ELIMINATION rule that allows us to prove something given a sentence that has it as the main logical operator.

In addition to the rules for each logical operator, we will also have a reiteration rule. If you already have shown something in the course of a proof, the reiteration rule allows you to repeat it on a new line. For instance:

1	A	
2	A	R 1

When we add a line to a proof, we write the rule that justifies that line. We also write the numbers of the lines to which the rule was applied. The reiteration rule above is justified by one line, the line that you are reiterating. So the ‘R 1’ on line 2 of the proof means that the line is justified by the reiteration rule (R) applied to line 1.

Obviously, the reiteration rule will not allow us to show anything *new*. For that, we will need more rules. The remainder of this section will give introduction and elimination rules for all of the sentential connectives. This will give us a complete proof system for SL. Later in the chapter, we introduce rules for quantifiers and identity.

All of the rules introduced in this chapter are summarized starting on p. 152.

Conjunction

Think for a moment: What would you need to show in order to prove $E \& F$?

Of course, you could show $E \& F$ by proving E and separately proving F . This holds even if the two conjuncts are not atomic sentences. If you can prove $[(A \vee J) \rightarrow V]$ and $[(V \rightarrow L) \leftrightarrow (F \vee N)]$, then you have effectively proven

$$[(A \vee J) \rightarrow V] \& [(V \rightarrow L) \leftrightarrow (F \vee N)].$$

So this will be our conjunction introduction rule, which we abbreviate $\&I$:

m	A	
n	B	
	A & B	$\&I\ m, n$

A line of proof must be justified by some rule, and here we have ‘ $\&I\ m, n$.’ This means: Conjunction introduction applied to line m and line n . These are variables, not real line numbers; m is some line and n is some other line. In an actual proof, the lines are numbered 1, 2, 3, ... and rules must be applied to specific line numbers. When we define the rule, however, we use variables to underscore the point that the rule may be applied to any two lines that are already in the proof. If you have K on line 8 and L on line 15, you can prove $(K \& L)$ at some later point in the proof with the justification ‘ $\&I\ 8, 15$.’

Now, consider the elimination rule for conjunction. What are you entitled to conclude from a sentence like $E \& F$? Surely, you are entitled to conclude E ; if

E & F were true, then E would be true. Similarly, you are entitled to conclude F . This will be our conjunction elimination rule, which we abbreviate $\&E$:

m	$\mathcal{A} \& \mathcal{B}$	
	\mathcal{A}	$\&E\ m$
	\mathcal{B}	$\&E\ m$

When you have a conjunction on some line of a proof, you can use $\&E$ to derive either of the conjuncts. The $\&E$ rule requires only one sentence, so we write one line number as the justification for applying it.

Even with just these two rules, we can provide some proofs. Consider this argument.

$$[(A \vee B) \rightarrow (C \vee D)] \& [(E \vee F) \rightarrow (G \vee H)]$$

$$\therefore [(E \vee F) \rightarrow (G \vee H)] \& [(A \vee B) \rightarrow (C \vee D)]$$

The main logical operator in both the premise and conclusion is conjunction. Since conjunction is symmetric, the argument is obviously valid. In order to provide a proof, we begin by writing down the premise. After the premises, we draw a horizontal line— everything below this line must be justified by a rule of proof. So the beginning of the proof looks like this:

$$1 \quad \underline{[(A \vee B) \rightarrow (C \vee D)] \& [(E \vee F) \rightarrow (G \vee H)]}$$

From the premise, we can get each of the conjuncts by $\&E$. The proof now looks like this:

1	$[(A \vee B) \rightarrow (C \vee D)] \& [(E \vee F) \rightarrow (G \vee H)]$	
2	$[(A \vee B) \rightarrow (C \vee D)]$	$\&E\ 1$
3	$[(E \vee F) \rightarrow (G \vee H)]$	$\&E\ 1$

The rule $\&I$ requires that we have each of the conjuncts available somewhere in the proof. They can be separated from one another, and they can appear in any order. So by applying the $\&I$ rule to lines 3 and 2, we arrive at the desired conclusion. The finished proof looks like this:

1	$[(A \vee B) \rightarrow (C \vee D)] \& [(E \vee F) \rightarrow (G \vee H)]$	
2	$[(A \vee B) \rightarrow (C \vee D)]$	&E 1
3	$[(E \vee F) \rightarrow (G \vee H)]$	&E 1
4	$[(E \vee F) \rightarrow (G \vee H)] \& [(A \vee B) \rightarrow (C \vee D)]$	&I 3, 2

This proof is trivial, but it shows how we can use rules of proof together to demonstrate the validity of an argument form. Also: Using a truth table to show that this argument is valid would have required a staggering 256 lines, since there are eight sentence letters in the argument.

Disjunction

If M were true, then $M \vee N$ would also be true. So the disjunction introduction rule (\vee I) allows us to derive a disjunction if we have one of the two disjuncts:

m	\mathcal{A}	
	$\mathcal{A} \vee \mathcal{B}$	\vee I m
	$\mathcal{B} \vee \mathcal{A}$	\vee I m

Notice that \mathcal{B} can be *any* sentence whatsoever. So the following is legitimate proof:

1	M	
2	$M \vee ((A \leftrightarrow B) \rightarrow (C \& D)) \leftrightarrow [E \& F]$	\vee I 1

It may seem odd that just by knowing M we can derive a conclusion that includes sentences like A , B , and the rest—sentences that have nothing to do with M . Yet the conclusion follows immediately by \vee I. This is as it should be: The truth conditions for the disjunction mean that, if \mathcal{A} is true, then $\mathcal{A} \vee \mathcal{B}$ is true regardless of what \mathcal{B} is. So the conclusion could not be false if the premise were true; the argument is valid.

Now consider the disjunction elimination rule. What can you conclude from $M \vee N$? You cannot conclude M . It might be M 's truth that makes $M \vee N$ true, as in the example above, but it might not. From $M \vee N$ alone, you cannot conclude anything about either M or N specifically. If you also knew that N was false, however, then you would be able to conclude M .

This is just disjunctive syllogism, it will be the disjunction elimination rule (\vee E).

$$\begin{array}{l|l}
 m & \mathcal{A} \vee \mathcal{B} \\
 n & \neg \mathcal{B} \\
 & \mathcal{A} \quad \vee\text{E } m, n
 \end{array}
 \qquad
 \begin{array}{l|l}
 m & \mathcal{A} \vee \mathcal{B} \\
 n & \neg \mathcal{A} \\
 & \mathcal{B} \quad \vee\text{E } m, n
 \end{array}$$

Conditional

Consider this argument:

$$\begin{array}{l}
 R \vee F \\
 \therefore \neg R \rightarrow F
 \end{array}$$

The argument is certainly a valid one. What should the conditional introduction rule be, such that we can draw this conclusion?

We begin the proof by writing down the premise of the argument and drawing a horizontal line, like this:

$$1 \quad \underline{R \vee F}$$

If we had $\neg R$ as a further premise, we could derive F by the $\vee\text{E}$ rule. We do not have $\neg R$ as a premise of this argument, nor can we derive it directly from the premise we do have— so we cannot simply prove F . What we will do instead is start a *subproof*, a proof within the main proof. When we start a subproof, we draw another vertical line to indicate that we are no longer in the main proof. Then we write in an assumption for the subproof. This can be anything we want. Here, it will be helpful to assume $\neg R$. Our proof now looks like this:

$$\begin{array}{l|l}
 1 & \underline{R \vee F} \\
 2 & \quad \underline{\neg R}
 \end{array}$$

It is important to notice that we are not claiming to have proven $\neg R$. We do not need to write in any justification for the assumption line of a subproof. You can think of the subproof as posing the question: What could we show *if* $\neg R$ were true? For one thing, we can derive F . So we do:

$$\begin{array}{l|l}
 1 & \underline{R \vee F} \\
 2 & \quad \underline{\neg R} \\
 3 & \quad \quad \underline{F} \quad \vee\text{E } 1, 2
 \end{array}$$

This has shown that *if* we had $\neg R$ as a premise, *then* we could prove F . In effect, we have proven $\neg R \rightarrow F$. So the conditional introduction rule (\rightarrow I) will allow us to close the subproof and derive $\neg R \rightarrow F$ in the main proof. Our final proof looks like this:

1	$R \vee F$			
2	<table style="border-collapse: collapse; margin-left: 20px;"> <tr> <td style="border-left: 1px solid black; padding-left: 10px;">$\neg R$</td> <td></td> </tr> </table>	$\neg R$		
$\neg R$				
3	<table style="border-collapse: collapse; margin-left: 20px;"> <tr> <td style="border-left: 1px solid black; padding-left: 10px; border-bottom: 1px solid black;">F</td> <td style="padding-left: 10px;">\veeE 1, 2</td> </tr> </table>	F	\vee E 1, 2	
F	\vee E 1, 2			
4	$\neg R \rightarrow F$	\rightarrow I 2-3		

Notice that the justification for applying the \rightarrow I rule is the entire subproof. Usually that will be more than just two lines.

It may seem as if the ability to assume anything at all in a subproof would lead to chaos: Does it allow you to prove any conclusion from any premises? The answer is no, it does not. Consider this proof:

1	\mathcal{A}			
2	<table style="border-collapse: collapse; margin-left: 20px;"> <tr> <td style="border-left: 1px solid black; padding-left: 10px;">\mathcal{B}</td> <td></td> </tr> </table>	\mathcal{B}		
\mathcal{B}				
3	<table style="border-collapse: collapse; margin-left: 20px;"> <tr> <td style="border-left: 1px solid black; padding-left: 10px; border-bottom: 1px solid black;">\mathcal{B}</td> <td style="padding-left: 10px;">R 2</td> </tr> </table>	\mathcal{B}	R 2	
\mathcal{B}	R 2			

It may seem as if this is a proof that you can derive any conclusions \mathcal{B} from any premise \mathcal{A} . When the vertical line for the subproof ends, the subproof is *closed*. In order to complete a proof, you must close all of the subproofs. And you cannot close the subproof and use the R rule again on line 4 to derive \mathcal{B} in the main proof. Once you close a subproof, you cannot refer back to individual lines inside it.

Closing a subproof is called *discharging* the assumptions of that subproof. So we can put the point this way: You cannot complete a proof until you have discharged all of the assumptions besides the original premises of the argument.

Of course, it is legitimate to do this:

1	\mathcal{A}			
2	<table style="border-collapse: collapse; margin-left: 20px;"> <tr> <td style="border-left: 1px solid black; padding-left: 10px;">\mathcal{B}</td> <td></td> </tr> </table>	\mathcal{B}		
\mathcal{B}				
3	<table style="border-collapse: collapse; margin-left: 20px;"> <tr> <td style="border-left: 1px solid black; padding-left: 10px; border-bottom: 1px solid black;">\mathcal{B}</td> <td style="padding-left: 10px;">R 2</td> </tr> </table>	\mathcal{B}	R 2	
\mathcal{B}	R 2			
4	$\mathcal{B} \rightarrow \mathcal{B}$	\rightarrow I 2-3		

This should not seem so strange, though. Since $\mathcal{B} \rightarrow \mathcal{B}$ is a tautology, it follows validly from any premises.

Put in a general form, the \rightarrow I rule looks like this:

$$\begin{array}{l|l} m & \mathcal{A} \quad \text{want } \mathcal{B} \\ n & \mathcal{B} \\ \hline & \mathcal{A} \rightarrow \mathcal{B} \quad \rightarrow\text{I } m-n \end{array}$$

When we introduce a subproof, we typically write what we want to derive in the column. This is just so that we do not forget why we started the subproof if it goes on for five or ten lines. There is no ‘want’ rule. It is a note to ourselves and not formally part of the proof.

Although it is always permissible to open a subproof with any assumption you please, there is some strategy involved in picking a useful assumption. Starting a subproof with an arbitrary, wacky assumption would just waste lines of the proof. In order to derive a conditional by the \rightarrow I, for instance, you must assume the antecedent of the conditional in a subproof.

Now consider the conditional elimination rule. Nothing follows from $M \rightarrow N$ alone, but if we have both $M \rightarrow N$ and M , then we can conclude N . This rule, modus ponens, will be the conditional elimination rule (\rightarrow E).

$$\begin{array}{l|l} m & \mathcal{A} \rightarrow \mathcal{B} \\ n & \mathcal{A} \\ \hline & \mathcal{B} \quad \rightarrow\text{E } m, n \end{array}$$

Now that we have rules for the conditional, consider this argument:

$$\begin{array}{l} P \rightarrow Q \\ Q \rightarrow R \\ \therefore P \rightarrow R \end{array}$$

We begin the proof by writing the two premises as assumptions. Since the main logical operator in the conclusion is a conditional, we can expect to use the \rightarrow I rule. For that, we need a subproof— so we write in the antecedent of the conditional as assumption of a subproof:

1	$P \rightarrow Q$
2	$Q \rightarrow R$
3	P

We made P available by assuming it in a subproof, allowing us to use \rightarrow E on the first premise. This gives us Q , which allows us to use \rightarrow E on the second premise. Having derived R , we close the subproof. By assuming P we were able to prove R , so we apply the \rightarrow I rule and finish the proof.

1	$P \rightarrow Q$	
2	$Q \rightarrow R$	
3	P	want R
4	Q	\rightarrow E 1, 3
5	R	\rightarrow E 2, 4
6	$P \rightarrow R$	\rightarrow I 3–5

Biconditional

The rules for the biconditional will be like double-barreled versions of the rules for the conditional.

In order to derive $W \leftrightarrow X$, for instance, you must be able to prove X by assuming W and prove W by assuming X . The biconditional introduction rule (\leftrightarrow I) requires two subproofs. The subproofs can come in any order, and the second subproof does not need to come immediately after the first— but schematically, the rule works like this:

m	\mathcal{A}	want \mathcal{B}
n	\mathcal{B}	
p	\mathcal{B}	want \mathcal{A}
q	\mathcal{A}	
	$\mathcal{A} \leftrightarrow \mathcal{B}$	\leftrightarrow I m – n , p – q

The biconditional elimination rule (\leftrightarrow E) lets you do a bit more than the conditional rule. If you have the left-hand subsentence of the biconditional, you can derive the right-hand subsentence. If you have the right-hand subsentence, you

can derive the left-hand subsentence. This is the rule:

$$\begin{array}{l|l}
 m & \mathcal{A} \rightarrow \mathcal{B} \\
 n & \mathcal{A} \\
 & \mathcal{B} \\
 \hline
 & \leftrightarrow E \ m, n
 \end{array}
 \qquad
 \begin{array}{l|l}
 m & \mathcal{A} \rightarrow \mathcal{B} \\
 n & \mathcal{B} \\
 & \mathcal{A} \\
 \hline
 & \leftrightarrow E \ m, n
 \end{array}$$

Negation

Here is a simple mathematical argument in English:

Assume there is some greatest natural number. Call it A .

That number plus one is also a natural number.

Obviously, $A + 1 > A$.

So there is a natural number greater than A .

This is impossible, since A is assumed to be the greatest natural number.

\therefore There is no greatest natural number.

This argument form is traditionally called a *reductio*. Its full Latin name is *reductio ad absurdum*, which means ‘reduction to absurdity.’ In a *reductio*, we assume something for the sake of argument— for example, that there is a greatest natural number. Then we show that the assumption leads to two contradictory sentences— for example, that A is the greatest natural number and that it is not. In this way, we show that the original assumption must have been false.

The basic rules for negation will allow for arguments like this. If we assume something and show that it leads to contradictory sentences, then we have proven the negation of the assumption. This is the negation introduction (\neg I) rule:

$$\begin{array}{l|l}
 m & \mathcal{A} \\
 n & \mathcal{B} \\
 n + 1 & \neg \mathcal{B} \\
 n + 2 & \neg \mathcal{A} \\
 \hline
 & \neg I \ m-n+1
 \end{array}
 \quad \text{for reductio}$$

For the rule to apply, the last two lines of the subproof must be an explicit contradiction: some sentence followed on the next line by its negation. We write ‘for reductio’ as a note to ourselves, a reminder of why we started the subproof. It is not formally part of the proof, and you can leave it out if you find it distracting.

To see how the rule works, suppose we want to prove the law of non-contradiction: $\neg(G \& \neg G)$. We can prove this without any premises by immediately starting a subproof. We want to apply $\neg I$ to the subproof, so we assume $(G \& \neg G)$. We then get an explicit contradiction by $\&E$. The proof looks like this:

1	$G \& \neg G$	for reductio
2	G	$\rightarrow E$ 1
3	$\neg G$	$\rightarrow E$ 1
4	$\neg(G \& \neg G)$	$\neg I$ 1-3

The $\neg E$ rule will work in much the same way. If we assume $\neg\mathcal{A}$ and show that it leads to a contradiction, we have effectively proven \mathcal{A} . So the rule looks like this:

m	$\neg\mathcal{A}$	for reductio
n	\mathcal{B}	
$n+1$	$\neg\mathcal{B}$	
$n+2$	\mathcal{A}	$\neg E$ $m-n+1$

6.2 Derived rules

The rules of the natural deduction system are meant to be systematic. There is an introduction and an elimination rule for each logical operator, but why these basic rules rather some others? Many natural deduction systems have a disjunction elimination rule that works like this:

m	$\mathcal{A} \vee \mathcal{B}$	
n	$\mathcal{A} \rightarrow \mathcal{C}$	
o	$\mathcal{B} \rightarrow \mathcal{C}$	
	\mathcal{C}	$\vee^* m, n, o$

It might seem as if there will be some proofs that we cannot do with our proof system, because we do not have this rule. Yet this is not the case. If you can do a proof with this rule, you can do a proof with the basic rules of the natural deduction system. Consider this proof:

1	$\mathcal{A} \vee \mathcal{B}$					
2	$\mathcal{A} \rightarrow \mathcal{C}$					
3	$\mathcal{B} \rightarrow \mathcal{C}$	want \mathcal{C}				
4	<table style="border-collapse: collapse; margin-left: 1em;"> <tr> <td style="border-left: 1px solid black; padding-left: 5px;">$\neg \mathcal{C}$</td> <td>for reductio</td> </tr> </table>	$\neg \mathcal{C}$	for reductio			
$\neg \mathcal{C}$	for reductio					
5	<table style="border-collapse: collapse; margin-left: 1em;"> <tr> <td style="border-left: 1px solid black; padding-left: 5px;"> <table style="border-collapse: collapse; margin-left: 1em;"> <tr> <td style="border-left: 1px solid black; padding-left: 5px;">\mathcal{A}</td> <td>for reductio</td> </tr> </table> </td> <td></td> </tr> </table>	<table style="border-collapse: collapse; margin-left: 1em;"> <tr> <td style="border-left: 1px solid black; padding-left: 5px;">\mathcal{A}</td> <td>for reductio</td> </tr> </table>	\mathcal{A}	for reductio		
<table style="border-collapse: collapse; margin-left: 1em;"> <tr> <td style="border-left: 1px solid black; padding-left: 5px;">\mathcal{A}</td> <td>for reductio</td> </tr> </table>	\mathcal{A}	for reductio				
\mathcal{A}	for reductio					
6	<table style="border-collapse: collapse; margin-left: 1em;"> <tr> <td style="border-left: 1px solid black; padding-left: 5px;">\mathcal{C}</td> <td>\rightarrowI 2, 5</td> </tr> </table>	\mathcal{C}	\rightarrow I 2, 5			
\mathcal{C}	\rightarrow I 2, 5					
7	<table style="border-collapse: collapse; margin-left: 1em;"> <tr> <td style="border-left: 1px solid black; padding-left: 5px;">$\neg \mathcal{C}$</td> <td>R 4</td> </tr> </table>	$\neg \mathcal{C}$	R 4			
$\neg \mathcal{C}$	R 4					
8	$\neg \mathcal{A}$	\neg I 5–7				
9	<table style="border-collapse: collapse; margin-left: 1em;"> <tr> <td style="border-left: 1px solid black; padding-left: 5px;">\mathcal{B}</td> <td>for reductio</td> </tr> </table>	\mathcal{B}	for reductio			
\mathcal{B}	for reductio					
10	<table style="border-collapse: collapse; margin-left: 1em;"> <tr> <td style="border-left: 1px solid black; padding-left: 5px;">\mathcal{C}</td> <td>\rightarrowI 3, 9</td> </tr> </table>	\mathcal{C}	\rightarrow I 3, 9			
\mathcal{C}	\rightarrow I 3, 9					
11	<table style="border-collapse: collapse; margin-left: 1em;"> <tr> <td style="border-left: 1px solid black; padding-left: 5px;">$\neg \mathcal{C}$</td> <td>R 4</td> </tr> </table>	$\neg \mathcal{C}$	R 4			
$\neg \mathcal{C}$	R 4					
12	\mathcal{B}	\vee E 1, 8				
13	$\neg \mathcal{B}$	\neg I 9–11				
14	\mathcal{C}	\neg E 4–13				

\mathcal{A} , \mathcal{B} , and \mathcal{C} are meta-variables. They are not symbols of SL, but stand-ins for arbitrary sentences of SL. So this is not, strictly speaking, a proof in SL. It is more like a recipe. It provides a pattern that can prove anything that the \vee^* rule can prove, using only the basic rules of SL. This means that the \vee^* is not really necessary. Adding it to the list of basic rules would not allow us to derive anything that we could not derive without it.

Nevertheless, the \vee^* rule would be convenient. It would allow us to do in one line what requires eleven lines and several nested subproofs with the basic rules. So we will add \vee^* to the proof system as a derived rule.

A DERIVED RULE is a rule of proof that does not make any new proofs possible. Anything that can be proven with a derived rule can be proven without it. You can think of a short proof using a derived rule as shorthand for a longer proof that uses only the basic rules. Anytime you use the \vee^* rule, you could always take ten extra lines and prove the same thing without it.

For the sake of convenience, we will add several other derived rules. One is *modus tollens* (MT).

$$\begin{array}{l|l}
 m & \mathcal{A} \rightarrow \mathcal{B} \\
 n & \neg \mathcal{B} \\
 & \neg \mathcal{A} \quad \text{MT } m, n
 \end{array}$$

We leave the proof of this rule as an exercise. Note that if we had already proven the MT rule, then the proof of the \forall^* rule could have been done in only five lines.

We also add hypothetical syllogism (HS) as a derived rule. We have already given a proof of it on p. 112.

$$\begin{array}{l|l}
 m & \mathcal{A} \rightarrow \mathcal{B} \\
 n & \mathcal{B} \rightarrow \mathcal{C} \\
 & \mathcal{A} \rightarrow \mathcal{C} \quad \text{HS } m, n
 \end{array}$$

6.3 Rules of replacement

Consider how you would prove this argument: $F \rightarrow (G \& H), \therefore F \rightarrow G$

Perhaps it is tempting to write down the premise and apply the $\&E$ rule to the conjunction $(G \& H)$. This is impermissible, however, because the basic rules of proof can only be applied to whole sentences. We need to get $(G \& H)$ on a line by itself. We can prove the argument in this way:

$$\begin{array}{l|l}
 1 & F \rightarrow (G \& H) \\
 \hline
 2 & \begin{array}{l|l} & F \\ \hline & G \& H \end{array} & \text{want } G \\
 3 & & \rightarrow E 1, 2 \\
 4 & \begin{array}{l|l} & G \\ \hline & G \end{array} & \& E 3 \\
 5 & F \rightarrow G & \rightarrow I 2-4
 \end{array}$$

We will now introduce some derived rules that may be applied to part of a sentence. These are called RULES OF REPLACEMENT, because they can be used to replace part of a sentence with a logically equivalent expression. One simple rule of replacement is commutivity (abbreviated Comm), which says that we can swap the order of conjuncts in a conjunction or the order of disjuncts in a disjunction. We define the rule this way:

$$\begin{aligned} (\mathcal{A} \& \mathcal{B}) &\iff (\mathcal{B} \& \mathcal{A}) \\ (\mathcal{A} \vee \mathcal{B}) &\iff (\mathcal{B} \vee \mathcal{A}) \quad \text{Comm} \end{aligned}$$

The bold arrow means that you can take a subformula on one side of the arrow and replace it with the subformula on the other side. The arrow is double-headed because rules of replacement work in both directions.

Consider this argument: $(M \vee P) \rightarrow (P \& M), \therefore (P \vee M) \rightarrow (M \& P)$

It is possible to give a proof of this using only the basic rules, but it will be long and inconvenient. With the Comm rule, we can provide a proof easily:

$$\begin{array}{l|l} 1 & (M \vee P) \rightarrow (P \& M) \\ \hline 2 & (P \vee M) \rightarrow (P \& M) \quad \text{Comm 1} \\ 3 & (P \vee M) \rightarrow (M \& P) \quad \text{Comm 2} \end{array}$$

Another rule of replacement is double negation (DN). With the DN rule, you can remove or insert a pair of negations anywhere in a sentence. This is the rule:

$$\neg\neg\mathcal{A} \iff \mathcal{A} \quad \text{DN}$$

Two more replacement rules are called De Morgan's Laws, named for the 19th-century British logician August De Morgan. (Although De Morgan did discover these laws, he was not the first to do so.) The rules capture useful the relations between negation, conjunction, and disjunction. Here are the rules, which we abbreviate DeM:

$$\begin{aligned} \neg(\mathcal{A} \vee \mathcal{B}) &\iff (\neg\mathcal{A} \& \neg\mathcal{B}) \\ \neg(\mathcal{A} \& \mathcal{B}) &\iff (\neg\mathcal{A} \vee \neg\mathcal{B}) \quad \text{DeM} \end{aligned}$$

Another replacement rule captures the relation between conditionals and disjunctions. This rule is sometimes called implication; we abbreviate the rule $\vee\rightarrow$, because it allows transformation between disjunctions and conditionals. It takes two forms:

$$(\mathcal{A} \rightarrow \mathcal{B}) \iff (\neg\mathcal{A} \vee \mathcal{B}) \quad \vee\rightarrow$$

Now consider this argument: $\neg(P \rightarrow Q), \therefore P \& \neg Q$

As always, we could prove this argument using only the basic rules. With rules of replacement, though, the proof is much simpler:

1	$\neg(P \rightarrow Q)$	
2	$\neg(\neg P \vee Q)$	$\vee \rightarrow 1$
3	$\neg\neg P \ \& \ \neg Q$	DeM 2
4	$P \ \& \ \neg Q$	DN 3

6.4 Rules for quantifiers

For proofs in QL, we use all of the basic rules of SL plus four new basic rules: both introduction and elimination rules for each of the quantifiers.

Since all of the derived rules of SL are derived from the basic rules, they will also hold in QL. We will add another derived rule, a replacement rule called quantifier negation.

Universal elimination

If you have $\forall xAx$, it is legitimate to infer that anything is an A . You can infer Aa, Ab, Az, Ad_3 — in short, you can infer Ac for any constant c . This is the general form of the universal elimination rule ($\forall E$):

m	$\forall \chi \mathcal{A}$	
	$\mathcal{A}[c \chi]$	$\forall E \ m$

$\mathcal{A}[c|\chi]$ is a substitution instance of $\forall \chi \mathcal{A}$. The symbols for a substitution instance are not symbols of QL, so you cannot write them in a proof. Instead, you write the substituted sentence with the constant c replacing all occurrences of the variable χ in \mathcal{A} . For example:

1	$\forall x(Mx \rightarrow Rxd)$	
2	$Ma \rightarrow Rad$	$\forall E \ 1$
3	$Md \rightarrow Rdd$	$\forall E \ 1$

Existential introduction

When is it legitimate to infer $\exists xAx$? If you know that something is an A — for instance, if you have Aa available in the proof.

This is the existential introduction rule ($\exists\text{I}$):

$$m \quad \left| \begin{array}{l} \mathcal{A} \\ \exists \chi \mathcal{A}[\chi||c] \end{array} \right. \quad \exists\text{I } m$$

It is important to notice that $\mathcal{A}[\chi||c]$ is not the same as a substitution instance. We write it with two bars to show that the variable χ does not need to replace all occurrences of the constant c . You can decide which occurrences to replace and which to leave in place. For example:

$$\begin{array}{l|l} 1 & Ma \rightarrow Rad \\ \hline 2 & \exists x(Ma \rightarrow Rax) \quad \exists\text{I } 1 \\ 3 & \exists x(Mx \rightarrow Rxd) \quad \exists\text{I } 1 \\ 4 & \exists x(Mx \rightarrow Rad) \quad \exists\text{I } 1 \\ 5 & \exists y \exists x(Mx \rightarrow Ryd) \quad \exists\text{I } 4 \\ 6 & \exists z \exists y \exists x(Mx \rightarrow Ryz) \quad \exists\text{I } 5 \end{array}$$

Universal introduction

A universal claim like $\forall xPx$ would be proven if every substitution instance of it had been proven, if every sentence Pa, Pb, \dots were available in a proof. Alas, there is no hope of proving *every* substitution instance. That would require proving $Pa, Pb, \dots, Pj_2, \dots, Ps_7, \dots$, and so on to infinity. There are infinitely many constants in QL, and so this process would never come to an end.

Consider a simple argument: $\forall xMx, \therefore \forall yMy$

It makes no difference to the meaning of the sentence whether we use the variable x or the variable y , so this argument is obviously valid. Suppose we begin in this way:

$$\begin{array}{l|l} 1 & \forall xMx \quad \text{want } \forall yMy \\ \hline 2 & Ma \quad \& \text{E } 1 \end{array}$$

We have derived Ma . Nothing stops us from using the same justification to derive $Mb, \dots, Mj_2, \dots, Ms_7, \dots$, and so on until we run out of space or patience. We have effectively shown the way to prove Mc for any constant c . From this, $\forall xMx$ follows.

1	$\forall x Mx$	
2	Ma	& E 1
3	$\forall y My$	& I 2

It is important here that a was just some arbitrary constant. We had not made any special assumptions about it. If Ma were a premise of the argument, then this would not show anything about *all* y . For example:

1	$\forall x Rxa$	
2	Raa	& E 1
3	$\forall y Ryy$	not allowed!

This is the schematic form of the universal introduction rule (\forall I):

m	\mathcal{A}	
	$\forall \chi \mathcal{A}[\chi c]^*$	\forall I m

* c must not occur in any undischarged assumptions.

Note that we can do this for any constant that does not occur in an undischarged assumption and for any variable.

Note also that the constant may not occur in any *undischarged* assumption, but it may occur as the assumption of a subproof that we have already closed. For example, we can prove $\forall z(Dz \rightarrow Dz)$ without any premises.

1	Df	want Df
2	Df	R 1
3	$Df \rightarrow Df$	\rightarrow I 1-2
4	$\forall z(Dz \rightarrow Dz)$	&I 3

Existential elimination

A sentence with an existential quantifier tells us that there is *some* member of the UD that satisfies a formula. For example, $\exists x Sx$ tells us (roughly) that there is at least one S . It does not tell us *which* member of the UD satisfies S , however. We cannot immediately conclude Sa , Sf_{23} , or any other substitution instance of the sentence. What can we do?

Suppose that we knew both $\exists xSx$ and $\forall x(Sx \rightarrow Tx)$. We could reason in this way:

Since $\exists xSx$, there is something that is an S . We do not know which constants refer to this thing, if any do, so call this thing Ω . From $\forall x(Sx \rightarrow Tx)$, it follows that if Ω is an S , then it is a T . Therefore Ω is a T . Because Ω is a T , we know that $\exists xTx$.

In this paragraph, we introduced a name for the thing that is an S . We called it Ω , so that we could reason about it and derive some consequences from there being an S . Since Ω is just a bogus name introduced for the purpose of the proof and not a genuine constant, we could not mention it in the conclusion. Yet we could derive a sentence that does not mention Ω ; namely, $\exists xTx$. This sentence does follow from the two premises.

We want the existential elimination rule to work in a similar way. Yet since Greek letters like Ω are not symbols of QL, we cannot use them in formal proofs. Instead, we will use constants of QL which do not otherwise appear in the proof. A constant that is used to stand in for whatever it is that satisfies an existential claim is called a PROXY. Reasoning with the proxy must all occur inside a subproof, and the proxy cannot be a constant that is doing work elsewhere in the proof.

This is the schematic form of the existential elimination rule ($\exists E$):

$$\begin{array}{l|l}
 m & \exists x\mathcal{A} \\
 n & \left| \begin{array}{l} \mathcal{A}[c|\chi]^* \\ \hline \mathcal{B} \end{array} \right. \\
 p & \left| \mathcal{B} \right. \\
 & \mathcal{B} \qquad \exists E \ m, \ n-p
 \end{array}$$

* The constant c must not appear in $\exists x\mathcal{A}$, in \mathcal{B} , or in any undischarged assumption.

With this rule, we can give a formal proof that $\exists xSx$ and $\forall x(Sx \rightarrow Tx)$ together entail $\exists xTx$. The structure of the proof is effectively the same as the English-language argument with which we began, except that the subproof uses the constant ' a ' rather than the bogus name Ω .

1	$\exists xSx$	
2	$\forall x(Sx \rightarrow Tx)$	want $\exists xTx$
3	Sa	
4	$Sa \rightarrow Ta$	$\forall E$ 2
5	Ta	$\rightarrow E$ 3, 4
6	$\exists xTx$	$\exists I$ 5
7	$\exists xTx$	$\exists E$ 1, 3-6

Quantifier negation

When translating from English to QL, we noted that $\neg\exists x\neg\mathcal{A}$ is logically equivalent to $\forall x\mathcal{A}$. In QL, they are provably equivalent. We can prove one half of the equivalence with a rather gruesome proof:

1	$\forall xAx$	want $\neg\exists x\neg Ax$
2	$\exists\neg Ax$	for reductio
3	$\neg Ac$	for $\exists E$
4	$\forall xAx$	for reductio
5	Ac	$\forall E$ 1
6	$\neg Ac$	R 3
7	$\neg\forall xAx$	$\neg I$ 4-6
8	$\forall xAx$	R 1
9	$\neg\forall xAx$	$\exists E$ 3-7
10	$\neg\exists\neg Ax$	$\neg I$ 2-9

In order to show that the two sentences are genuinely equivalent, we need a second proof that assumes $\neg\exists x\neg\mathcal{A}$ and derives $\forall x\mathcal{A}$. We leave that proof as an exercise for the reader.

It will often be useful to translate between quantifiers by adding or subtracting negations in this way, so we add two derived rules for this purpose. These rules are called quantifier negation (QN):

$$\begin{aligned} \neg \forall x \mathcal{A} &\iff \exists x \neg \mathcal{A} \\ \neg \exists x \mathcal{A} &\iff \forall x \neg \mathcal{A} \quad \text{QN} \end{aligned}$$

Since QN is a replacement rule, it can be used on whole sentences or on subformulae.

6.5 Rules for identity

The identity predicate is not part of QL, but we add it when we need to symbolize certain sentences. For proofs involving identity, we add two rules of proof.

Suppose you know that many things that are true of a are also true of b . For example: $Aa \ \& \ Ab$, $Ba \ \& \ Bb$, $\neg Ca \ \& \ \neg Cb$, $Da \ \& \ Db$, $\neg Ea \ \& \ \neg Eb$, and so on. This would not be enough to justify the conclusion $a = b$. (See p. 97.) In general, there are no sentences that do not already contain the identity predicate that could justify the conclusion $a = b$. This means that the identity introduction rule will not justify $a = b$ or any other identity claim containing two different constants.

However, it is always true that $a = a$. In general, no premises are required in order to conclude that something is identical to itself. So this will be the identity introduction rule, abbreviated =I:

$$\left| \begin{array}{l} c = c \end{array} \right. \quad =I$$

Notice that the =I rule does not require referring to any prior lines of the proof. For any constant c , you can write $c = c$ on any point with only the =I rule as justification.

If you have shown that $a = b$, then anything that is true of a must also be true of b . For any sentence with a in it, you can replace some or all of the occurrences of a with b and produce an equivalent sentence. For example, if you already know Raa , then you are justified in concluding Rab , Rba , Rbb . Recall that $\mathcal{A}[a||b]$ is the sentence produced by replacing a in \mathcal{A} with b . This is not the same as a substitution instance, because b may replace some or all occurrences of a . The identity elimination rule (=E) justifies replacing terms with other terms that are identical to it:

$$\left| \begin{array}{l} m \quad a = b \\ n \quad \mathcal{A} \\ \mathcal{A}[a||b] \quad =E \ m, \ n \\ \mathcal{A}[b||a] \quad =E \ m, \ n \end{array} \right.$$

To see the rules in action, consider this proof:

1	$\forall x \forall y x = y$	
2	$\exists x Bx$	
3	$\forall x (Bx \rightarrow \neg Cx)$	want $\neg \exists x Cx$
4	Be	
5	$e = f$	$\forall E$ 1
6	Bf	$=E$ 5, 4
7	$Bf \rightarrow \neg Cf$	$\forall E$ 3
8	$\neg Cf$	$\rightarrow E$ 7, 6
9	$\neg Cf$	$\exists E$ 2, 4-8
10	$\forall x \neg Cx$	$\forall I$ 9
11	$\neg \exists x Cx$	QN 10

6.6 Proof strategy

There is no simple recipe for proofs, and there is no substitute for practice. Here, though, are some rules of thumb and strategies to keep in mind.

Work backwards from what you want. The ultimate goal is to derive the conclusion. Look at the conclusion and ask what the introduction rule is for its main logical operator. This gives you an idea of what should happen *just before* the last line of the proof. Then you can treat this line as if it were your goal. Ask what you could do to derive this new goal.

For example: If your conclusion is a conditional $\mathcal{A} \rightarrow \mathcal{B}$, plan to use the $\rightarrow I$ rule. This requires starting a subproof in which you assume \mathcal{A} . In the subproof, you want to derive \mathcal{B} .

Work forwards from what you have. When you are starting a proof, look at the premises; later, look at the sentences that you have derived so far. Think about the elimination rules for the main operators these sentences. These will tell you what your options are.

For example: If you have $\forall x \mathcal{A}$, think about instantiating it for any constant that might be helpful. If you have $\exists x \mathcal{A}$ and intend to use the $\exists E$ rule, then you

should assume $\mathcal{A}[c|x]$ for some c that is not in use and then derive a conclusion that does not contain c .

For a short proof, you might be able to eliminate the premises and introduce the conclusion. A long proof is formally just a number of short proofs linked together, so you can fill the gap by alternately working back from the conclusion and forward from the premises.

Change what you are looking at. Replacement rules can often make your life easier. If a proof seems impossible, try out some different substitutions.

For example: It is often difficult to prove a disjunction using the basic rules. If you want to show $\mathcal{A} \vee \mathcal{B}$, it is often easier to show $\neg\mathcal{A} \rightarrow \mathcal{B}$ and use the $\vee\rightarrow$ rule.

Showing $\neg\exists x\mathcal{A}$ can also be hard, and it is often easier to show $\forall x\neg\mathcal{A}$ and use the QN rule.

Some replacement rules should become second nature. If you see a negated disjunction, for instance, you should immediately think of DeMorgan's rule.

Do not forget indirect proof. If you cannot find a way to show something directly, try assuming its negation.

Remember that most proofs can be done either indirectly or directly. One way might be easier— or perhaps one sparks your imagination more than the other— but either one is formally legitimate.

Repeat as necessary. Once you have decided how you might be able to get to the conclusion, ask what you might be able to do with the premises. Then consider the target sentences again and ask how you might reach them.

Persist. Try different things. If one approach fails, then try something else.

6.7 Proof-theoretic concepts

We will use the symbol ' \vdash ' to indicate that a proof is possible. This symbol is called the *turnstile*. Sometimes it is called a *single turnstile*, to underscore the fact that this is not the double turnstile symbol (\models) that we used to represent semantic entailment in ch. 5.

When we write $\{\mathcal{A}_1, \mathcal{A}_2, \dots\} \vdash \mathcal{B}$, this means that it is possible to give a proof of \mathcal{B} with $\mathcal{A}_1, \mathcal{A}_2, \dots$ as premises. With just one premise, we leave out the curly braces, so $\mathcal{A} \vdash \mathcal{B}$ means that there is a proof of \mathcal{B} with \mathcal{A} as a premise. Naturally, $\vdash \mathcal{C}$ means that there is a proof of \mathcal{C} that has no premises.

Often, logical proofs are called *derivations*. So $\mathcal{A} \vdash \mathcal{B}$ can be read as ‘ \mathcal{B} is derivable from \mathcal{A} .’

A THEOREM is a sentence that is derivable without any premises; i.e., \mathcal{T} is a theorem if and only if $\vdash \mathcal{T}$.

It is not too hard to show that something is a theorem—you just have to give a proof of it. How could you show that something is *not* a theorem? If its negation is a theorem, then you could provide a proof. For example, it is easy to prove $\neg(Pa \ \& \ \neg Pa)$, which shows that $(Pa \ \& \ \neg Pa)$ cannot be a theorem. For a sentence that is neither a theorem nor the negation of a theorem, however, there is no easy way to show this. You would have to demonstrate not just that certain proof strategies fail, but that no proof is possible. Even if you fail in trying to prove a sentence in a thousand different ways, perhaps the proof is just too long and complex for you to make out.

Two sentences \mathcal{A} and \mathcal{B} are PROVABLY EQUIVALENT if and only if each can be derived from the other; i.e., $\mathcal{A} \vdash \mathcal{B}$ and $\mathcal{B} \vdash \mathcal{A}$

It is relatively easy to show that two sentences are provably equivalent—it just requires a pair of proofs. Showing that sentences are *not* provably equivalent would be much harder. It would be just as hard as showing that a sentence is not a theorem. (In fact, these problems are interchangeable. Can you think of a sentence that would be a theorem if and only if \mathcal{A} and \mathcal{B} were provably equivalent?)

The set of sentences $\{\mathcal{A}_1, \mathcal{A}_2, \dots\}$ is PROVABLY INCONSISTENT if and only if a contradiction is derivable from it; i.e., for some sentence \mathcal{B} , $\{\mathcal{A}_1, \mathcal{A}_2, \dots\} \vdash \mathcal{B}$ and $\{\mathcal{A}_1, \mathcal{A}_2, \dots\} \vdash \neg \mathcal{B}$.

It is easy to show that the set is provably inconsistent: You just need to assume the sentences in the set and prove a contradiction. Showing that a set is *not* provably inconsistent will be much harder. It would require more than just providing a proof or two; it would require showing that proofs of a certain kind are *impossible*.

6.8 Proofs and models

As you might already suspect, there is a connection between *theorems* and *tautologies*.

There is a formal way of showing that a sentence is a theorem: Prove it. For each line, we can check to see if that line follows by the cited rule. It may be hard to produce a twenty line proof, but it is not so hard to check each line of the proof and confirm that it is legitimate— and if each line of the proof individually is legitimate, then the whole proof is legitimate. Showing that a sentence is a tautology, though, requires reasoning in English about all possible models. There is no formal way of checking to see if the reasoning is sound. Given a choice between showing that a sentence is a theorem and showing that it is a tautology, it would be easier to show that it is a theorem.

Contrawise, there is no formal way of showing that a sentence is *not* a theorem. We would need to reason in English about all possible proofs. Yet there is a formal method for showing that a sentence is not a tautology. We need only construct a model in which the sentence is false. Given a choice between showing that a sentence is not a theorem and showing that it is not a tautology, it would be easier to show that it is not a tautology.

Fortunately, a sentence is a theorem if and only if it is a tautology. If we provide a proof of $\vdash \mathcal{A}$ and thus show that it is a theorem, it follows that \mathcal{A} is a tautology; i.e., $\models \mathcal{A}$. Similarly, if we construct a model in which \mathcal{A} is false and thus show that it is not a tautology, it follows that \mathcal{A} is not a theorem.

In general, $\mathcal{A} \vdash \mathcal{B}$ if and only if $\mathcal{A} \models \mathcal{B}$. As such:

- ▷ An argument is *valid* if and only if *the conclusion is derivable from the premises*.
- ▷ Two sentences are *logically equivalent* if and only if they are *provably equivalent*.
- ▷ A set of sentences is *consistent* if and only if it is *not provably inconsistent*.

You can pick and choose when to think in terms of proofs and when to think in terms of models, doing whichever is easier for a given task. Table 6.1 summarizes when it is best to give proofs and when it is best to give models.

In this way, proofs and models give us a versatile toolkit for working with arguments. If we can translate an argument into QL, then we can measure its logical weight in a purely formal way. If it is deductively valid, we can give a formal proof; if it is invalid, we can provide a formal counterexample.

	YES	NO
Is \mathcal{A} a tautology?	prove $\vdash \mathcal{A}$	give a model in which \mathcal{A} is false
Is \mathcal{A} a contradiction?	prove $\vdash \neg \mathcal{A}$	give a model in which \mathcal{A} is true
Is \mathcal{A} contingent?	give a model in which \mathcal{A} is true and another in which \mathcal{A} is false	prove $\vdash \mathcal{A}$ or $\vdash \neg \mathcal{A}$
Are \mathcal{A} and \mathcal{B} equivalent?	prove $\mathcal{A} \vdash \mathcal{B}$ and $\mathcal{B} \vdash \mathcal{A}$	give a model in which \mathcal{A} and \mathcal{B} have different truth values
Is the set \mathbb{A} consistent?	give a model in which all the sentences in \mathbb{A} are true	taking the sentences in \mathbb{A} , prove \mathcal{B} and $\neg \mathcal{B}$
Is the argument ' $\mathcal{P}, \therefore \mathcal{C}$ ' valid?	prove $\mathcal{P} \vdash \mathcal{C}$	give a model in which \mathcal{P} is true and \mathcal{C} is false

Table 6.1: Sometimes it is easier to show something by providing proofs than it is by providing models. Sometimes it is the other way round. It depends on what you are trying to show.

6.9 Soundness and completeness

This toolkit is incredibly convenient. It is also intuitive, because it seems natural that provability and semantic entailment should agree. Yet, do not be fooled by the similarity of the symbols ' \models ' and ' \vdash .' The fact that these two are really interchangeable is not a simple thing to prove.

Why should we think that an argument that *can be proven* is necessarily a *valid* argument? That is, why think that $\mathcal{A} \vdash \mathcal{B}$ implies $\mathcal{A} \models \mathcal{B}$?

This is the problem of SOUNDNESS. A proof system is SOUND if there are no proofs of invalid arguments. Demonstrating that the proof system is sound would require showing that *any* possible proof is the proof of a valid argument. It would not be enough simply to succeed when proving of many valid arguments and to fail when trying to prove invalid ones.

Fortunately, there is a way of approaching this in a step-wise fashion. If using the $\&E$ rule on the last line of a proof could never change a valid argument into an invalid one, then using the rule many times could not make an argument invalid. Similarly, if using the $\&E$ and $\forall E$ rules individually on the last line of a proof could never change a valid argument into an invalid one, then using them in combination could not either.

The strategy is to show for every rule of inference that it alone could not make a valid argument into an invalid one. It follows that the rules used in combination would not make a valid argument invalid. Since a proof is just a series of lines, each justified by a rule of inference, this would show that every provable argument is valid.

Consider, for example, the $\&E$ rule. Suppose we use it to add $\mathcal{A} \& \mathcal{B}$ to a valid argument. In order for the rule to apply, \mathcal{A} and \mathcal{B} must already be available in the proof. Since the argument so far is valid, \mathcal{A} and \mathcal{B} are either premises of the argument or valid consequences of the premises. As such, any model in which the premises are true must be a model in which \mathcal{A} and \mathcal{B} are true. According to the definition of TRUTH IN QL, this means that $\mathcal{A} \& \mathcal{B}$ is also true in such a model. Therefore, $\mathcal{A} \& \mathcal{B}$ validly follows from the premises. This means that using the $\&E$ rule to extend a valid proof produces another valid proof.

In order to show that the proof system is sound, we would need to show this for the other inference rules. Since the derived rules are consequences of the basic rules, it would suffice to provide similar arguments for the 16 other basic rules. This tedious exercise falls beyond the scope of this book.

Given a proof that the proof system is sound, it follows that every theorem is a tautology.

It is still possible to ask: Why think that *every* valid argument is an argument that can be proven? That is, why think that $\mathcal{A} \models \mathcal{B}$ implies $\mathcal{A} \vdash \mathcal{B}$?

This is the problem of COMPLETENESS. A proof system is COMPLETE if there is a proof of every valid argument. Completeness for a language like QL was first proven by Kurt Gödel in 1929. The proof is beyond the scope of this book.

The important point is that, happily, the proof system for QL is both sound and complete. This is not the case for all proof systems and all formal languages. Because it is true of QL, we can choose to give proofs or construct models—whichever is easier for the task at hand.

Summary of definitions

- ▷ A sentence \mathcal{A} is a THEOREM if and only if $\vdash \mathcal{A}$.
- ▷ Two sentences \mathcal{A} and \mathcal{B} are PROVABLY EQUIVALENT if and only if $\mathcal{A} \vdash \mathcal{B}$ and $\mathcal{B} \vdash \mathcal{A}$.
- ▷ $\{\mathcal{A}_1, \mathcal{A}_2, \dots\}$ is PROVABLY INCONSISTENT if and only if, for some sentence \mathcal{B} , $\{\mathcal{A}_1, \mathcal{A}_2, \dots\} \vdash (\mathcal{B} \& \neg \mathcal{B})$.

Practice Exercises

★ **Part A** Provide a justification (rule and line numbers) for each line of proof that requires one.

1	$W \rightarrow \neg B$
2	$A \& W$
3	$B \vee (J \& K)$
4	W
5	$\neg B$
6	$J \& K$
7	K

1	$L \leftrightarrow \neg O$
2	$L \vee \neg O$
3	$\neg L$
4	$\neg O$
5	L
6	$\neg L$
7	L

1	$Z \rightarrow (C \& \neg N)$
2	$\neg Z \rightarrow (N \& \neg C)$
3	$\neg(N \vee C)$
4	$\neg N \& \neg C$
5	Z
6	$C \& \neg N$
7	C
8	$\neg C$
9	$\neg Z$
10	$N \& \neg C$
11	N
12	$\neg N$
13	$N \vee C$

Part B Give a proof for each argument in SL.

1. $Q \rightarrow (Q \& \neg Q), \therefore \neg Q$
2. $J \rightarrow \neg J, \therefore \neg J$
3. $E \vee F, F \vee G, \neg F, \therefore E \& G$
4. $A \leftrightarrow B, B \leftrightarrow C, \therefore A \leftrightarrow C$
5. $M \vee (N \rightarrow M), \therefore \neg M \rightarrow \neg N$
6. $S \leftrightarrow T, \therefore S \leftrightarrow (T \vee S)$
7. $(M \vee N) \& (O \vee P), N \rightarrow P, \neg P, \therefore M \& O$
8. $(Z \& K) \vee (K \& M), K \rightarrow D, \therefore D$

Part C Show that each of the following sentences is a theorem in SL.

1. $O \rightarrow O$

2. $N \vee \neg N$
3. $\neg(P \& \neg P)$
4. $(A \rightarrow C) \rightarrow \neg(A \rightarrow \neg C)$
5. $J \leftrightarrow [J \vee (L \& \neg L)]$

Part D Show that each of the following pairs of sentences are provably equivalent in SL.

1. $\neg\neg\neg\neg G, G$
2. $T \rightarrow S, \neg S \rightarrow \neg T$
3. $D \leftrightarrow E, (D \rightarrow E) \& (E \rightarrow D)$
4. $\neg G \leftrightarrow H, \neg(G \leftrightarrow H)$
5. $U \rightarrow I, \neg(U \& \neg I)$

Part E Provide proofs to show each of the following.

1. $M \& (\neg N \rightarrow \neg M) \vdash (N \& M) \vee \neg M$
2. $\{C \rightarrow (E \& G), \neg C \rightarrow G\} \vdash G$
3. $\{(Z \& K) \leftrightarrow (Y \& M), D \& (D \rightarrow M)\} \vdash Y \rightarrow Z$
4. $\{(W \vee X) \vee (Y \vee Z), X \rightarrow Y, \neg Z\} \vdash W \vee Y$

Part F For the following, provide proofs using only the basic rules. The proofs will be longer than proofs of the same claims would be using the derived rules.

1. Show that MT is a legitimate derived rule. Using only the basic rules, prove the following: $\mathcal{A} \rightarrow \mathcal{B}, \neg \mathcal{B}, \therefore \neg \mathcal{A}$
2. Using only the basic rules, prove the following instance of DeMorgan's Laws: $(\neg A \& \neg B), \therefore \neg(A \vee B)$
3. Without using the QN rule, prove $\neg \exists x \neg \mathcal{A} \vdash \forall x \mathcal{A}$

★ **Part G** Provide a justification (rule and line numbers) for each line of proof that requires one.

1	$\forall x(\exists yLxy \rightarrow \forall zLzx)$
2	Lab
3	$\exists yLay \rightarrow \forall zLza$
4	$\exists yLay$
5	$\forall zLza$
6	Lca
7	$\exists yLcy \rightarrow \forall zLzc$
8	$\exists yLcy$
9	$\forall zLzc$
10	Lcc
11	$\forall xLxx$

1	$\forall x(Jx \rightarrow Kx)$	1	$\forall x\exists y(Rxy \vee Ryx)$
2	$\exists x\forall yLxy$	2	$\forall x\neg Rmx$
3	$\forall xJx$	3	$\exists y(Rmy \vee Rym)$
4	Ja	4	$Rma \vee Ram$
5	$Ja \rightarrow Ka$	5	$\neg Rma$
6	Ka	6	Ram
7	$\forall yLay$	7	$\exists xRxm$
8	Laa	8	$\exists xRxm$
9	$Ka \& Laa$	1	$\neg(\exists xMx \vee \forall x\neg Mx)$
10	$\exists x(Kx \& Lxx)$	2	$\neg\exists xMx \& \neg\forall x\neg Mx$
11	$\exists x(Kx \& Lxx)$	3	$\neg\exists xMx$
		4	$\forall x\neg Mx$
		5	$\neg\forall x\neg Mx$
		6	$\exists xMx \vee \forall x\neg Mx$

Part H Provide a proof of the argument about Billy on p. 65.

Part I Look back at Part C on p. 76. Provide proofs to show that each of the argument forms is valid in QL.

Part J Aristotle and his successors identified other syllogistic forms. Symbolize each of the following argument forms in QL and add the additional assumptions ‘There is an A ’ and ‘There is a B .’ Then prove that the supplemented argument forms are valid in QL.

Darapti: All B s are C s. All A s are B s. \therefore Some A is C .

Felapton: No B s are C s. All A s are B s. \therefore Some A is not C .

Barbari: All B s are C s. All A s are B s. \therefore Some A is C .

Camestros: All C s are B s. No A s are B s. \therefore Some A is not C .

Celarent: No B s are C s. All A s are B s. \therefore Some A is not C .

Cesaro: No C s are B s. All A s are B s. \therefore Some A is not C .

Fapesmo: All B s are C s. No A s are B s. \therefore Some C is not A .

Part K Provide a proof of each claim.

1. $\forall x\forall yGxy \vdash \exists xGxx$
2. $\forall x\forall y(Gxy \rightarrow Gyx) \vdash \forall x\forall y(Gxy \leftrightarrow Gyx)$
3. $\{\forall x(Ax \rightarrow Bx), \exists xAx\} \vdash \exists xBx$
4. $\{Na \rightarrow \forall x(Mx \leftrightarrow Ma), Ma, \neg Mb\} \vdash \neg Na$
5. $\vdash \forall z(Pz \vee \neg Pz)$
6. $\vdash \forall xRxx \rightarrow \exists x\exists yRxy$
7. $\vdash \forall y\exists x(Qy \rightarrow Qx)$

Part L Show that each pair of sentences is provably equivalent.

1. $\forall x(Ax \rightarrow \neg Bx), \neg\exists x(Ax \& Bx)$
2. $\forall x(\neg Ax \rightarrow Bx), \forall xAx \vee Bx$
3. $\exists xPx \rightarrow Qc, \forall x(Px \rightarrow Qc)$
4. $Rca \leftrightarrow \forall xRxa, \forall x(Rca \leftrightarrow Rxa)$

Part M Show that each of the following is provably inconsistent.

1. $\{Sa \rightarrow Tm, Tm \rightarrow Sa, Tm \& \neg Sa\}$
2. $\{\exists xRxa, \forall x\forall yRyx\}$
3. $\{\neg\exists x\exists yLxy, Laa\}$
4. $\{\forall x(Px \rightarrow Qx), \forall z(Pz \rightarrow Rz), \forall yPy, \neg Qa \& \neg Rb\}$

★ **Part N** Write a symbolization key for the following argument, translate it, and prove it:

There is someone who likes everyone who likes everyone that he likes.
Therefore, there is someone who likes himself.

Part O Provide a proof of each claim.

1. $\{Pa \vee Qb, Qb \rightarrow b = c, \neg Pa\} \vdash Qc$
2. $\{m = n \vee n = o, An\} \vdash Am \vee Ao$
3. $\{\forall xx = m, Rma\} \vdash \exists xRxx$
4. $\neg\exists xx \neq m \vdash \forall x\forall y(Px \rightarrow Py)$
5. $\forall x\forall y(Rxy \rightarrow x = y) \vdash Rab \rightarrow Rba$
6. $\{\exists xJx, \exists x\neg Jx\} \vdash \exists x\exists yx \neq y$
7. $\{\forall x(x = n \leftrightarrow Mx), \forall x(Ox \& \neg Mx)\} \vdash On$
8. $\{\exists xDx, \forall x(x = p \leftrightarrow Dx)\} \vdash Dp$
9. $\{\exists x[Kx \& \forall y(Ky \rightarrow x = y) \& Bx], Kd\} \vdash Bd$
10. $\vdash Pa \rightarrow \forall x(Px \vee x \neq a)$

Part P Look back at Part E on p. 78. For each argument: If it is valid in QL, give a proof. If it is invalid, construct a model to show that it is invalid.

* **Part Q** For each of the following pairs of sentences: If they are logically equivalent in QL, give proofs to show this. If they are not, construct a model to show this.

1. $\forall xPx \rightarrow Qc, \forall x(Px \rightarrow Qc)$
2. $\forall xPx \& Qc, \forall x(Px \& Qc)$
3. $Qc \vee \exists xQx, \exists x(Qc \vee Qx)$
4. $\forall x\forall y\forall zBxyz, \forall xBxxx$
5. $\forall x\forall yDxy, \forall y\forall xDxy$
6. $\exists x\forall yDxy, \forall y\exists xDxy$

* **Part R** For each of the following arguments: If it is valid in QL, give a proof. If it is invalid, construct a model to show that it is invalid.

1. $\forall x\exists yRxy, \therefore \exists y\forall xRxy$
2. $\exists y\forall xRxy, \therefore \forall x\exists yRxy$
3. $\exists x(Px \& \neg Qx), \therefore \forall x(Px \rightarrow \neg Qx)$
4. $\forall x(Sx \rightarrow Ta), Sd, \therefore Ta$
5. $\forall x(Ax \rightarrow Bx), \forall x(Bx \rightarrow Cx), \therefore \forall x(Ax \rightarrow Cx)$
6. $\exists x(Dx \vee Ex), \forall x(Dx \rightarrow Fx), \therefore \exists x(Dx \& Fx)$
7. $\forall x\forall y(Rxy \vee Ryx), \therefore Rjj$
8. $\exists x\exists y(Rxy \vee Ryx), \therefore Rjj$
9. $\forall xPx \rightarrow \forall xQx, \exists x\neg Px, \therefore \exists x\neg Qx$
10. $\exists xMx \rightarrow \exists xNx, \neg\exists xNx, \therefore \forall x\neg Mx$

Part S

1. If you know that $\mathcal{A} \vdash \mathcal{B}$, what can you say about $(\mathcal{A} \& \mathcal{C}) \vdash \mathcal{B}$? Explain your answer.
2. If you know that $\mathcal{A} \vdash \mathcal{B}$, what can you say about $(\mathcal{A} \vee \mathcal{C}) \vdash \mathcal{B}$? Explain your answer.

Appendix A

Symbolic notation

In the history of formal logic, different symbols have been used at different times and by different authors. Often, authors were forced to use notation that their printers could typeset.

In one sense, the symbols used for various logical constants is arbitrary. There is nothing written in heaven that says that ‘ \neg ’ must be the symbol for truth-functional negation. We might have specified a different symbol to play that part. Once we have given definitions for well-formed formulae (wff) and for truth in our logic languages, however, using ‘ \neg ’ is no longer arbitrary. That is the symbol for negation in this textbook, and so it is the symbol for negation when writing sentences in our languages SL or QL.

This appendix presents some common symbols, so that you can recognize them if you encounter them in an article or in another book.

Negation Two commonly used symbols are the *hoe*, ‘ \neg ’, and the *swung dash*, ‘ \sim .’ In some more advanced formal systems it is necessary to distinguish between two kinds of negation; the distinction is sometimes represented by using both ‘ \neg ’ and ‘ \sim .’

Disjunction The symbol ‘ \vee ’ is typically used to symbolize inclusive disjunction.

Conjunction Conjunction is often symbolized with the *ampersand*, ‘ $\&$.’ The ampersand is actually a decorative form of the Latin word ‘et’ which means ‘and’; it is commonly used in English writing. As a symbol in a formal sys-

summary of symbols	
negation	\neg, \sim
conjunction	$\&, \wedge, \cdot$
disjunction	\vee
conditional	\rightarrow, \supset
biconditional	\leftrightarrow, \equiv

tem, the ampersand is not the word ‘and’; its meaning is given by the formal semantics for the language. Perhaps to avoid this confusion, some systems use a different symbol for conjunction. For example, ‘ \wedge ’ is a counterpart to the symbol used for disjunction. Sometimes a single dot ‘ \cdot ’ is used. In some older texts, there is no symbol for conjunction at all; ‘ A and B ’ is simply written ‘ AB .’

Material Conditional There are two common symbols for the material conditional: the *arrow*, ‘ \rightarrow ’, and the *hook*, ‘ \supset .’

Material Biconditional The *double-headed arrow*, ‘ \leftrightarrow ’, is used in systems that use the arrow to represent the material conditional. Systems that use the hook for the conditional typically use the *triple bar*, ‘ \equiv ’, for the biconditional.

Quantifiers The universal quantifier is typically symbolized as an upside-down A, ‘ \forall ’, and the existential quantifier as a backwards E, ‘ \exists .’ In some texts, the quantifiers are symbolized with larger versions of the symbols used for conjunction and disjunction. Although quantified expressions cannot be translated into expressions without quantifiers, there is a conceptual connection between the universal quantifier and conjunction and between the existential quantifier and disjunction. Consider the sentence $\exists xPx$, for example. It means that *either* the first member of the UD is a P , *or* the second one is, *or* the third one is, Such a system uses the symbol ‘ \bigvee ’ instead of ‘ \exists .’

Polish notation

This section briefly discusses sentential logic in Polish notation, a system of notation introduced in the late 1920s by the Polish logician Jan Łukasiewicz.

Lower case letters are used as sentence letters. The letter N is used for negation. A is used for disjunction, K for conjunction, C for the conditional, E for the biconditional. (‘ A ’ is for alternation, another name for logical disjunction. ‘ E ’ is for equivalence.)

In Polish notation, a binary connective is written *before* the two sentences that it connects. For example, the sentence $A \& B$ of SL would be written Kab in Polish notation.

The sentences $\neg A \rightarrow B$ and $\neg(A \rightarrow B)$ are very different; the main logical operator of the first is the conditional, but the main connective of the second is negation. In SL, we show this by putting parentheses around the conditional

in the second sentence. In Polish notation, parentheses are never required. The left-most connective is always the main connective. The first sentence would simply be written $CNab$ and the second $NCab$.

This feature of Polish notation means that it is possible to evaluate sentences simply by working through the symbols from right to left. If you were constructing a truth table for $NKab$, for example, you would first consider the truth-values assigned to a and b , then consider their conjunction, and then negate the result. The general rule for what to evaluate next in SL is not nearly so simple. In SL, the truth table for $\neg(A \& B)$ requires looking at A and B , then looking in the middle of the sentence at the conjunction, and then at the beginning of the sentence at the negation. Because the order of operations can be specified more mechanically in Polish notation, variants of Polish notation are used as the internal structure for many computer programming languages.

Appendix B

Solutions to selected exercises

Many of the exercises may be answered correctly in different ways. Where that is the case, the solution here represents one possible correct answer.

Chapter 1 Part C

1. consistent
2. inconsistent
3. consistent
4. consistent

Chapter 1 Part D

1, 2, 3, 6, 8, and 10 are possible.

Chapter 2 Part A

1. $\neg M$
2. $M \vee \neg M$
3. $G \vee C$

4. $\neg C \ \& \ \neg G$
5. $C \rightarrow (\neg G \ \& \ \neg M)$
6. $M \vee (C \vee G)$

Chapter 2 Part C

1. $E_1 \ \& \ E_2$
2. $F_1 \rightarrow S_1$
3. $F_1 \vee E_1$
4. $E_2 \ \& \ \neg S_2$
5. $\neg E_1 \ \& \ \neg E_2$
6. $E_1 \ \& \ E_2 \ \& \ \neg(S_1 \vee S_2)$
7. $S_2 \rightarrow F_2$
8. $(\neg E_1 \rightarrow \neg E_2) \ \& \ (E_1 \rightarrow E_2)$
9. $S_1 \leftrightarrow \neg S_2$
10. $(E_2 \ \& \ F_2) \rightarrow S_2$
11. $\neg(E_2 \ \& \ F_2)$
12. $(F_1 \ \& \ F_2) \leftrightarrow (\neg E_1 \ \& \ \neg E_2)$

Chapter 2 Part D

- A:** Alice is a spy.
B: Bob is a spy.
C: The code has been broken.
G: The German embassy will be in an uproar.

1. $A \ \& \ B$
2. $(A \vee B) \rightarrow C$
3. $\neg(A \vee B) \rightarrow \neg C$
4. $G \vee C$
5. $(C \vee \neg C) \ \& \ G$
6. $(A \vee B) \ \& \ \neg(A \ \& \ B)$

Chapter 2 Part G

1. (a) no (b) no
2. (a) no (b) yes
3. (a) yes (b) yes
4. (a) no (b) no

5. (a) yes (b) yes
6. (a) no (b) no
7. (a) no (b) yes
8. (a) no (b) yes
9. (a) no (b) no

Chapter 3 Part A

1. tautology
2. contradiction
3. contingent
4. tautology
5. tautology
6. contingent
7. tautology
8. contradiction
9. tautology
10. contradiction
11. tautology
12. contingent
13. contradiction
14. contingent
15. tautology
16. tautology
17. contingent
18. contingent

Chapter 3 Part B

2, 5, 6, 8, and 9 are logically equivalent.

Chapter 3 Part C

1, 3, 6, 7, and 8 are consistent.

Chapter 3 Part D

3, 5, 8, and 10 are valid.

Chapter 3 Part E

1. \mathcal{A} and \mathcal{B} have the same truth value on every line of a complete truth table, so $\mathcal{A} \leftrightarrow \mathcal{B}$ is true on every line. It is a tautology.
2. The sentence is false on some line of a complete truth table. On that line, \mathcal{A} and \mathcal{B} are true and \mathcal{C} is false. So the argument is invalid.
3. Since there is no line of a complete truth table on which all three sentences are true, the conjunction is false on every line. So it is a contradiction.
4. Since \mathcal{A} is false on every line of a complete truth table, there is no line on which \mathcal{A} and \mathcal{B} are true and \mathcal{C} is false. So the argument is valid.
5. Since \mathcal{C} is true on every line of a complete truth table, there is no line on which \mathcal{A} and \mathcal{B} are true and \mathcal{C} is false. So the argument is valid.
6. Not much. $(\mathcal{A} \vee \mathcal{B})$ is a tautology if \mathcal{A} and \mathcal{B} are tautologies; it is a contradiction if they are contradictions; it is contingent if they are contingent.
7. \mathcal{A} and \mathcal{B} have different truth values on at least one line of a complete truth table, and $(\mathcal{A} \vee \mathcal{B})$ will be true on that line. On other lines, it might be true or false. So $(\mathcal{A} \vee \mathcal{B})$ is either a tautology or it is contingent; it is *not* a contradiction.

Chapter 3 Part F

1. $\neg A \rightarrow B$
2. $\neg(A \rightarrow \neg B)$
3. $\neg[(A \rightarrow B) \rightarrow \neg(B \rightarrow A)]$

Chapter 4 Part B

1. $Za \ \& \ Zb \ \& \ Zc$
2. $Rb \ \& \ \neg Ab$
3. $Lcb \rightarrow Mb$
4. $(Ab \ \& \ Ac) \rightarrow (Lab \ \& \ Lac)$
5. $\exists x(Rx \ \& \ Zx)$
6. $\forall x(Ax \rightarrow Rx)$
7. $\forall x[Zx \rightarrow (Mx \vee Ax)]$
8. $\exists x(Rx \ \& \ \neg Ax)$
9. $\exists x(Rx \ \& \ Lcx)$
10. $\forall x[(Mx \ \& \ Zx) \rightarrow Lbx]$
11. $\forall x[(Mx \ \& \ Lax) \rightarrow Lxa]$
12. $\exists xRx \rightarrow Ra$
13. $\forall x(Ax \rightarrow Rx)$
14. $\forall x[(Mx \ \& \ Lcx) \rightarrow Lax]$
15. $\exists x(Mx \ \& \ Lxb \ \& \ \neg Lbx)$

Chapter 4 Part F

1. $\neg\exists xTx$
2. $\forall x(Mx \rightarrow Sx)$
3. $\exists x\neg Sx$
4. $\exists x[Cx \& \neg\exists yByx]$
5. $\neg\exists xBxx$
6. $\neg\exists x(Cx \& \neg Sx \& Tx)$
7. $\exists x(Cx \& Tx) \& \exists x(Mx \& Tx) \& \neg\exists x(Cx \& Mx \& Tx)$
8. $\forall x[Cx \rightarrow \forall y(\neg Cy \rightarrow Bxy)]$
9. $\forall x((Cx \& Mx) \rightarrow \forall y[(\neg Cy \& \neg My) \rightarrow Bxy])$

Chapter 4 Part H

1. $\forall x(Cxp \rightarrow Dx)$
2. $Cjp \& Fj$
3. $\exists x(Cxp \& Fx)$
4. $\neg\exists xSxj$
5. $\forall x[(Cxp \& Fx) \rightarrow Dx]$
6. $\neg\exists x(Cxp \& Mx)$
7. $\exists x(Cjx \& Sxe \& Fj)$
8. $Spe \& Mp$
9. $\forall x[(Sxp \& Mx) \rightarrow \neg\exists yCyx]$
10. $\exists x(Sxj \& \exists yCyx \& Fj)$
11. $\forall x[Dx \rightarrow \exists y(Sxy \& Fy \& Dy)]$
12. $\forall x[(Mx \& Dx) \rightarrow \exists y(Cxy \& Dy)]$

Chapter 4 Part J

1. $\forall x(Cx \rightarrow Bx)$
2. $\neg\exists xWx$
3. $\exists x\exists y(Cx \& Cy \& x \neq y)$
4. $\exists x\exists y(Jx \& Ox \& Jy \& Oy \& x \neq y)$
5. $\forall x\forall y\forall z[(Jx \& Ox \& Jy \& Oy \& Jz \& Oz) \rightarrow (x = y \vee x = z \vee y = z)]$
6. $\exists x\exists y(Jx \& Bx \& Jy \& By \& \forall z[(Jz \& Bz) \rightarrow (x = z \vee y = z)])$
7. $\exists x_1\exists x_2\exists x_3\exists x_4[Dx_1 \& Dx_2 \& Dx_3 \& Dx_4 \& x_1 \neq x_2 \& x_1 \neq x_3 \& x_1 \neq x_4 \& x_2 \neq x_3 \& x_2 \neq x_4 \& x_3 \neq x_4 \& \neg\exists y(Dy \& y \neq x_1 \& y \neq x_2 \& y \neq x_3 \& y \neq x_4)]$
8. $\exists x(Dx \& Cx \& \forall y[(Dy \& Cy) \rightarrow x = y] \& Bx)$
9. $\forall x[(Ox \& Jx) \rightarrow Wx] \& \exists x[Mx \& \forall y(My \rightarrow x = y) \& Wx]$
10. $\exists x(Dx \& Cx \& \forall y[(Dy \& Cy) \rightarrow x = y] \& Wx) \rightarrow \exists x(Wx \leftrightarrow x = y)$

11. wide scope: $\neg\exists x[Mx \& \forall y(My \rightarrow x = y) \& Jx]$
 narrow scope: $\exists x[Mx \& \forall y(My \rightarrow x = y) \& \neg Jx]$
12. wide scope: $\neg\exists x\exists z(Dx \& Cx \& Mz \& \forall y[(Dy \& Cy) \rightarrow x = y] \& \forall y[(My \rightarrow z = y) \& x = z])$
 narrow scope: $\exists x\exists z(Dx \& Cx \& Mz \& \forall y[(Dy \& Cy) \rightarrow x = y] \& \forall y[(My \rightarrow z = y) \& x \neq z])$

Chapter 5 Part A

2, 3, 4, 6, 8, and 9 are true in the model.

Chapter 5 Part B

2, 4, 5, and 7 are true in the model.

Chapter 5 Part C

1. The sentence is true in this model:

$$\begin{aligned} \text{UD} &= \{\text{Stan}\} \\ \text{extension}(D) &= \{\text{Stan}\} \\ \text{referent}(a) &= \text{Stan} \\ \text{referent}(b) &= \text{Stan} \end{aligned}$$

And it is false in this model:

$$\begin{aligned} \text{UD} &= \{\text{Stan}\} \\ \text{extension}(D) &= \emptyset \\ \text{referent}(a) &= \text{Stan} \\ \text{referent}(b) &= \text{Stan} \end{aligned}$$

2. The sentence is true in this model:

$$\begin{aligned} \text{UD} &= \{\text{Stan}\} \\ \text{extension}(T) &= \{\langle \text{Stan}, \text{Stan} \rangle\} \\ \text{referent}(h) &= \text{Stan} \end{aligned}$$

And it is false in this model:

$$\begin{aligned} \text{UD} &= \{\text{Stan}\} \\ \text{extension}(T) &= \emptyset \\ \text{referent}(h) &= \text{Stan} \end{aligned}$$

3. The sentence is true in this model:

$$\begin{aligned} \text{UD} &= \{\text{Stan, Ollie}\} \\ \text{extension}(P) &= \{\text{Stan}\} \\ \text{referent}(m) &= \text{Stan} \end{aligned}$$

And it is false in this model:

$$\begin{aligned} \text{UD} &= \{\text{Stan}\} \\ \text{extension}(P) &= \emptyset \\ \text{referent}(m) &= \text{Stan} \end{aligned}$$

Chapter 5 Part D

There are many possible correct answers. Here are some:

1. Making the first sentence true and the second false:

$$\begin{aligned} \text{UD} &= \{\text{alpha}\} \\ \text{extension}(J) &= \{\text{alpha}\} \\ \text{extension}(K) &= \emptyset \\ \text{referent}(a) &= \text{alpha} \end{aligned}$$

2. Making the first sentence true and the second false:

$$\begin{aligned} \text{UD} &= \{\text{alpha, omega}\} \\ \text{extension}(J) &= \{\text{alpha}\} \\ \text{referent}(m) &= \text{omega} \end{aligned}$$

3. Making the first sentence false and the second true:

$$\begin{aligned} \text{UD} &= \{\text{alpha, omega}\} \\ \text{extension}(R) &= \{\langle \text{alpha, alpha} \rangle\} \end{aligned}$$

4. Making the first sentence false and the second true:

$$\begin{aligned} \text{UD} &= \{\text{alpha, omega}\} \\ \text{extension}(P) &= \{\text{alpha}\} \\ \text{extension}(Q) &= \emptyset \\ \text{referent}(c) &= \text{alpha} \end{aligned}$$

5. Making the first sentence true and the second false:

$$\begin{aligned} \text{UD} &= \{\text{iota}\} \\ \text{extension}(P) &= \emptyset \\ \text{extension}(Q) &= \emptyset \end{aligned}$$

6. Making the first sentence false and the second true:

$$\begin{aligned} \text{UD} &= \{\text{iota}\} \\ \text{extension}(P) &= \emptyset \\ \text{extension}(Q) &= \{\text{iota}\} \end{aligned}$$

7. Making the first sentence true and the second false:

$$\begin{aligned} \text{UD} &= \{\text{iota}\} \\ \text{extension}(P) &= \emptyset \\ \text{extension}(Q) &= \{\text{iota}\} \end{aligned}$$

8. Making the first sentence true and the second false:

$$\text{UD} = \{\text{alpha}, \text{omega}\}$$

$$\text{extension}(R) = \{\langle \text{alpha}, \text{omega} \rangle, \langle \text{omega}, \text{alpha} \rangle\}$$

9. Making the first sentence false and the second true:

$$\text{UD} = \{\text{alpha}, \text{omega}\}$$

$$\text{extension}(R) = \{\langle \text{alpha}, \text{alpha} \rangle, \langle \text{alpha}, \text{omega} \rangle\}$$

Chapter 5 Part G

1. There are many possible answers. Here is one:

$$\text{UD} = \{\text{Harry}, \text{Sally}\}$$

$$\text{extension}(R) = \{\langle \text{Sally}, \text{Harry} \rangle\}$$

$$\text{referent}(a) = \text{Harry}$$

2. There are no predicates or constants, so we only need to give a UD. Any UD with 2 members will do.
3. We need to show that it is impossible to construct a model in which these are both true. Suppose $\exists x x \neq a$ is true in a model. There is something in the universe of discourse that is *not* the referent of a . So there are at least two things in the universe of discourse: $\text{referent}(a)$ and this other thing. Call this other thing β — we know $a \neq \beta$. But if $a \neq \beta$, then $\forall x \forall y x = y$ is false. So the first sentence must be false if the second sentence is true. As such, there is no model in which they are both true. Therefore, they are inconsistent.

Chapter 6 Part A

1	$W \rightarrow \neg B$		1	$L \leftrightarrow \neg O$	
2	$A \& W$		2	$L \vee \neg O$	
3	$B \vee (J \& K)$		3	$\neg L$	
4	W	& E 2	4	$\neg O$	$\vee E$ 2, 3
5	$\neg B$	$\rightarrow E$ 1, 4	5	L	$\leftrightarrow E$ 1, 4
6	$J \& K$	$\vee E$ 3, 5	6	$\neg L$	R 3
7	K	& E 6	7	L	$\neg E$ 3-6

1	$Z \rightarrow (C \& \neg N)$		
2	$\neg Z \rightarrow (N \& \neg C)$		
3	<hr style="width: 100%;"/> $\neg(N \vee C)$		
4	$\neg N \& \neg C$	DeM 3	
5	Z		
6	<hr style="width: 100%;"/> $C \& \neg N$	\rightarrow E 1, 5	
7	C	$\&$ E 6	
8	$\neg C$	$\&$ E 4	
9	$\neg Z$	\neg I 5-8	
10	$N \& \neg C$	\rightarrow E 2, 9	
11	N	$\&$ E 10	
12	$\neg N$	$\&$ E 4	
13	$N \vee C$	\neg E 3-12	

Chapter 6 Part G

1	$\forall x(\exists yLxy \rightarrow \forall zLzx)$		1	$\forall x(Jx \rightarrow Kx)$	
2	Lab		2	$\exists x\forall yLxy$	
3	<hr style="width: 100%;"/> $\exists yLay \rightarrow \forall zLza$	\forall E 1	3	$\forall xJx$	
4	$\exists yLay$	\exists I 2	4	Ja	\forall E 3
5	$\forall zLza$	\rightarrow E 3, 4	5	$Ja \rightarrow Ka$	\forall E 1
6	Lca	\forall E 5	6	Ka	\rightarrow E 5, 4
7	$\exists yLcy \rightarrow \forall zLzc$	\forall E 1	7	$\forall yLay$	
8	$\exists yLcy$	\exists I 5	8	<hr style="width: 100%;"/> Laa	\forall E 7
9	$\forall zLzc$	\rightarrow E 7, 8	9	$Ka \& Laa$	$\&$ I 6, 8
10	Lcc	\forall E 9	10	$\exists x(Kx \& Lxx)$	\exists I 9
11	$\forall xLxx$	\forall I 10	11	$\exists x(Kx \& Lxx)$	\exists E 2, 7-10

<table style="border-collapse: collapse; width: 100%;"> <tr><td style="border-right: 1px solid black; padding-right: 5px;">1</td><td style="padding-right: 5px;">$\forall x \exists y (Rxy \vee Ryx)$</td><td></td></tr> <tr><td style="border-right: 1px solid black; padding-right: 5px;">2</td><td style="padding-right: 5px;">$\forall x \neg Rmx$</td><td></td></tr> <tr><td style="border-right: 1px solid black; padding-right: 5px;">3</td><td style="padding-right: 5px;">$\exists y (Rmy \vee Rym)$</td><td style="padding-left: 5px;">$\forall E$ 1</td></tr> <tr><td style="border-right: 1px solid black; padding-right: 5px;">4</td><td style="padding-right: 5px;">$Rma \vee Ram$</td><td></td></tr> <tr><td style="border-right: 1px solid black; padding-right: 5px;">5</td><td style="padding-right: 5px;">$\neg Rma$</td><td style="padding-left: 5px;">$\forall E$ 2</td></tr> <tr><td style="border-right: 1px solid black; padding-right: 5px;">6</td><td style="padding-right: 5px;">Ram</td><td style="padding-left: 5px;">$\vee E$ 4, 5</td></tr> <tr><td style="border-right: 1px solid black; padding-right: 5px;">7</td><td style="padding-right: 5px;">$\exists x Rxm$</td><td style="padding-left: 5px;">$\exists I$ 6</td></tr> <tr><td style="border-right: 1px solid black; padding-right: 5px;">8</td><td style="padding-right: 5px;">$\exists x Rxm$</td><td style="padding-left: 5px;">$\exists E$ 3, 4-7</td></tr> </table>	1	$\forall x \exists y (Rxy \vee Ryx)$		2	$\forall x \neg Rmx$		3	$\exists y (Rmy \vee Rym)$	$\forall E$ 1	4	$Rma \vee Ram$		5	$\neg Rma$	$\forall E$ 2	6	Ram	$\vee E$ 4, 5	7	$\exists x Rxm$	$\exists I$ 6	8	$\exists x Rxm$	$\exists E$ 3, 4-7	<table style="border-collapse: collapse; width: 100%;"> <tr><td style="border-right: 1px solid black; padding-right: 5px;">1</td><td style="padding-right: 5px;">$\neg(\exists x Mx \vee \forall x \neg Mx)$</td><td></td></tr> <tr><td style="border-right: 1px solid black; padding-right: 5px;">2</td><td style="padding-right: 5px;">$\neg \exists x Mx \ \& \ \neg \forall x \neg Mx$</td><td style="padding-left: 5px;">DeM 1</td></tr> <tr><td style="border-right: 1px solid black; padding-right: 5px;">3</td><td style="padding-right: 5px;">$\neg \exists x Mx$</td><td style="padding-left: 5px;">$\& E$ 2</td></tr> <tr><td style="border-right: 1px solid black; padding-right: 5px;">4</td><td style="padding-right: 5px;">$\forall x \neg Mx$</td><td style="padding-left: 5px;">QN 3</td></tr> <tr><td style="border-right: 1px solid black; padding-right: 5px;">5</td><td style="padding-right: 5px;">$\neg \forall x \neg Mx$</td><td style="padding-left: 5px;">$\& E$ 2</td></tr> <tr><td style="border-right: 1px solid black; padding-right: 5px;">6</td><td style="padding-right: 5px;">$\exists x Mx \vee \forall x \neg Mx$</td><td style="padding-left: 5px;">$\neg E$ 1-5</td></tr> </table>	1	$\neg(\exists x Mx \vee \forall x \neg Mx)$		2	$\neg \exists x Mx \ \& \ \neg \forall x \neg Mx$	DeM 1	3	$\neg \exists x Mx$	$\& E$ 2	4	$\forall x \neg Mx$	QN 3	5	$\neg \forall x \neg Mx$	$\& E$ 2	6	$\exists x Mx \vee \forall x \neg Mx$	$\neg E$ 1-5
1	$\forall x \exists y (Rxy \vee Ryx)$																																										
2	$\forall x \neg Rmx$																																										
3	$\exists y (Rmy \vee Rym)$	$\forall E$ 1																																									
4	$Rma \vee Ram$																																										
5	$\neg Rma$	$\forall E$ 2																																									
6	Ram	$\vee E$ 4, 5																																									
7	$\exists x Rxm$	$\exists I$ 6																																									
8	$\exists x Rxm$	$\exists E$ 3, 4-7																																									
1	$\neg(\exists x Mx \vee \forall x \neg Mx)$																																										
2	$\neg \exists x Mx \ \& \ \neg \forall x \neg Mx$	DeM 1																																									
3	$\neg \exists x Mx$	$\& E$ 2																																									
4	$\forall x \neg Mx$	QN 3																																									
5	$\neg \forall x \neg Mx$	$\& E$ 2																																									
6	$\exists x Mx \vee \forall x \neg Mx$	$\neg E$ 1-5																																									

Chapter 6 Part N

Regarding the translation of this argument, see p. 67.

1	$\exists x \forall y [\forall z (Lxz \rightarrow Lyz) \rightarrow Lxy]$	
2	$\forall y [\forall z (Laz \rightarrow Lyz) \rightarrow Lay]$	
3	$\forall z (Laz \rightarrow Laz) \rightarrow Laa$	$\forall E$ 2
4	$\neg \exists x Lxx$	for reductio
5	$\forall x \neg Lxx$	QN 4
6	$\neg Laa$	$\forall E$ 5
7	$\neg \forall z (Laz \rightarrow Laz)$	MT 5, 6
8	Lab	
9	Lab	R 8
10	$Lab \rightarrow Lab$	$\rightarrow I$ 8-9
11	$\forall z (Laz \rightarrow Laz)$	$\forall I$ 10
12	$\neg \forall z (Laz \rightarrow Laz)$	R 7
13	$\exists x Lxx$	$\neg E$ 4-12
14	$\exists x Lxx$	$\exists E$ 1, 2-13

Chapter 6 Part Q

2, 3, and 5 are logically equivalent.

Chapter 6 Part R

2, 4, 5, 7, and 10 are valid. Here are complete answers for some of them:

1. $UD = \{\text{mocha, freddo}\}$
 $\text{extension}(R) = \{\langle \text{mocha, freddo} \rangle, \langle \text{freddo, mocha} \rangle\}$
2.

1	$\exists y \forall x Rxy$	want $\forall x \exists y Rxy$
2	$\forall x Rxa$	
3	Rba	$\forall E$ 2
4	$\exists y Rby$	$\exists I$ 3
5	$\forall x \exists y Rxy$	$\forall I$ 4
6	$\forall x \exists y Rxy$	$\exists E$ 1, 2-5

Appendix C

Quick reference

SENTENTIAL CONNECTIVES (chapter 2)

It is not the case that P .	$\neg P$
Either P , or Q .	$(P \vee Q)$
Neither P , nor Q .	$\neg(P \vee Q)$ or $(\neg P \ \& \ \neg Q)$
Both P , and Q .	$(P \ \& \ Q)$
If P , then Q .	$(P \rightarrow Q)$
P only if Q .	$(P \rightarrow Q)$
P if and only if Q .	$(P \leftrightarrow Q)$
Unless P , Q . P unless Q .	$(P \vee Q)$

PREDICATES (chapter 4)

All F s are G s.	$\forall x(Fx \rightarrow Gx)$
Some F s are G s.	$\exists x(Fx \ \& \ Gx)$
Not all F s are G s.	$\neg\forall x(Fx \rightarrow Gx)$ or $\exists x(Fx \ \& \ \neg Gx)$
No F s are G s.	$\forall x(Fx \rightarrow \neg Gx)$ or $\neg\exists x(Fx \ \& \ Gx)$

IDENTITY (section 4.6)

Only j is G .	$Gj \ \& \ \forall x(Gx \leftrightarrow x = j)$
Everything besides j is G .	$\forall x(x \neq j \rightarrow Gx)$
The F is G .	$\exists x(Fx \ \& \ \forall y(Fy \rightarrow x = y) \ \& \ Gx)$

'The F is not G ' can be translated two ways:

It is not the case that the F is G . (wide)	$\neg\exists x(Fx \ \& \ \forall y(Fy \rightarrow x = y) \ \& \ Gx)$
The F is non- G . (narrow)	$\exists x(Fx \ \& \ \forall y(Fy \rightarrow x = y) \ \& \ \neg Gx)$

Using identity to symbolize quantities

There are at least _____ *F*s.

- one** $\exists xFx$
two $\exists x_1\exists x_2(Fx_1 \& Fx_2 \& x_1 \neq x_2)$
three $\exists x_1\exists x_2\exists x_3(Fx_1 \& Fx_2 \& Fx_3 \& x_1 \neq x_2 \& x_1 \neq x_3 \& x_2 \neq x_3)$
four $\exists x_1\exists x_2\exists x_3\exists x_4(Fx_1 \& Fx_2 \& Fx_3 \& Fx_4 \& x_1 \neq x_2 \& x_1 \neq x_3 \& x_1 \neq x_4 \& x_2 \neq x_3 \& x_2 \neq x_4 \& x_3 \neq x_4)$
n $\exists x_1 \cdots \exists x_n(Fx_1 \& \cdots \& Fx_n \& x_1 \neq x_2 \& \cdots \& x_{n-1} \neq x_n)$

There are at most _____ *F*s.

One way to say ‘at most n things are F ’ is to put a negation sign in front of one of the symbolizations above and say ‘at least $n+1$ things are F .’ Equivalently:

- one** $\forall x_1\forall x_2[(Fx_1 \& Fx_2) \rightarrow x_1 = x_2]$
two $\forall x_1\forall x_2\forall x_3[(Fx_1 \& Fx_2 \& Fx_3) \rightarrow (x_1 = x_2 \vee x_1 = x_3 \vee x_2 = x_3)]$
three $\forall x_1\forall x_2\forall x_3\forall x_4[(Fx_1 \& Fx_2 \& Fx_3 \& Fx_4) \rightarrow (x_1 = x_2 \vee x_1 = x_3 \vee x_1 = x_4 \vee x_2 = x_3 \vee x_2 = x_4 \vee x_3 = x_4)]$
n $\forall x_1 \cdots \forall x_{n+1}[(Fx_1 \& \cdots \& Fx_{n+1}) \rightarrow (x_1 = x_2 \vee \cdots \vee x_n = x_{n+1})]$

There are exactly _____ *F*s.

One way to say ‘exactly n things are F ’ is to conjoin two of the symbolizations above and say ‘at least n things are F ’ & ‘at most n things are F .’ The following equivalent formulae are shorter:

- zero** $\forall x\neg Fx$
one $\exists x[Ax \& \neg\exists y(Ay \& x \neq y)]$
two $\exists x_1\exists x_2[Fx_1 \& Fx_2 \& x_1 \neq x_2 \& \neg\exists y(Fy \& y \neq x_1 \& y \neq x_2)]$
three $\exists x_1\exists x_2\exists x_3[Fx_1 \& Fx_2 \& Fx_3 \& x_1 \neq x_2 \& x_1 \neq x_3 \& x_2 \neq x_3 \& \neg\exists y(Fy \& y \neq x_1 \& y \neq x_2 \& y \neq x_3)]$
n $\exists x_1 \cdots \exists x_n[Fx_1 \& \cdots \& Fx_n \& x_1 \neq x_2 \& \cdots \& x_{n-1} \neq x_n \& \neg\exists y(Fy \& y \neq x_1 \& \cdots \& y \neq x_n)]$

Specifying the size of the UD

Removing F from the symbolizations above produces sentences that talk about the size of the UD. For instance, ‘there are at least 2 things (in the UD)’ may be symbolized as $\exists x\exists y(x \neq y)$.

	YES	NO
Is \mathcal{A} a tautology?	prove $\vdash \mathcal{A}$	give a model in which \mathcal{A} is false
Is \mathcal{A} a contradiction?	prove $\vdash \neg \mathcal{A}$	give a model in which \mathcal{A} is true
Is \mathcal{A} contingent?	give a model in which \mathcal{A} is true and another in which \mathcal{A} is false	prove $\vdash \mathcal{A}$ or $\vdash \neg \mathcal{A}$
Are \mathcal{A} and \mathcal{B} equivalent?	prove $\mathcal{A} \vdash \mathcal{B}$ and $\mathcal{B} \vdash \mathcal{A}$	give a model in which \mathcal{A} and \mathcal{B} have different truth values
Is the set \mathbb{A} consistent?	give a model in which all the sentences in \mathbb{A} are true	taking the sentences in \mathbb{A} , prove \mathcal{B} and $\neg \mathcal{B}$
Is the argument ' $\mathcal{P}, \therefore \mathcal{C}$ ' valid?	prove $\mathcal{P} \vdash \mathcal{C}$	give a model in which \mathcal{P} is true and \mathcal{C} is false

Basic Rules of Proof

REITERATION

$$\begin{array}{l|l} m & \mathcal{A} \\ & \mathcal{A} \quad \text{R } m \end{array}$$

CONJUNCTION INTRODUCTION

$$\begin{array}{l|l} m & \mathcal{A} \\ n & \mathcal{B} \\ & \mathcal{A} \& \mathcal{B} \quad \& \text{I } m, n \end{array}$$

CONJUNCTION ELIMINATION

$$\begin{array}{l|l} m & \mathcal{A} \& \mathcal{B} \\ & \mathcal{A} \quad \& \text{E } m \\ & \mathcal{B} \quad \& \text{E } m \end{array}$$

DISJUNCTION INTRODUCTION

$$\begin{array}{l|l} m & \mathcal{A} \\ & \mathcal{A} \vee \mathcal{B} \quad \vee \text{I } m \\ & \mathcal{B} \vee \mathcal{A} \quad \vee \text{I } m \end{array}$$

DISJUNCTION ELIMINATION

$$\begin{array}{l|l} m & \mathcal{A} \vee \mathcal{B} \\ n & \neg \mathcal{B} \\ & \mathcal{A} \quad \vee \text{E } m, n \end{array}$$

$$\begin{array}{l|l} m & \mathcal{A} \vee \mathcal{B} \\ n & \neg \mathcal{A} \\ & \mathcal{B} \quad \vee \text{E } m, n \end{array}$$

CONDITIONAL INTRODUCTION

$$\begin{array}{l|l} m & \begin{array}{l|l} & \mathcal{A} \quad \text{want } \mathcal{B} \\ & \mathcal{B} \end{array} \\ n & \mathcal{A} \rightarrow \mathcal{B} \quad \rightarrow \text{I } m-n \end{array}$$

CONDITIONAL ELIMINATION

$$\begin{array}{l|l} m & \mathcal{A} \rightarrow \mathcal{B} \\ n & \mathcal{A} \\ & \mathcal{B} \quad \rightarrow \text{E } m, n \end{array}$$

BICONDITIONAL INTRODUCTION

$$\begin{array}{l|l} m & \begin{array}{l|l} & \mathcal{A} \quad \text{want } \mathcal{B} \\ & \mathcal{B} \end{array} \\ n & \begin{array}{l|l} & \mathcal{B} \quad \text{want } \mathcal{A} \\ & \mathcal{A} \end{array} \\ q & \mathcal{A} \leftrightarrow \mathcal{B} \quad \leftrightarrow \text{I } m-n, p-q \end{array}$$

BICONDITIONAL ELIMINATION

$$\begin{array}{l|l} m & \mathcal{A} \leftrightarrow \mathcal{B} \\ n & \mathcal{B} \\ & \mathcal{A} \quad \leftrightarrow \text{E } m, n \end{array}$$

$$\begin{array}{l|l} m & \mathcal{A} \leftrightarrow \mathcal{B} \\ n & \mathcal{A} \\ & \mathcal{B} \quad \leftrightarrow \text{E } m, n \end{array}$$

NEGATION INTRODUCTION

$$\begin{array}{l|l} m & \begin{array}{l|l} & \mathcal{A} \quad \text{for reductio} \\ & \mathcal{B} \\ & \neg \mathcal{B} \end{array} \\ n-1 & \begin{array}{l|l} & \mathcal{B} \\ & \neg \mathcal{B} \end{array} \\ n & \neg \mathcal{A} \quad \neg \text{I } m-n \end{array}$$

NEGATION ELIMINATION

$$\begin{array}{l|l} m & \begin{array}{l|l} & \neg \mathcal{A} \quad \text{for reductio} \\ & \mathcal{B} \\ & \neg \mathcal{B} \end{array} \\ n-1 & \begin{array}{l|l} & \mathcal{B} \\ & \neg \mathcal{B} \end{array} \\ n & \mathcal{A} \quad \neg \text{E } m-n \end{array}$$

Quantifier Rules

EXISTENTIAL INTRODUCTION

$$m \left| \begin{array}{l} \mathcal{A} \\ \hline \exists \chi \mathcal{A}[\chi|c]^* \end{array} \right. \quad \exists I \ m$$

* χ may replace some or all occurrences of c in \mathcal{A} .

EXISTENTIAL ELIMINATION

$$m \left| \begin{array}{l} \exists \chi \mathcal{A} \\ n \left| \begin{array}{l} \mathcal{A}[c|\chi] \\ \hline \mathcal{B} \end{array} \right. \\ p \left| \begin{array}{l} \mathcal{B} \\ \hline \mathcal{B}^\dagger \end{array} \right. \end{array} \right. \quad \exists E \ m, n-p$$

† The constant c must not appear in $\exists \chi \mathcal{A}$, in \mathcal{B} , or in any undischarged assumption.

UNIVERSAL INTRODUCTION

$$m \left| \begin{array}{l} \mathcal{A} \\ \hline \forall \chi \mathcal{A}[\chi|c]^\ddagger \end{array} \right. \quad \forall I \ m$$

‡ c must not occur in any undischarged assumptions.

UNIVERSAL ELIMINATION

$$m \left| \begin{array}{l} \forall \chi \mathcal{A} \\ \hline \mathcal{A}[c|\chi] \end{array} \right. \quad \forall E \ m$$

Identity Rules

$$\left| \begin{array}{l} c = c \end{array} \right. \quad =I$$

$$m \left| \begin{array}{l} c = d \\ n \left| \begin{array}{l} \mathcal{A} \\ \hline \mathcal{A}[c|d] \end{array} \right. \\ \hline \mathcal{A}[d|c] \end{array} \right. \quad =E \ m, n$$

Derived Rules

THE OTHER DISJUNCTION RULE

$$m \left| \begin{array}{l} \mathcal{A} \vee \mathcal{B} \\ n \left| \begin{array}{l} \mathcal{A} \rightarrow \mathcal{C} \\ \hline \mathcal{B} \rightarrow \mathcal{C} \end{array} \right. \\ p \left| \begin{array}{l} \mathcal{C} \end{array} \right. \end{array} \right. \quad \vee^* \ m, n, p$$

MODUS TOLLENS

$$m \left| \begin{array}{l} \mathcal{A} \rightarrow \mathcal{B} \\ n \left| \begin{array}{l} \neg \mathcal{B} \\ \hline \neg \mathcal{A} \end{array} \right. \end{array} \right. \quad \text{MT } m, n$$

HYPOTHETICAL SYLLOGISM

$$m \left| \begin{array}{l} \mathcal{A} \rightarrow \mathcal{B} \\ n \left| \begin{array}{l} \mathcal{B} \rightarrow \mathcal{C} \\ \hline \mathcal{A} \rightarrow \mathcal{C} \end{array} \right. \end{array} \right. \quad \text{HS } m, n$$

Replacement Rules

COMMUTIVITY (Comm)
 $(\mathcal{A} \& \mathcal{B}) \iff (\mathcal{B} \& \mathcal{A})$
 $(\mathcal{A} \vee \mathcal{B}) \iff (\mathcal{B} \vee \mathcal{A})$

DEMORGAN (DeM)
 $\neg(\mathcal{A} \vee \mathcal{B}) \iff (\neg \mathcal{A} \& \neg \mathcal{B})$
 $\neg(\mathcal{A} \& \mathcal{B}) \iff (\neg \mathcal{A} \vee \neg \mathcal{B})$

DOUBLE NEGATION (DN)
 $\neg \neg \mathcal{A} \iff \mathcal{A}$

IMPLICATION ($\vee \rightarrow$)
 $(\mathcal{A} \rightarrow \mathcal{B}) \iff (\neg \mathcal{A} \vee \mathcal{B})$

QUANTIFIER NEGATION (QN)
 $\neg \forall \chi \mathcal{A} \iff \exists \chi \neg \mathcal{A}$
 $\neg \exists \chi \mathcal{A} \iff \forall \chi \neg \mathcal{A}$

In the Introduction to his volume *Symbolic Logic*, Charles Lutwidge Dodson advised: “When you come to any passage you don’t understand, *read it again*: if you *still* don’t understand it, *read it again*: if you fail, even after *three* readings, very likely your brain is getting a little tired. In that case, put the book away, and take to other occupations, and next day, when you come to it fresh, you will very likely find that it is *quite* easy.”

The same might be said for this volume, although the reader is forgiven if they take a break for snacks after *two* readings.

about the author:

P.D. Magnus is an assistant professor of philosophy in Albany, New York. His primary research is in the philosophy of science, concerned especially with the underdetermination of theory by data.